

Automated Temporal Equilibrium Analysis: Verification and Synthesis of Multi-Player Games

Julian Gutierrez^a, Muhammad Najib^b, Giuseppe Perelli^c, Michael Wooldridge^d

^a*Faculty of Information Technology, Monash University*

^b*Department of Computer Science, University of Kaiserslautern*

^c*Department of Computer, Automatic, and Management Engineering, Sapienza University of Rome*

^d*Department of Computer Science, University of Oxford*

Abstract

In the context of multi-agent systems, the *rational verification* problem is concerned with checking which temporal logic properties will hold in a system when its constituent agents are assumed to behave rationally and strategically in pursuit of individual objectives. Typically, those objectives are expressed as temporal logic formulae which the relevant agent desires to see satisfied. Unfortunately, rational verification is computationally complex, and requires specialised techniques in order to obtain practically useable implementations. In this paper, we present such a technique. This technique relies on a reduction of the rational verification problem to the solution of a collection of parity games. Our approach has been implemented in the *Equilibrium Verification Environment (EVE)* system. The EVE system takes as input a model of a concurrent/multi-agent system represented using the Simple Reactive Modules Language (SRML), where agent goals are represented as Linear Temporal Logic (LTL) formulae, together with a claim about the equilibrium behaviour of the system, also expressed as an LTL formula. EVE can then check whether the LTL claim holds on some (or every) computation of the system that could arise through agents choosing Nash equilibrium strategies; it can also check whether a system has a Nash equilibrium, and synthesise individual strategies for players in the multi-player game. After

*Corresponding author: Julian Gutierrez.

Email addresses: julian.gutierrez@monash.edu (Julian Gutierrez), najib@cs.uni-kl.de (Muhammad Najib), perelli@diag.uniroma1.it (Giuseppe Perelli), michael.wooldridge@cs.ox.ac.uk (Michael Wooldridge)

presenting our basic framework, we describe our new technique and prove its correctness. We then describe our implementation in the EVE system, and present experimental results which show that EVE performs favourably in comparison to other existing tools that support rational verification.

Keywords: Multi-agent systems, Temporal logic, Nash equilibrium, Bisimulation invariance, Rational verification, Model checking, Synthesis.

1. Introduction

The deployment of AI technologies in a wide range of application areas over the past decade has brought the problem of *verifying* such systems into sharp focus. Verification is the problem of ensuring that a particular system is correct with respect to some specification. The most successful approach to automated formal verification is that of *model checking* [1]. With this approach, we first derive a finite state abstract model of the system \mathcal{S} being studied; a common approach involves representing the system as a directed graph in which vertices correspond to states of the system, and edges correspond to the execution of program instructions, or the performance of actions; branching in the graph represents either input from the environment, or choices available to components of the system. With this approach, the directed graph is typically referred to as a labelled transition system, or Kripke structure: each path through the transition system represents a possible execution or computation of the system \mathcal{S} . Correctness properties of interest are expressed as formulae φ of propositional temporal logic—the most popular such logics for this purpose are Linear Temporal Logic (LTL) and the Computation Tree Logic (CTL). In the case of properties φ expressed as LTL formulae, we typically want to check whether φ is satisfied on some or all possible computations of \mathcal{S} , that is, on some or all possible paths through the transition system/Kripke structure representing \mathcal{S} .

Great advances have been made in model checking since the approach was first proposed in the early 1980s, and the technique is now widely used in industry. Nevertheless, the verification of practical software systems is by no means a solved problem, and remains the subject of intense ongoing research. The verification of AI systems,

24 however, raises a distinctive new set of challenges. The present paper is concerned
25 with the problem of verifying *multi-agent systems*, which are AI systems consisting of
26 multiple interacting semi-autonomous software components known as *agents* [2, 3].
27 Software agents were originally proposed in the late 1980s, but it is only over the past
28 decade that the software agent paradigm has been widely adopted. At the time of
29 writing, software agents are ubiquitous: we have software agents in our phone (*e.g.*,
30 Siri), processing requests online, automatically trading in global markets, controlling
31 complex navigation systems (*e.g.*, those in self-driving cars), and even carrying out
32 tasks on our behalf in our homes (*e.g.*, Alexa). Typically, these agents do not work in
33 isolation: they may interact with humans or with other software agents. The field of
34 multi-agent systems is concerned with understanding and engineering systems that
35 have these characteristics.

36 We typically assume that agents are acting in pursuit of goals or preferences that
37 are delegated to them by their users. However, whether an agent is able to achieve
38 its goal, or the extent to which it can bring about its preferences, will be directly
39 influenced by the behaviour of other agents. Thus, to act optimally, an agent must
40 reason *strategically*, taking into account the goals/preferences of other agents, and the
41 fact that they too will be acting strategically in the pursuit of these, taking into account
42 the goals/preferences of other agents and their own strategic behaviour. *Game theory*
43 is the mathematical theory of strategic interaction, and as such, it provides a natural
44 set of tools for reasoning about multi-agent systems [4].

45 With respect to the problem of verifying multi-agent systems, the relevance of
46 game theory is as follows. Suppose we are interested in whether a multi-agent system
47 \mathcal{S} , populated by self-interested agents, might exhibit some property represented by an
48 LTL formula φ . We can, of course, directly apply standard model checking techniques,
49 to determine whether φ holds on some or all computations of \mathcal{S} . However, given that
50 our agents are assumed to act rationally, whether φ holds on some or all computations
51 is not relevant if the computations in question involve irrational choices on behalf of
52 some agents in the system. A much more relevant question, therefore, is whether φ
53 holds on some or all computations *that could result from agents in the system making*
54 *rational choices*. This raises the question of what counts as a rational choice by the

55 agents in the system, and for this game theory provides a number of answers, in the
56 form of *solution concepts* such as Nash equilibrium [4, 3]. Thus, from the point of view
57 of game theory, *correct behaviour* would correspond to *rational behaviour* according
58 to some game theoretic solution concept, which is another way of saying that agents
59 in the system will act *optimally* with respect to their preferences/goals, under the
60 assumption that other agents do the same.

61 This approach to reasoning about the behaviour of multi-agent AI systems es-
62 tablishes a natural connection between multi-agent systems and multi-player games:
63 agents correspond to players, computations of the multi-agent system correspond to
64 plays of the game, individual agent behaviours correspond to player strategies (which
65 define how players make choices in the system over time), and correct behaviour
66 would correspond to rational behaviour—in our case, player behaviour that is con-
67 sistent with the set of Nash equilibria of the multi-player game, whenever such a
68 set is non-empty. Our main interest in this paper is the development of the theory,
69 algorithms, and tools for the automated game theoretic analysis of concurrent and
70 multi-agent systems, and in particular, the analysis of temporal logic properties that
71 will hold in a multi-agent system under the assumption that players choose strategies
72 which form a Nash equilibrium¹.

73 The connection between AI systems (modelled as multi-agent systems) and multi-
74 player games is well-established, but one may still wonder why *correct behaviour* for
75 the AI system should correspond to *rational behaviour* in the multi-player game. This
76 is a legitimate question, especially, because game theory offers very many different
77 notions of rationality, and therefore of optimal behaviour in the system/game. For
78 instance, solution concepts such as subgame-perfect Nash equilibrium (SPNE) and
79 strong Nash equilibrium (SNE) are refinements of Nash equilibrium where the notion
80 of rationality needs to satisfy stronger requirements. Consequently, there may be
81 executions of a multi-agent system that would correspond to a Nash equilibrium of
82 the associated multi-player game (thus, regarded as correct behaviours of the multi-

¹Although in this work we focus on Nash equilibrium, a similar methodology may be applied using refinements of Nash equilibrium and other solution concepts.

83 agent system), but which do not correspond to a subgame-perfect Nash equilibrium
 84 or to a strong Nash equilibrium of the associated multi-player game. We do not argue
 85 that Nash equilibrium is the only solution concept of relevance in the game theoretic
 86 analysis of multi-agent systems, but we believe (as do many others [3, 5, 6]) that Nash
 87 equilibrium is a natural and appropriate starting point for such an analysis. Taking
 88 Nash equilibrium as our baseline notion of rationality in multi-player games, and
 89 therefore of correctness in multi-agent systems, we focus our study on two problems
 90 related to the temporal equilibrium analysis of multi-agent systems [7, 8], as we now
 91 explain.

92 *Synthesis and Rational Verification.* The two main problems of interest to us are the
 93 *rational verification* and *automated synthesis* problems for concurrent and multi-agent
 94 systems modelled as multi-player games. In the *rational verification* problem, we de-
 95 sire to check which temporal logic properties are satisfied by the system/game *in*
 96 *equilibrium*, that is, temporal logic properties satisfied by executions of the multi-
 97 agent system generated by strategies that form a Nash equilibrium. A little more for-
 98 mally, let P_1, \dots, P_n be the agents in our concurrent and multi-agent system, and let
 99 $\text{NE}(P_1, \dots, P_n)$ denote the set of all executions, hereafter called runs, of the system
 100 that could be generated by agents selecting strategies that form a Nash equilibrium.
 101 Finally, let φ be an LTL formula. Then, in the rational verification problem, we want
 102 to know whether for some/every run $\pi \in \text{NE}(P_1, \dots, P_n)$ we have $\pi \models \varphi$.

103 In the automated *synthesis* problem, on the other hand, we additionally desire to
 104 *construct* a profile of strategies for players so that the resulting profile is an equilib-
 105 rium of the multi-player game, and induces a run that satisfies a given property of
 106 interest, again expressed as a temporal logic formula. That is, we are given the system
 107 P_1, \dots, P_n , and a temporal logic property φ , and we are asked to compute Nash equi-
 108 librium strategies $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$, one for each player in the game, that would result
 109 in φ being satisfied in the run $\pi(\vec{\sigma})$ that would be generated when these strategies are
 110 enacted.

111 *Our Approach.* In this paper, we present a new approach to the rational verifica-
112 tion and automated synthesis problems for concurrent and multi-agent systems. In
113 particular, we develop a novel technique that can be used for both rational verifica-
114 tion and automated synthesis using a reduction to the solution of a collection of
115 *parity games*. The technique can be efficiently implemented making use of power-
116 ful techniques for parity games and temporal logic synthesis and verification, and has
117 been deployed in the *Equilibrium Verification Environment* (EVE [9]), which supports
118 high-level descriptions of systems/games using the *Simple Reactive Modules Language*
119 (SRML [10, 7]) and temporal logic specifications given by Linear Temporal Logic for-
120 mulae [11].

121 The central decision problem that we consider is that of NON-EMPTYNESS, the prob-
122 lem of checking if the set of Nash equilibria in a multi-player game is empty; as we will
123 later show, rational verification and synthesis can be reduced to this problem. If we
124 consider concurrent and multi-player games in which players have goals expressed
125 as temporal logic formulae, this problem is known to be 2EXPTIME-complete for a
126 wide range of system representations and temporal logic languages. For instance, for
127 games with perfect information played on labelled graphs, the problem is 2EXPTIME-
128 complete when goals are given as LTL formulae [12], and 2EXPTIME-hard when goals
129 are given in CTL [13]. The problem is 2EXPTIME-complete even if succinct represen-
130 tations [14, 15] or only two-player games [16] are considered, and becomes undecid-
131 able if imperfect information and more than two players are allowed [17], showing
132 the very high complexity of solving this problem, from both practical and theoretical
133 viewpoints.

134 A common feature of the results above mentioned is that—modulo minor variations—
135 their solutions are, in the end, reduced to the construction of an alternating parity
136 automaton over *infinite trees* (APT [18]) which are then checked for non-emptiness.
137 Here, we present a novel, simpler, and more direct technique for checking the ex-
138 istence of Nash equilibria in games where players have goals expressed in LTL. In
139 particular, our technique does not rely on the solution of an APT. Instead, we reduce
140 the problem to the solution of (a collection of) parity games [19], which are widely
141 used for synthesis and verification problems.

142 Formally, a parity game is a two-player zero-sum turn-based game given by a
143 labelled finite graph $H = (V_0, V_1, E, \alpha)$ such that $V = V_0 \cup V_1$ is a set of states
144 partitioned into Player 0 (V_0) and Player 1 (V_1) states, respectively, $E \subseteq V \times V$ is
145 a set of edges/transitions, and $\alpha : V \rightarrow \mathbb{N}$ is a labelling priority function. Player 0
146 wins if the smallest priority that occurs infinitely often in the infinite play is even.
147 Otherwise, player 1 wins. It is known that solving a parity game (checking which
148 player has a winning strategy) is in $\text{NP} \cap \text{coNP}$ [20], and can be solved in quasi-
149 polynomial time [21]².

150 Our technique uses parity games in the following way. We take as input a game G
151 (representing a concurrent and multi-agent system) and build a parity game H whose
152 sets of states and transitions are doubly exponential in the size of the input but with
153 priority function only exponential in the size of the input game. Using a determin-
154 istic Streett automaton on *infinite words* (DSW [22]), we then solve the parity game,
155 leading to a decision procedure that is, overall, in 2EXPTIME, and, therefore, given
156 the hardness results we mentioned above, essentially optimal.

157 *Context.* Games have several dimensions: for example, they may be cooperative or
158 non-cooperative; have perfect or imperfect information; have perfect or imperfect
159 recall; be stochastic or not; amongst many other features. Each of these aspects will
160 have a modelling and computational impact on the work to be developed, and so it is
161 important to be precise about the nature of the games we are studying, and therefore
162 the assumptions underpinning our approach.

163 Our framework considers non-cooperative multi-player general-sum games with
164 perfect information, with Nash equilibrium as the main game-theoretic solution con-
165 cept. The games are played on finite structures (state-transition structures induced
166 by high-level SRML descriptions), with players having goals (preferences over plays)
167 given by LTL formulae and deterministic strategies represented by finite-state ma-
168 chines with output (Moore machines, sometimes referred to as transducers). Because

²Despite more than 30 years of research, and promising practical performance for algorithms to solve them, it remains unknown whether parity games can be solved in polynomial time.

169 of the features of our framework – chiefly, the fact that players have LTL goals and
170 games are played on finite structures – considering deterministic strategies modelled
171 as finite-state machines does not represent a restriction: in our framework, anything
172 that a player can achieve with a perfect-recall strategy can also be achieved with a
173 finite-state machine strategy (see, *e.g.*, [15] for the formal results).

174 Finally, we note that our games have equilibria that are *bisimulation invariant*: that
175 is, bisimilar structures have the same set of Nash equilibria. This is a highly desirable
176 property, and to the best of our knowledge, in this respect our work is unique in the
177 computer science and multi-agent systems literatures.

178 *The EVE System.* The technique outlined above and described in detail in this pa-
179 per has been successfully implemented in the *Equilibrium Verification Environment*
180 (EVE) system [23]. EVE takes as input a model of a concurrent and multi-agent
181 system, in which agents are specified using the Simple Reactive Modules Language
182 (SRML) [10, 7], and preferences for agents are defined by associating with each agent
183 a goal, represented as a formula of LTL [11]. Note that we believe our choice of the
184 Reactive Modules language is a very natural one [24]: The language is both widely
185 used in practical model checking systems, such as MOCHA [25] and PRISM [26], and
186 close to real-world (declarative) programming models and specification languages.

187 Now, given a specification of a multi-agent system and player preferences, the
188 EVE system can: (i) check for the existence of a Nash equilibrium in a multi-player
189 game; (ii) check whether a given LTL formula is satisfied on some or every Nash
190 equilibrium of the system; and (iii) synthesise individual player strategies in the game.
191 As we will show in the paper, EVE performs favourably compared with other existing
192 tools that support rational verification. Moreover, EVE is the first and only tool for
193 automated temporal equilibrium analysis for a model of multi-player games where
194 Nash equilibria are preserved under bisimilarity³.

195 Note that our approach may be used to model a wide range of multi-agent systems.

³Other tools to compute Nash equilibria exist, but they do not use our model of strategies. A comparison with those other techniques for equilibrium analysis are discussed later.

196 For example, as shown in [7], it is easy to capture multi-agent STRIPS systems [27].

197 *Structure of the paper.* The remainder of this article is structured as follows.

- 198 • Section 2 presents the relevant background on games, logic, and automata.
- 199 • In Section 3, we formalise the main problem of interest and give a high-level
200 description of the core decision procedure for temporal equilibrium analysis
201 developed in this paper.
- 202 • In Sections 4, 5, and 6, we describe in detail our main decision procedure for
203 temporal equilibrium analysis, prove its correctness, and show that it is essen-
204 tially optimal with respect to computational complexity.
- 205 • In Section 7, we show how to use our main decision procedure to do rational
206 verification and automated synthesis of logic-based multi-player games.
- 207 • In Section 8, we describe the EVE system, and give detailed experimental results
208 which demonstrate that EVE performs favourably in comparison with other
209 tools that support rational verification.
- 210 • In Section 9, we conclude, discuss relevant related work, and propose some
211 avenues for future work.

212 The source code for EVE is available online⁴, and the system can also be accessed via
213 the web⁵.

214 2. Preliminaries

Games. A *concurrent (multi-player) game structure* (CGS) is a tuple

$$\mathcal{M} = (\mathbb{N}, (Ac_i)_{i \in \mathbb{N}}, St, s_0, tr)$$

215 where $\mathbb{N} = \{1, \dots, n\}$ is a set of *players*, each Ac_i is a set of *actions*, St is a set
216 of *states*, with a designated *initial* state s_0 . With each player $i \in \mathbb{N}$ and each state

⁴See <https://github.com/eve-mas/eve-parity>

⁵See <http://eve.cs.ox.ac.uk/>

217 $s \in \text{St}$, we associate a non-empty set $\text{Ac}_i(s)$ of *available* actions that, intuitively, i
 218 can perform when in state s . We refer to a profile of actions $\vec{a} = (a_1, \dots, a_n) \in \vec{\text{Ac}} =$
 219 $\text{Ac}_1 \times \dots \times \text{Ac}_n$ as a *direction*. A direction \vec{a} is available in state s if for all i we have
 220 $a_i \in \text{Ac}_i(s)$. Write $\vec{\text{Ac}}(s)$ for the set of available directions in state s . For a given set
 221 of players $A \subseteq \mathbb{N}$ and an action profile \vec{a} , we let \vec{a}_A and \vec{a}_{-A} be two tuples of actions,
 222 respectively, one for each player in A and one for each player in $\mathbb{N} \setminus A$. We also write
 223 \vec{a}_i for $\vec{a}_{\{i\}}$ and \vec{a}_{-i} for $\vec{a}_{\mathbb{N} \setminus \{i\}}$. Furthermore, for two directions \vec{a} and \vec{a}' , we write
 224 $(\vec{a}_A, \vec{a}'_{-A})$ to denote the direction where the actions for players in A are taken from
 225 \vec{a} and the actions for players in $\mathbb{N} \setminus A$ are taken from \vec{a}' . Finally, tr is a *deterministic*
 226 *transition function*, which associate each state s and every available direction \vec{a} in s a
 227 state $s' \in \text{St}$.

228 Whenever there is \vec{a} such that $\text{tr}(s, \vec{a}) = s'$, we say that s' is *accessible* from s . A
 229 *path* $\pi = s_0, s_1, \dots \in \text{St}^\omega$ is an infinite sequence of states such that, for every $k \in \mathbb{N}$,
 230 s_{k+1} is accessible from s_k . By π_k we refer to the $(k+1)$ -th state in π and by $\pi_{\leq k}$ to
 231 the (finite) prefix of π up to the $(k+1)$ -th element. An *action profile run* is an infinite
 232 sequence $\eta = \vec{a}_0, \vec{a}_1, \dots$ of action profiles. Note that, since \mathcal{M} is deterministic (*i.e.*,
 233 the transition function tr is deterministic), for a given state s_0 , an action profile run
 234 uniquely determines the path π in which, for every $k \in \mathbb{N}$, $\pi_{k+1} = \text{tr}(\pi_k, \vec{a}_k)$.

235 A CGS is a type of concurrent system. As such, behaviourally equivalent CGSs
 236 should give rise to strategically equivalent games. However, that is not always the
 237 case. A comprehensive study of this issue can be found in [28, 29] where the strate-
 238 gic power of games is compared using one of the most important behavioural (also
 239 called observational) equivalences in concurrency, namely bisimilarity, which is usu-
 240 ally defined over Kripke structures or labelled transition systems (see, *e.g.*, [30, 31]).
 241 However, the equivalence can be uniformly defined for general CGSs, where direc-
 242 tions play the role of, for instance, actions in transition systems. Formally, let $M =$
 243 $(\mathbb{N}, (\text{Ac}_i)_{i \in \mathbb{N}}, \text{St}, s_0, \text{tr})$ and $M' = (\mathbb{N}, (\text{Ac}'_i)_{i \in \mathbb{N}}, \text{St}', s'_0, \text{tr}')$ be two CGSs, and $\lambda :$
 244 $\text{St} \rightarrow \text{AP}$ and $\lambda' : \text{St}' \rightarrow \text{AP}$ be two labelling functions over a set of propositional
 245 variables AP . A *bisimulation*, denoted by \sim , between states $s^* \in \text{St}$ and $t^* \in \text{St}'$ is
 246 a non-empty binary relation $R \subseteq \text{St} \times \text{St}'$, such that $s^* R t^*$ and for all $s, s' \in \text{St}$,
 247 $t, t' \in \text{St}'$, and $\vec{a} \in \vec{\text{Ac}}$:

- 248 • $s R t$ implies $\lambda(s) = \lambda'(t)$,
- 249 • $s R t$ and $\text{tr}(s, \vec{a}) = s'$ implies $\text{tr}(t, \vec{a}) = t''$ for some $t'' \in \text{St}'$ with $s' R t''$,
- 250 • $s R t$ and $\text{tr}(t, \vec{a}) = t'$ implies $\text{tr}(s, \vec{a}) = s''$ for some $s'' \in \text{St}$ with $s'' R t'$.

251 Then, if there is a bisimulation between two states s^* and t^* , we say that they are
 252 *bisimilar* and write $s^* \sim t^*$ in such a case. We also say that CGSs M and M' are
 253 *bisimilar* (in symbols $M \sim M'$) if $s_0 \sim s'_0$. Bisimilar structures satisfy the same set
 254 of temporal logic properties, a desirable property that will be relevant later.

255 A CGS defines the dynamic structure of a game, but lacks a central aspect of games
 256 in the sense of game theory: preferences, which give games their strategic structure.
 257 A *multi-player game* is obtained from a structure \mathcal{M} by associating each player with a
 258 goal. In this paper, we consider multi-player games with parity and Linear Temporal
 259 Logic (LTL) goals.

LTL [11] extends classical propositional logic with two operators, **X** (“next”) and
U (“until”), that can be used to express properties of paths. The syntax of LTL is
 defined with respect to a set AP of propositional variables as follows:

$$\varphi ::= \top \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U} \varphi$$

260 where $p \in \text{AP}$. The remaining classical logical connectives are defined in terms of
 261 \neg and \vee in the usual way. Two key derived LTL operators are **F** (“eventually”) and
 262 **G** (“always”), which are defined in terms of **U** as follows: $\mathbf{F}\varphi = \top \mathbf{U} \varphi$ and $\mathbf{G}\varphi =$
 263 $\neg \mathbf{F} \neg \varphi$.

We interpret formulae of LTL with respect to tuples (π, t, λ) , where π is a path
 over some multi-player game, $t \in \mathbb{N}$ is a temporal index into π , and $\lambda : \text{St} \rightarrow 2^{\text{AP}}$
 is a labelling function, that indicates which propositional variables are true in every

state. Formally, the semantics of LTL is given by the following rules:

$$\begin{aligned}
(\pi, t, \lambda) &\models \top \\
(\pi, t, \lambda) &\models p \quad \text{iff} \quad p \in \lambda(\pi_t) \\
(\pi, t, \lambda) &\models \neg\varphi \quad \text{iff} \quad \text{it is not the case that } (\pi, t, \lambda) \models \varphi \\
(\pi, t, \lambda) &\models \varphi \vee \psi \quad \text{iff} \quad (\pi, t, \lambda) \models \varphi \text{ or } (\pi, t, \lambda) \models \psi \\
(\pi, t, \lambda) &\models \mathbf{X}\varphi \quad \text{iff} \quad (\pi, t+1, \lambda) \models \varphi \\
(\pi, t, \lambda) &\models \varphi \mathbf{U} \psi \quad \text{iff} \quad \text{for some } t' \geq t : ((\pi, t', \lambda) \models \psi \text{ and} \\
&\quad \text{for all } t \leq t'' < t' : (\pi, t'', \lambda) \models \varphi).
\end{aligned}$$

264 If $(\pi, 0, \lambda) \models \varphi$, we write $\pi \models \varphi$ and say that π *satisfies* φ .

Definition 1. A (*concurrent multi-player*) LTL *game* is a tuple

$$\mathcal{G}_{\text{LTL}} = (\mathcal{M}, \lambda, (\gamma_i)_{i \in \mathbb{N}})$$

265 where $\lambda : \text{St} \rightarrow 2^{\text{AP}}$ is a labelling function on the set of states St of \mathcal{M} , and each γ_i
266 is the goal of player i , given as an LTL formula over AP.

267 To define multi-player games with parity goals we consider priority functions.
268 Let $\alpha : \text{St} \rightarrow \mathbb{N}$ be a priority function. A path π satisfies $\alpha : \text{St} \rightarrow \mathbb{N}$, and write
269 $\pi \models \alpha$ in that case, if the minimum number occurring infinitely often in the infinite
270 sequence $\alpha(\pi_0), \alpha(\pi_1), \alpha(\pi_2), \dots$ is even.

271 Observe that parity conditions are *prefix-independent*, that is, for every path π and
272 a finite sequence $h \in \text{St}^*$, it holds that $h \cdot \pi \models \alpha$ if and only if $\pi \models \alpha$.

Definition 2. A (*concurrent multi-player*) Parity *game* is a tuple

$$\mathcal{G}_{\text{PAR}} = (\mathcal{M}, (\alpha_i)_{i \in \mathbb{N}})$$

273 where $\alpha_i : \text{St} \rightarrow \mathbb{N}$ is the goal of player i , given as a priority function over St .

274 Hereafter, for statements regarding either LTL or Parity games⁶, we will simply
275 denote the underlying structure as \mathcal{G} . Games are played by each player i selecting

⁶To simplify notations, note that, hereafter, by “Parity game” we denote the concurrent and multi-player extension defined here of the well-known two-player turn-based parity games in the literature.

276 a *strategy* σ_i that will define how to make choices over time. Formally, for a given
277 game \mathcal{G} , a strategy $\sigma_i = (S_i, s_i^0, \delta_i, \tau_i)$ for player i is a finite state machine with
278 output (a transducer), where S_i is a finite and non-empty set of *internal states*, s_i^0
279 is the *initial state*, $\delta_i : S_i \times \vec{Ac} \rightarrow S_i$ is a deterministic *internal transition function*,
280 and $\tau_i : S_i \rightarrow Ac_i$ an *action function*. Note that strategies are required to output
281 actions that are available to the agent in the current state. To enforce this, we assume
282 that the current state $s \in \text{St}$ in the arena is encoded in the internal state s_i in S_i
283 of agent i and that the action $\tau_i(s_i)$ taken by the action function belongs to $Ac_i(s)$.
284 Let Σ_i be the set of strategies for player i . A strategy is *memoryless* in \mathcal{G} from s if
285 $S_i = \text{St}$, $s_i^0 = s$, and $\delta_i = \text{tr}$. Once every player i has selected a strategy σ_i , a *strategy*
286 *profile* $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$ results and the game has an *outcome*, a path in \mathcal{M} , which
287 we will denote by $\pi(\vec{\sigma})$. Because strategies are deterministic, $\pi(\vec{\sigma})$ is the unique path
288 induced by $\vec{\sigma}$, that is, the infinite sequence s_0, s_1, s_2, \dots such that

- 289 • $s_{k+1} = \text{tr}(s_k, (\tau_1(s_1^k), \dots, \tau_n(s_n^k)))$, and
- 290 • $s_i^{k+1} = \delta_i(s_i^k, (\tau_1(s_1^k), \dots, \tau_n(s_n^k)))$, for all $k \geq 0$.

291 Note that the path induced by the strategy profile $\vec{\sigma}(\sigma_1, \dots, \sigma_n)$ from state s_0
292 corresponds to the one generated by the finite transducer $T_{\vec{\sigma}}$ obtained from the com-
293 position of the strategies σ_i 's in $\vec{\sigma}$, with input set St and output set \vec{Ac} , where the
294 initial input is s_0 . Since such transducer is finite, the generated path π is *ultimately*
295 *periodic*, that is, there exists $p, r \in \mathbb{N}$ such that $\pi_k = \pi_{k+r}$ for every $p \leq k$. This means
296 that, after the prefix $\pi_{\leq p}$, the path loops indefinitely over the sequence $\pi_{p+1} \dots \pi_{p+r}$.

Nash equilibrium. Since the outcome of a game determines if a player goal is sat-
isfied, we can define a preference relation \succeq_i over outcomes for each player i . Let w_i
be γ_i if \mathcal{G} is an LTL game, and be α_i if \mathcal{G} is a Parity game. Then, for two strategy
profiles $\vec{\sigma}$ and $\vec{\sigma}'$ in \mathcal{G} , we have

$$\pi(\vec{\sigma}) \succeq_i \pi(\vec{\sigma}') \text{ if and only if } \pi(\vec{\sigma}') \models w_i \text{ implies } \pi(\vec{\sigma}) \models w_i.$$

On this basis, we can define the concept of Nash equilibrium [4] for a multi-player
game with LTL or parity goals: given a game \mathcal{G} , a strategy profile $\vec{\sigma}$ is a *Nash equilib-*

rium of \mathcal{G} if, for every player i and strategy $\sigma'_i \in \Sigma_i$, we have

$$\pi(\vec{\sigma}) \succeq_i \pi((\vec{\sigma}_{-i}, \sigma'_i))$$

297 where $(\vec{\sigma}_{-i}, \sigma'_i)$ denotes $(\sigma_1, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots, \sigma_n)$, the strategy profile where
 298 the strategy of player i in $\vec{\sigma}$ is replaced by σ'_i . Let $\text{NE}(\mathcal{G})$ denote the set of Nash
 299 equilibria of \mathcal{G} . In [28, 29] we showed that, using the model of strategies defined
 300 above, the existence of Nash equilibria is preserved across bisimilar systems. This is
 301 in contrast to other models of strategies considered in the concurrent games literature,
 302 which do not preserve Nash equilibria. Because of this, hereafter, we say that $\{\Sigma_i\}_{i \in \mathbb{N}}$
 303 is a set of *bisimulation-invariant strategies* and that $\text{NE}(\mathcal{G})$ is the set of bisimulation-
 304 invariant Nash equilibrium profiles of \mathcal{G} .

Automata. A *deterministic automaton on infinite words* is a tuple

$$\mathcal{A} = (\text{AP}, Q, q^0, \rho, \mathcal{F})$$

where Q is a finite set of states, $\rho : Q \times \text{AP} \rightarrow Q$ is a transition function, q^0 is an initial state, and \mathcal{F} is an acceptance condition. We mainly use *parity* and *Streett* acceptance conditions. A parity condition \mathcal{F} is a partition $\{F_1, \dots, F_n\}$ of Q , where n is the *index* of the parity condition and any $[1, n] \ni k$ is a *priority*. We use a *priority function* $\alpha : Q \rightarrow \mathbb{N}$ that maps states to priorities such that $\alpha(q) = k$ if and only if $q \in F_k$. For a run $\pi = q^0, q^1, q^2 \dots$, let $\text{inf}(\pi)$ denote the set of states occurring infinitely often in the run:

$$\text{inf}(\pi) = \{q \in Q \mid q = q^i \text{ for infinitely many } i\}$$

A run π is accepted by a deterministic parity word (DPW) automaton with condition \mathcal{F} if the minimum priority that occurs infinitely often is even, *i.e.*, if the following condition is satisfied:

$$\left(\min_{k \in [1, n]} (\text{inf}(\pi) \cap F_k \neq \emptyset) \right) \bmod 2 = 0.$$

305 A Streett condition \mathcal{F} is a set of pairs $\{(E_1, C_1), \dots, (E_n, C_n)\}$ where $E_k \subseteq Q$ and
 306 $C_k \subseteq Q$ for all $k \in [1, n]$. A run π is accepted by a deterministic Streett word (DSW)
 307 automaton \mathcal{S} with condition \mathcal{F} if π either visits E_k finitely many times or visits C_k
 308 infinitely often, *i.e.*, if for every k either $\text{inf}(\pi) \cap E_k = \emptyset$ or $\text{inf}(\pi) \cap C_k \neq \emptyset$.




Figure 1: Example of a 4×4 grid world.

309 **Example.** In order to illustrate the usage of our framework, consider the following
 310 example. Suppose we have two robots/agents inhabiting a grid world (an abstraction
 311 of some environment, *e.g.*, a warehouse) with dimensions $n \times n$. Initially, the agents
 312 are located at some corners of the grid; The agents are each able to move around the
 313 grid in directions *north*, *south*, *east*, and *west*. The goal of each agent is to reach the
 314 opposite corner. For instance, if agent i 's initial position is $(0, 0)$, then the goal is to
 315 reach position $(n - 1, n - 1)$. A number of obstacles may also appear on the grid. The
 316 agents are not allowed to move into a coordinate occupied by an obstacle or outside
 317 the grid world. To make it clearer, consider the configuration shown in Figure 1; a
 318 (grey) filled square depicts an obstacle. Agent 1, depicted by ■, can only move west
 319 to $(2, 3)$, whereas agent 2, depicted by ○, can only move east to $(1, 0)$.

320 In this example we make the following assumptions: (1) at each timestep, each
 321 agent has to make a move, that is, it cannot stay at the same position for two consec-
 322 utive timesteps, and it can only move at most one step; (2) the goal of each agent is, as
 323 stated previously, to eventually reach the opposite corner of her initial position. From
 324 system design point of view, the question that may be asked is: can we synthesise a
 325 strategy profile such that it induces a stable (Nash equilibrium) run and at the same
 326 time ensures that the agents never crash into each other?

327 Checking the existence of such strategy profile is not trivial. For instance, the
 328 configuration in Figure 1 does not admit any safe Nash equilibrium runs, that is, where
 329 all agents get their goals achieved without crashing into each other. Player ○ can




Figure 2: A 4×4 grid world with safe Nash equilibrium.

330 reach $(3, 3)$ without crashing into \blacksquare , since \blacksquare can safely “wait” by moving back and
 331 forth between $(0, 3)$ and $(1, 3)$ until \circ reaches $(3, 3)$. However, there is no similar
 332 safe “waiting zone” for \circ to get out of \blacksquare ’s way. On the other hand, the configuration
 333 in Figure 2, admits safe Nash equilibrium; \circ and \blacksquare have safe waiting zones $(0, 0)$
 334 and $(1, 0)$, and $(0, 3)$ and $(1, 3)$, respectively. Clearly, such a reasoning is not always
 335 straightforward, especially when the setting is more complex, and therefore, having
 336 a tool to verify and synthesise such scenario is desirable. Later in Section 8.5 we will
 337 discuss how to encode and check such systems using our tool.

338 3. A Decision Procedure using Parity Games

339 We are now in a position to formally state the NON-EMPTYNESS problem:

340 *Given:* An LTL Game \mathcal{G}_{LTL} .

341 *Question:* Is it the case that $\text{NE}(\mathcal{G}_{\text{LTL}}) \neq \emptyset$?

342 As indicated before, we solve both verification and synthesis through a reduction
 343 to the above problem. The technique we develop consists of three steps. First, we
 344 build a Parity game \mathcal{G}_{PAR} from an input LTL game \mathcal{G}_{LTL} . Then—using a characteri-
 345 sation of Nash equilibrium (presented later) that separates players in the game into
 346 those that achieve their goals in a Nash equilibrium (the “winners”, W) and those that
 347 do not achieve their goals (the “losers”, L)—for each set of players in the game, we
 348 eliminate nodes and paths in \mathcal{G}_{PAR} which cannot be a part of a Nash equilibrium, thus

349 producing a modified Parity game, $\mathcal{G}_{\text{PAR}}^{-L}$. Finally, in the third step, we use Streett au-
 350 tomata on infinite words to check if the obtained Parity game witnesses the existence
 351 of a Nash equilibrium. The overall algorithm is presented in Algorithm 1 which also
 352 includes some comments pointing to the relevant Sections/Theorems. The first step
 353 is contained in line 3, while the third step is in lines 12–14. The rest of the algorithm
 354 is concerned with the second step. In the sections that follow, we will describe each
 355 step of the algorithm and, in particular, what are and how to compute $\text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 356 and $\mathcal{G}_{\text{PAR}}^{-L}$, two key constructions used in our decision procedure.

Algorithm 1: Nash equilibrium via Parity games

```

1 Input: An LTL game  $\mathcal{G}_{\text{LTL}} = (\mathbb{N}, (\mathcal{A}c_i)_{i \in \mathbb{N}}, \text{St}, s_0, \text{tr}, \lambda, (\gamma_i)_{i \in \mathbb{N}})$ .
2 Output: “Yes” if  $\text{NE}(\mathcal{G}_{\text{LTL}}) \neq \emptyset$ ; “No” otherwise.
3  $\mathcal{G}_{\text{PAR}} \leftarrow \mathcal{G}_{\text{LTL}}$ ; /* from Section 4 (Theorem 1) */
4 foreach  $W \subseteq \mathbb{N}$  do
5   foreach  $j \in L = \mathbb{N} \setminus W$  do
6     Compute  $\text{Pun}_j(\mathcal{G}_{\text{PAR}})$ ; /* from Section 5 (Theorem 2) */
7   end
8   Compute  $\mathcal{G}_{\text{PAR}}^{-L}$ 
9   foreach  $i \in W$  do
10    Compute  $\mathcal{A}_i$  and  $\mathcal{S}_i$  from  $\mathcal{G}_{\text{PAR}}^{-L}$ 
11  end
12  if  $\mathcal{L}(\times_{i \in W} (\mathcal{S}_i)) \neq \emptyset$ ; /* from Section 5 (Theorem 3) */
13    then
14      return “Yes”
15    end
16 end
17 return “No”

```

357 **Complexity.** The procedure presented above runs in doubly exponential time, match-
 358 ing the *optimal* upper bound of the problem. In the first step we obtain a doubly ex-
 359 ponential blowup. The underlying structure \mathcal{M} of the obtained Parity game \mathcal{G}_{PAR}

360 is doubly exponential in the size of the goals of the input LTL game \mathcal{G}_{LTL} , but the
 361 priority functions set $(\alpha_i)_{i \in \mathbb{N}}$ is only (singly) exponential. Then, in the second step,
 362 reasoning takes only polynomial time in the size of the underlying concurrent game
 363 structure of \mathcal{G}_{PAR} , but exponential time in both the number of players and the size of
 364 the priority functions set. Finally, the third step takes only polynomial time, leading
 365 to an overall 2EXPTIME complexity.

366 4. From LTL to Parity

367 We now describe how to realise line 3 of Algorithm 1, and in doing so we prove a
 368 strong correspondence between the set of Nash equilibria of the input LTL game \mathcal{G}_{LTL}
 369 and the set of Nash equilibria of its associated Parity game \mathcal{G}_{PAR} . This result al-
 370 lows us to shift reasoning on the set of Nash equilibria of \mathcal{G}_{LTL} into reasoning on
 371 the set of Nash equilibria of \mathcal{G}_{PAR} . The basic idea behind this step of the decision
 372 procedure is to transform all LTL goals $(\gamma_i)_{i \in \mathbb{N}}$ in \mathcal{G}_{LTL} into a collection of DPWs,
 373 denoted by $(\mathcal{A}_{\gamma_i})_{i \in \mathbb{N}}$, that will be used to build the underlying CGS of \mathcal{G}_{PAR} . We
 374 construct \mathcal{G}_{PAR} as follows.

375 In general, using the results in [32, 33], from any LTL formula φ over AP one can
 376 build a DPW $\mathcal{A}_\varphi = \langle 2^{\text{AP}}, Q, q^0, \rho, \alpha \rangle$ such that, $\mathcal{L}(\mathcal{A}_\varphi) = \{\pi \in (2^{\text{AP}})^\omega : \pi \models \varphi\}$,
 377 that is, the language accepted by \mathcal{A}_φ is exactly the set of words over 2^{AP} that are
 378 models of φ . The size of Q is doubly exponential in $|\varphi|$ and the size of the range
 379 of α is singly exponential in $|\varphi|$. Using this construction we can define, for each LTL
 380 goal γ_i , a DPW \mathcal{A}_{γ_i} .

381 **Definition 3.** Let $\mathcal{G}_{\text{LTL}} = (\mathcal{M}, \lambda, (\gamma_i)_{i \in \mathbb{N}})$ be an LTL game whose underlying CGS
 382 is $\mathcal{M} = (\mathbb{N}, (\text{Ac}_i)_{i \in \mathbb{N}}, \text{St}, s_0, \text{tr})$, and let $\mathcal{A}_{\gamma_i} = \langle 2^{\text{AP}}, Q_i, q_i^0, \rho_i, \alpha_i \rangle$ be the DPW
 383 corresponding to player i 's goal γ_i in \mathcal{G}_{LTL} . The *Parity game* \mathcal{G}_{PAR} associated to \mathcal{G}_{LTL} is
 384 $\mathcal{G}_{\text{PAR}} = (\mathcal{M}', (\alpha'_i)_{i \in \mathbb{N}})$, where $\mathcal{M}' = (\mathbb{N}, (\text{Ac}_i)_{i \in \mathbb{N}}, \text{St}', s'_0, \text{tr}')$ and $(\alpha'_i)_{i \in \mathbb{N}}$ are as
 385 follows:

- 386 • $\text{St}' = \text{St} \times \prod_{i \in \mathbb{N}} Q_i$ and $s'_0 = (s_0, q_1^0, \dots, q_n^0)$;
- 387 • for each state $(s, q_1, \dots, q_n) \in \text{St}'$ and action profile \vec{a} ,
 388 $\text{tr}'((s, q_1, \dots, q_n), \vec{a}) = (\text{tr}(s, \vec{a}), \rho_1(q_1, \lambda(s)), \dots, \rho_n(q_n, \lambda(s)))$;

$$\bullet \alpha'_i(s, q_1, \dots, q_n) = \alpha_i(q_i).$$

Intuitively, the game \mathcal{G}_{PAR} is the product of the LTL game \mathcal{G}_{LTL} and the collection of parity (word) automata \mathcal{A}_{γ_i} that recognise the models of each player's goal. Informally, the game executes in parallel the original LTL game together with the automata built on top of the LTL goals. At every step of the game, the first component of the product state follows the transition function of the original game \mathcal{G}_{LTL} , while the “automata” components are updated according to the labelling of the current state of \mathcal{G}_{LTL} . As a result, the execution in \mathcal{G}_{PAR} is made, component by component, by the original execution, say π , in the LTL game \mathcal{G}_{LTL} , paired with the unique runs of the DPWs \mathcal{A}_{γ_i} generated when reading the word $\lambda(\pi)$.

Observe that in the translation from \mathcal{G}_{LTL} to its associated \mathcal{G}_{PAR} the set of actions for each player is unchanged. This, in turn, means that the set of strategies in both \mathcal{G}_{LTL} and \mathcal{G}_{PAR} is the same, since for every state $s \in \text{St}$ and action profile \vec{a} , it follows that \vec{a} is available in s if and only if it is available in $(s, q_1, \dots, q_n) \in \text{St}'$, for all $(q_1, \dots, q_n) \in \times_{i \in \mathbb{N}} Q_i$. Using this correspondence between strategies in \mathcal{G}_{LTL} and strategies in \mathcal{G}_{PAR} , we can prove the following Lemma, which states an invariance result between \mathcal{G}_{LTL} and \mathcal{G}_{PAR} with respect to the satisfaction of players' goals.

Lemma 1 (Goals satisfaction invariance). *Let \mathcal{G}_{LTL} be an LTL game and \mathcal{G}_{PAR} its associated Parity game. Then, for every strategy profile $\vec{\sigma}$ and player i , it is the case that $\pi(\vec{\sigma}) \models \gamma_i$ in \mathcal{G}_{LTL} if and only if $\pi(\vec{\sigma}) \models \alpha_i$ in \mathcal{G}_{PAR} .*

Proof. We prove the statement by double implication. To show the left to right implication, assume that $\pi(\vec{\sigma}) \models \gamma_i$ in \mathcal{G}_{LTL} , for any player $i \in \mathbb{N}$, and let π denote the infinite path generated by $\vec{\sigma}$ in \mathcal{G}_{LTL} ; thus, we have that $\lambda(\pi) \models \gamma_i$. On the other hand, let π' denote the infinite path generated in \mathcal{G}_{PAR} by the same strategy profile $\vec{\sigma}$. Observe that the first component of π' is exactly π . Moreover, consider the $(i + 1)$ -th component ρ_i of π' . By the definition of \mathcal{G}_{PAR} , it holds that ρ_i is the run executed by the automaton \mathcal{A}_{γ_i} when the word $\lambda(\pi)$ is read. By the definition of the labelling function of \mathcal{G}_{PAR} , it holds that the parity of π' according to α'_i corresponds to the one recognised by \mathcal{A}_{γ_i} in ρ_i . Thus, since we know that $\lambda(\pi) \models \gamma_i$, it follows that ρ_i is accepting in \mathcal{A}_{γ_i} and therefore $\pi' \models \alpha_i$, which implies that $\pi(\vec{\sigma}) \models \alpha_i$ in \mathcal{G}_{PAR} . For

419 the other direction, observe that all implications used above are equivalences. Using
 420 those equivalences one can reason backwards to prove the statement. \square

421 Using Lemma 1 we can then show that the set of Nash Equilibria for any LTL
 422 game exactly corresponds to the set of Nash equilibria of its associated Parity game.
 423 Formally, we have the following invariance result between games.

424 **Theorem 1** (Nash equilibrium invariance). *Let \mathcal{G}_{LTL} be an LTL game and \mathcal{G}_{PAR} its*
 425 *associated Parity game. Then, $\text{NE}(\mathcal{G}_{\text{LTL}}) = \text{NE}(\mathcal{G}_{\text{PAR}})$.*

426 *Proof.* The proof proceeds by double inclusion. First, assume that a strategy pro-
 427 file $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{LTL}})$ is a Nash Equilibrium in \mathcal{G}_{LTL} and, by contradiction, it is not a Nash
 428 Equilibrium in \mathcal{G}_{PAR} . Observe that, due to Lemma 1, we know that the set of players
 429 that get their goals satisfied by $\pi(\vec{\sigma})$ in \mathcal{G}_{LTL} (the “winners”, W) is the same set of play-
 430 ers that get their goals satisfied by $\pi(\vec{\sigma})$ in \mathcal{G}_{PAR} . Then, there is player $j \in L = N \setminus W$
 431 and a strategy σ'_j such that $\pi((\vec{\sigma}_{-j}, \sigma'_j)) \models \alpha_j$ in \mathcal{G}_{PAR} . Then, due to Lemma 1, we
 432 have that $\pi((\vec{\sigma}_{-j}, \sigma'_j)) \models \gamma_j$ in \mathcal{G}_{LTL} and so σ'_j would be a beneficial deviation for
 433 player j in \mathcal{G}_{LTL} too—a contradiction. On the other hand, for every $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{PAR}})$,
 434 we can reason in a symmetric way and conclude that $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{LTL}})$. \square

435 5. Characterising Nash Equilibria

436 Thanks to Theorem 1, we can focus our attention on Parity games, since a tech-
 437 nique for solving such games will also provide a technique for solving their associated
 438 LTL games. To do this we characterise the set of Nash equilibria in the Parity game
 439 construction \mathcal{G}_{PAR} in our algorithm. The existence of Nash Equilibria in LTL games
 440 can be characterised in terms of punishment strategies and memoryful reasoning [34].
 441 We will show that a similar characterisation holds here in a parity games framework,
 442 where only memoryless reasoning is required. To do this, we first introduce the notion
 443 of punishment strategies and regions formally, as well as some useful definitions and
 444 notations. In what follows, given a (memoryless) strategy profile $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$
 445 defined on a state $s \in \text{St}$ of a Parity game \mathcal{G}_{PAR} , that is, such that $s_i^0 = s$ for every

446 $i \in \mathbb{N}$, we write $\mathcal{G}_{\text{PAR}}, \vec{\sigma}, s \models \alpha_i$ if $\pi(\vec{\sigma}) \models \alpha_i$ in \mathcal{G}_{PAR} . Moreover, if $s = s_0$ is the
 447 initial state of the game, we omit it and simply write $\mathcal{G}_{\text{PAR}}, \vec{\sigma} \models \alpha_i$ in such a case.

448 **Definition 4** (Punishment strategies and regions). For a Parity game \mathcal{G}_{PAR} and a
 449 player $i \in \mathbb{N}$, we say that $\vec{\sigma}_{-i}$ is a *punishment (partial) strategy profile* against i in a
 450 state s if, for all strategies $\sigma'_i \in \Sigma_i$, it is the case that $\mathcal{G}_{\text{PAR}}, (\vec{\sigma}_{-i}, \sigma'_i), s \not\models \alpha_i$. A state
 451 s is *punishing* for i if there exists a punishment (partial) strategy profile against i in s .
 452 By $\text{Pun}_i(\mathcal{G}_{\text{PAR}})$ we denote the set of punishing states, the *punishment region*, for i in
 453 \mathcal{G}_{PAR} .

454 To understand the meaning of a punishment (partial) strategy profile, it is useful
 455 to think of a modification of the game \mathcal{G}_{PAR} , in which player i still has its goal α_i ,
 456 while the rest of the players are collectively playing in an adversarial mode, *i.e.*, try-
 457 ing to make sure that i does not achieve α_i . This scenario is represented by a two-
 458 player zero-sum game in which the winning strategies of the (coalition) player, de-
 459 noted by $-i$, correspond (one-to-one) to the punishment strategies in the original
 460 game \mathcal{G}_{PAR} . As described in [34], knowing the set of punishment (partial) strategy
 461 profiles in a given game is important to compute its set of Nash Equilibria. For this
 462 reason, it is useful to compute the set $\text{Pun}_i(\mathcal{G}_{\text{PAR}})$, that is, the set of states in the
 463 game from which a given player i can be punished. (*e.g.*, to deter undesirable unilat-
 464 eral player deviations). To do this, we reduce the problem to computing a winning
 465 strategy in a turn-based two-player zero-sum parity game, whose definition is as fol-
 466 lows.

Definition 5. For a (concurrent multi-player) Parity game

$$\mathcal{G}_{\text{PAR}} = (\mathbb{N}, \text{St}, (\text{Ac}_i)_{i \in \mathbb{N}}, s_0, \text{tr}, (\alpha_i)_{i \in \mathbb{N}})$$

467 and player $j \in \mathbb{N}$, the *sequentialisation* of \mathcal{G}_{PAR} with respect to player j is the (turn-
 468 based two-player) parity game $\mathcal{G}_{\text{PAR}}^j = \langle V_0, V_1, E, \alpha \rangle$ where

- 469 • $V_0 = \text{St}$ and $V_1 = \text{St} \times \vec{\text{Ac}}_{-j}$;
- 470 • $E = \{(s, (s, \vec{a}_{-j})) \in \text{St} \times (\text{St} \times \vec{\text{Ac}}_{-j})\} \cup \{((s, \vec{a}_{-j}), s') \in (\text{St} \times \vec{\text{Ac}}_{-j}) \times \text{St} :$



Figure 3: Sequentialisation of a game. On the left, a representation of a transition from s_1 to s_2 using action profile (\vec{a}_{-j}, a_j) . On the right, the two states s_1 and s_2 are assigned to Player 0 in the parity game, which are interleaved with a state of Player 1 corresponding to the choice of \vec{a}_{-j} by coalition $-j$ in the original game.

$$\exists a'_j \in \text{Ac}_j. s' = \text{tr}(s, (\vec{a}_{-j}, a'_j));$$

• $\alpha : V_0 \cup V_1 \rightarrow \mathbb{N}$ is such that

$$\alpha(s) = \alpha_j(s) + 1 \text{ and } \alpha(s, \vec{a}_{-j}) = \alpha_j(s) + 1.$$

The formal connection between the notion of punishment in \mathcal{G}_{PAR} and the set of winning strategies in $\mathcal{G}_{\text{PAR}}^j$ is established in the following theorem, where by $\text{Win}_0(\mathcal{G}_{\text{PAR}}^j)$ we denote the winning region of Player 0 in $\mathcal{G}_{\text{PAR}}^j$, that is, the states from which Player 0, representing the set of players $-j = \mathbb{N} \setminus \{j\}$ (the coalition of players not including j), has a memoryless winning strategy against player j in the two-player zero-sum parity game $\mathcal{G}_{\text{PAR}}^j$.

Theorem 2. *For all states $s \in \text{St}$, it is the case that $s \in \text{Pun}_j(\mathcal{G}_{\text{PAR}})$ if and only if $s \in \text{Win}_0(\mathcal{G}_{\text{PAR}}^j)$. In other words, it holds that $\text{Pun}_j(\mathcal{G}_{\text{PAR}}) = \text{Win}_0(\mathcal{G}_{\text{PAR}}^j) \cap \text{St}$.*

Proof. The proof goes by double inclusion. From left to right, assume $s \in \text{Pun}_j(\mathcal{G}_{\text{PAR}})$ and let $\vec{\sigma}_{-j}$ be a punishment strategy profile against player j in s , i.e., such that $\mathcal{G}_{\text{PAR}}, (\vec{\sigma}_{-j}, \sigma'_j), s \not\models \alpha_j$, for every strategy $\sigma'_j \in \Sigma_j$ of player j . We now define a strategy σ_0 for player 0 in $\mathcal{G}_{\text{PAR}}^j$ that is winning in s . In order to do this, first observe that, for every finite path $\pi'_{\leq k} \in V^* \cdot V_0$ in $\mathcal{G}_{\text{PAR}}^j$ starting from s , there is a unique finite sequence of action profiles $\vec{a}_{-j}^0, \dots, \vec{a}_{-j}^k$ and a sequence $\pi_{\leq k} = s^0, \dots, s^{k+1}$ of states in St^* such that

$$\pi'_{\leq k} = s^0, (s^0, \vec{a}_{-j}^0), \dots, s^k, (s^k, \vec{a}_{-j}^k), \dots, s^{k+1}.$$

Now, for every path $\pi'_{\leq k}$ of this form that is consistent with $\vec{\sigma}_{-j}$, i.e., the sequence $\vec{a}_{-j}^0, \dots, \vec{a}_{-j}^{k-1}$ is generated by $\vec{\sigma}_{-j}$, define $\sigma_0(\pi'_{\leq k}) = (s^{k+1}, \vec{a}_{-j}^{k+1})$, where \vec{a}_{-j}^{k+1} is the action profile selected by $\vec{\sigma}_{-j}$. To prove that σ_0 is winning, consider a strategy σ_1 for Player 1 and the infinite path $\pi' = \pi((\sigma_0, \sigma_1))$ generated by (σ_0, σ_1) . It is

486 not hard to see that the sequence π'_{odd} of odd positions in π' belongs to a path π in
 487 \mathcal{G}_{PAR} and it is consistent with $\vec{\sigma}_{-j}$. Thus, since $\vec{\sigma}_{-j}$ is a punishment strategy, π'_{odd}
 488 does not satisfy α_j . Moreover, observe that the parity of the sequence π'_{even} of even
 489 positions equals that of π'_{odd} . Thus, we have that $\text{Inf}(\lambda'(\pi')) + 1 = \text{Inf}(\lambda'(\pi'_{\text{odd}})) +$
 490 $1 \cup \text{Inf}(\lambda'(\pi'_{\text{even}})) + 1 = \text{Inf}(\lambda(\pi))$ and so π' is winning for player 0 in $\mathcal{G}_{\text{PAR}}^j$ and σ_0
 491 is a winning strategy.

492 From right to left, let $s \in \text{St} \cap \text{Win}_0(\mathcal{G}_{\text{PAR}}^j)$ and let σ_0 be a winning strat-
 493 egy for Player 0 in $\mathcal{G}_{\text{PAR}}^j$, and assume σ_0 is memoryless. Now, for every player i ,
 494 with $i \neq j$, define the memoryless strategy σ_i in \mathcal{G}_{PAR} such that, for every $s' \in \text{St}$,
 495 if $\sigma_0(s') = (s', \vec{a}_{-j})$, then $\sigma_i(s') = (\vec{a}_{-j})_i$ ⁷, i.e., the action that player i takes in
 496 σ_0 at s' . Now, consider the (memoryless) strategy profile $\vec{\sigma}_{-j}$ given by the com-
 497 position of all strategies σ_i , and consider a play π in \mathcal{G}_{PAR} , starting from s , that
 498 is consistent with $\vec{\sigma}_{-j}$. Thus, there exists a play π' in $\mathcal{G}_{\text{PAR}}^i$, consistent with σ_0 ,
 499 such that $\pi = \pi'_{\text{odd}}$. Moreover, since $\pi'_{\text{odd}} = \pi'_{\text{even}}$, we have that $\text{Inf}(\lambda'(\pi')) =$
 500 $\text{Inf}(\lambda'(\pi'_{\text{odd}})) \cup \text{Inf}(\lambda'(\pi'_{\text{even}})) = \text{Inf}(\lambda(\pi)) - 1$. Since π' is winning for Player 0, we
 501 know that $\pi \not\models \alpha_j$ and so $\vec{\sigma}_{-j}$ is a punishment strategy against Player j in s . \square

502 Definition 5 and Theorem 2 not only make a bridge from the notion of punish-
 503 ment strategy to the notion of winning strategy for two-player zero-sum games, but
 504 also provide a way to understand how to compute punishment regions as well as
 505 how to synthesise an actual punishment strategy in Parity games. In this way, by
 506 computing winning regions and winning strategies in these games we can solve the
 507 *synthesis* problem for individual players in the original game with LTL goals, one of
 508 the problems we are interested in. Thus, from Definition 5 and Theorem 2, we have
 509 the following corollary.

510 **Corollary 1.** *Computing $\text{Pun}_i(\mathcal{G}_{\text{PAR}})$ can be done in polynomial time with respect*
 511 *to the size of the underlying graph of the game \mathcal{G}_{PAR} and exponential in the size of the*
 512 *priority function α_i , that is, to the size of the range of α_i . Moreover, there is a memoryless*
 513 *strategy $\vec{\sigma}_i$ that is a punishment against player i in every state $s \in \text{Pun}_i(\mathcal{G}_{\text{PAR}})$.*

⁷By an abuse of notation, we let $\sigma_i(s')$ be the value of $\tau_i(s')$.




Figure 4: Representation of the strategy σ_i . At the beginning, player i follows the transducer \mathbb{T}_η that generates the action profile run η . The strategy adheres to it until a unilateral deviation from player j occurs, here represented at the k -th step of the play. Once the deviation has occurred, and the game entered a state s' , player i starts executing the strategy $\sigma_i^{\text{pun}j}$, to employ the punishment strategy against player j .

514 As described in [34], in any (infinite) run *sustained* by a Nash equilibrium $\vec{\sigma}$ in
515 deterministic and pure strategies, that is, in $\pi(\vec{\sigma})$, it is the case that all players that
516 do not get their goals achieved in $\pi(\vec{\sigma})$ can deviate from such a (Nash equilibrium)
517 run only to states where they can be punished by the coalition consisting of all other
518 players in the game. To formalise this idea in the present setting, we need one more
519 concept about punishments, defined next.

520 **Definition 6.** An action profile run $\eta = \vec{a}_0, \vec{a}_1, \dots \in \vec{A}c^\omega$ is *punishing-secure* in s for
521 player j if, for all $k \in \mathbb{N}$ and a'_j , we have $\text{tr}(\pi_j, ((\vec{a}_k)_{-j}, a'_j)) \in \text{Pun}_j(\mathcal{G}_{\text{PAR}})$, where
522 π is the only play in \mathcal{G}_{PAR} starting from s and generated by η .

523 Using the above definition, we can characterise the set of Nash equilibria of a
524 given game. Recall that strategies are formalised as transducers, *i.e.*, as finite state
525 machines with output, so such Nash equilibria strategy profiles produce runs which
526 are *ultimately periodic*. Moreover, since in every run π there are players who get their
527 goals achieved in π (and therefore do not have an incentive to deviate from π) and
528 players who do not get their goals achieve in π (and therefore may have an incentive
529 to deviate from π), we will also want to explicitly refer to such players. To do that, the
530 following notation will be useful: Let $W(\mathcal{G}_{\text{PAR}}, \vec{\sigma}) = \{i \in \mathbb{N} : \mathcal{G}_{\text{PAR}}, \vec{\sigma} \models \alpha_i\}$ denote
531 the set of player that get their goals achieved in $\pi(\vec{\sigma})$. We also write $W(\mathcal{G}_{\text{PAR}}, \pi) =$
532 $\{i \in \mathbb{N} : \mathcal{G}_{\text{PAR}}, \pi \models \alpha_i\}$.

533 **Theorem 3** (Nash equilibrium characterisation). *For a Parity game \mathcal{G}_{PAR} , there is a*
534 *Nash Equilibrium strategy profile $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{PAR}})$ if and only if there is an ultimately*

535 periodic action profile run η such that, for every player $j \in L = \mathbb{N} \setminus W(\mathcal{G}_{\text{PAR}}, \pi)$, the
 536 run η is punishing-secure for j in state s_0 , where π is the unique path generated by η
 537 from s_0 .

538 *Proof.* The proof is by double implication. From left to right, for $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{PAR}})$, let
 539 η be the ultimately periodic sequence of action profiles generated by $\vec{\sigma}$. Moreover,
 540 assume for a contradiction that η is not punishing-secure for some $j \in L$. By the
 541 definition of punishment-secure, there is $k \in \mathbb{N}$ and action $a'_j \in \text{Ac}_j$ for player j
 542 such that $s' = \text{tr}(\pi_k, ((\vec{a}_k)_{-j}, a'_j)) \notin \text{Pun}_j(\mathcal{G}_{\text{PAR}})$. Now, consider the strategy σ'_j that
 543 follows η up to the $(k-1)$ -th step, executes action a'_j on step k to get into state s' , and
 544 applies a strategy that achieves α_j from that point onwards. Note that such a strategy
 545 is guaranteed to exist since $s' \notin \text{Pun}_j(\mathcal{G}_{\text{PAR}})$. Therefore, $\mathcal{G}_{\text{PAR}}, (\vec{\sigma}_{-j}, \sigma'_j) \models \alpha_j$
 546 and so σ'_j is a beneficial deviation for player j , a contradiction to $\vec{\sigma}$ being a Nash
 547 equilibrium.

548 From right to left, we need to define a Nash equilibrium $\vec{\sigma}$ assuming only the
 549 existence of η . First, recall that η can be generated by a finite transducer $\mathbb{T}_\eta =$
 550 $(Q_\eta, q_\eta^0, \delta_\eta, \tau_\eta)$ where $\delta_\eta : Q_\eta \rightarrow Q_\eta$ and $\tau_\eta : Q_\eta \rightarrow \vec{\text{Ac}}$. Moreover, for every
 551 player i and deviating player j , with $i \neq j$, there is a (memoryless) strategy $\sigma_i^{\text{pun}_j}$ to
 552 punish player j in every state in $\text{Pun}_j(\mathcal{G}_{\text{PAR}})$. By suitably combining the transducer
 553 with the punishment strategies, we define the following strategy $\sigma_i = (Q_i, q_i^0, \delta_i, \tau_i)$
 554 for player i where

- 555 • $Q_i = \text{St} \times Q_\eta \times (L \cup \{\top\})$ and $q_i^0 = (s^0, q_\eta^0, \top)$;
- 556 • $\delta_i = Q_i \times \vec{\text{Ac}} \rightarrow Q_i$ is defined as

$$557 \quad \delta_i((s, q, \top), \vec{a}) = \begin{cases} (\text{tr}(s, \vec{a}), \delta_\eta(q), \top), & \text{if } a = \tau_\eta(q) \\ (\text{tr}(s, \vec{a}), \delta_\eta(q), j), & a_{-j} = (\tau_\eta(q))_{-j} \text{ and } \vec{a}_j \neq (\tau_\eta(q))_j^8 \\ \perp, & \text{otherwise} \end{cases}$$
- 558 • $\tau_i : Q_i \rightarrow \text{Ac}_i$ is such that

⁸For completeness, the function δ_i is assumed to take an available action. However, this is not important, as it is clear from the proof we never use this case.

559 $-\tau_i(s, q, \top) = (\tau_\eta(q))_i$, and

560 $-\tau_i(s, q, j) = \sigma_i^{\text{pun}j}(s)$.

561 To understand how strategy σ_i works, observe that its set of internal states is given
 562 by the following triple. The first component is a state of the game, remembering
 563 the position of the execution. The second component is a state of the transducer
 564 \top_η , which is used to employ the execution of the action profile run η . The third
 565 component is either the symbol \top , used to flag that no deviation has occurred, or the
 566 name of a losing player j , used to remember that such a player has deviated from η .
 567 At the beginning of the play, strategy σ_i starts executing the actions prescribed by
 568 the transducer \top_η . It sticks to it until some losing player j performs a deviation. In
 569 such a case, the third component of the internal state of σ_i switches to remember the
 570 deviating player. Moreover, from that point on, it starts executing the punishment
 571 strategy $\sigma_i^{\text{pun}j}$. Recall that parity conditions are prefix-independent. Therefore, no
 572 matter the result of the execution, if all the players start playing according to the
 573 punishment strategy $\sigma_i^{\text{pun}j}$, the resulting path will not satisfy the parity condition
 574 α_j . Now, define σ to be the collection of all σ_i . It remains to prove that $\vec{\sigma}$ is a Nash
 575 Equilibrium.

576 First, observe that since $\vec{\sigma}$ produces exactly η , we have $W(\mathcal{G}_{\text{PAR}}, \vec{\sigma}) = W(\mathcal{G}_{\text{PAR}}, \eta)$,
 577 that is, the players that get their goals achieved in $\pi(\vec{\sigma})$ and η are the same. Thus,
 578 only players in L could have a beneficial deviation. Now, consider a player $j \in L$
 579 and a strategy σ'_j and let $k \in \mathbb{N}$ be the minimum (first) step where σ'_j produces
 580 an outcome that differs from σ_j when executed along with $\vec{\sigma}_{-j}$. We write π' for
 581 $\pi((\vec{\sigma}_{-j}, \sigma'_j))$. Thus, we have $\pi_h = \pi'_h$ for all $h \leq k$ and $\pi_{k+1} \neq \pi'_{k+1}$. Hence $\pi'_{k+1} =$
 582 $\text{tr}(\pi'_k, (\eta_k)_{-j}, a'_j) = \text{tr}(\pi_k, (\eta_k)_{-j}, a'_j) \in \text{Pun}_j(\mathcal{G}_{\text{PAR}})$ and $\mathcal{G}_{\text{PAR}}, (\vec{\sigma}_{-j}, \sigma'_j) \not\models \alpha_j$,
 583 since σ_{-j} is a punishment strategy from π'_{k+1} . Thus, there is no beneficial deviation
 584 for j and $\vec{\sigma}$ is a Nash equilibrium. \square

585 6. Computing Nash Equilibria

586 Theorem 3 allows us to reduce the problem of finding a Nash equilibrium to finding
 587 a path in the game satisfying certain properties, which we will show how to check

588 using DPW and DSW automata. To do this, let us fix a given set $W \subseteq N$ of players
 589 in a given game \mathcal{G}_{PAR} , which are assumed to get their goals achieved. Now, due to
 590 Theorem 3, we have that an action profile run η corresponds to a Nash equilibrium
 591 with W being the set of “winners” in the game if, and only if, the following two
 592 properties are satisfied:

- 593 • η is punishment-secure for j in s^0 , for all $j \in L = N \setminus W$;
- 594 • $\mathcal{G}_{\text{PAR}}, \pi \models \alpha_i$, for every $i \in W$;

595 where π is, as usual, the path generated by η from s^0 .

596 To check the existence of such η , we have to check these two properties. First,
 597 note that, for η to be punishment-secure for every losing player $j \in L$, the game
 598 has to remain in the punishment region of each j . This means that an acceptable
 599 action profile run needs to generate a path that is, at every step, contained in the
 600 intersection $\bigcap_{j \in L} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$. Thus, to find a Nash equilibrium, we can remove all
 601 states not in such an intersection. We also need to remove some edges from the game.
 602 Indeed, consider a state s and a partial action profile \vec{a}_{-j} . It might be the case that
 603 $\text{tr}(s, (\vec{a}_{-j}, a'_j)) \notin \text{Pun}_j(\mathcal{G}_{\text{PAR}})$, for some $a'_j \in \text{Ac}_j$. Therefore, an action profile run
 604 that executes the partial profile \vec{a}_{-j} over s cannot be punishment-secure, and so all
 605 outgoing edges from (s, \vec{a}_{-j}) , can also be removed. After doing this for every $j \in L$,
 606 we obtain $\mathcal{G}_{\text{PAR}}^{-L}$, the game resulting from \mathcal{G}_{PAR} after the removal of the states and
 607 edges just described. As a consequence, $\mathcal{G}_{\text{PAR}}^{-L}$ has all and only the paths that can be
 608 generated by an action profile run that is punishment-secure for every $j \in L$.

609 The only thing that remains to be done is to check whether there exists a path in
 610 $\mathcal{G}_{\text{PAR}}^{-L}$ that satisfies all players in W . To do this, we use DPW and DSW automata.
 611 Since players goals are parity conditions, a path satisfying player i is an accepting
 612 run of the DPW \mathcal{A}^i where the set of states and transitions are exactly those of $\mathcal{G}_{\text{PAR}}^{-L}$
 613 and the acceptance condition is given by α_i . Then, in order to find a path satisfying
 614 the goals of all players in W , we can solve the emptiness problem of the automaton
 615 intersection $\times_{i \in W} \mathcal{A}^i$. However, observe that each \mathcal{A}^i differs from each other only in
 616 its acceptance condition α_i . Moreover, each parity condition $\alpha = (F_1, \dots, F_n)$ can be
 617 regarded as a Street condition of the form $((E_1, C_1), \dots, (E_m, C_m))$ with $m = \lceil \frac{n}{2} \rceil$

618 and $(E_i, C_i) = (F_{2i+1}, \bigcup_{j \leq i} F_{2j})$, for every $0 \leq i < m$. Therefore, the intersection
619 language of $\times_{i \in W} \mathcal{A}^i$ can be recognized by a Street automaton over the same set of
620 states and transitions and the concatenation of all the Streett conditions determined
621 by the parity conditions of the players in W . The overall translation is a DSW au-
622 tomaton with a number of Streett pairs being logarithmic in the number of its states,
623 whose emptiness can be solved in polynomial time [35]. Finally, as we fixed W at the
624 *beginning*, all we need to do is to use the procedure just described for each $W \subseteq N$, if
625 needed (see Algorithm 1).⁹

626 Concerning the complexity analysis, consider again Algorithm 1 and denote by
627 n the number of agents and $|\text{St}_{\text{LTL}}|$ the number of states. Observe that Line 3 of the
628 algorithm builds a Parity game \mathcal{G}_{PAR} by making the product construction between
629 \mathcal{G}_{LTL} and all the DPW automata \mathcal{A}_{γ_i} , whose state space is $2^{2^{|\gamma_i|}}$, and the number of
630 priorities is $2^{|\gamma_i|}$. Thus, the number of states of \mathcal{G}_{PAR} is $|\text{St}_{\text{PAR}}| = |\text{St}_{\text{LTL}}| \cdot 2^{2^{|\gamma_1|}} \cdot \dots \cdot$
631 $2^{2^{|\gamma_n|}}$. Now, on the one hand, Line 6 requires to solve a parity game on the state-graph
632 of \mathcal{G}_{PAR} with 2^{γ_i} priorities. This is solved by applying Zielonka's algorithm [37], that
633 works in time $(|\text{St}_{\text{PAR}}|)^2 \cdot (|\text{St}_{\text{PAR}}|)^{2^{\gamma_i}}$, thus polynomial in the state space of \mathcal{G}_{PAR}
634 and doubly exponential in the size of objectives γ_i 's. On the other hand, Line 12
635 calls for the Non-Emptiness procedure of a DSW whose number of Street pairs is
636 linear in the sum of priorities of the automata $\mathcal{A}_{\gamma_1}, \dots, \mathcal{A}_{\gamma_n}$ and so logarithmic in its
637 state-space (that is doubly exponential in the size of the objectives). Such procedure
638 is polynomial in the state space of the automaton [35, Corollary 10.8] and therefore
639 polynomial in $|\text{St}_{\text{PAR}}|$. Finally, consider the loops of Line 4 and Line 5,
640 respectively. The first is on all the possible subsets of agents, and thus of length 2^n .
641 The second is on all the possible agents, and thus of length n . This sums up to an
642 overall complexity for Algorithm 1 of:

$$2^n \cdot n \cdot ((|\text{St}_{\text{PAR}}|)^2 \cdot (|\text{St}_{\text{PAR}}|)^{\sum_{i \in N} 2^{\gamma_i}} + |\text{St}_{\text{PAR}}|).$$

643 Recall that $|\text{St}_{\text{PAR}}|$ is linear in the set of states of the \mathcal{G}_{LTL} and doubly exponential

⁹Some previous techniques, e.g. [36], to the computation of pure Nash equilibria are not optimal as they have exponential space complexity in the number of players $|N|$.

644 in every objective γ_i 's of the agents. Thus, the procedure is *polynomial* in $|\text{St}_{\text{LTL}}|$,
 645 exponential in N , and doubly exponential in the size of the formulas $|\gamma_1|, \dots, |\gamma_N|$.

646 7. Synthesis and Verification

647 We now show how to solve the synthesis and verification problems using NON-
 648 EMPTINESS. For *synthesis*, the solution is already contained in the proof of Theorem 3,
 649 so we only need to sketch out the approach here. Note that, in the computation of
 650 punishing regions, the algorithm builds, for every player i and potential deviator j ,
 651 a (memoryless) strategy that player i can play in the collective strategy profile $\vec{\sigma}_{-j}$
 652 in order to punish player j , should player j wishes to deviate. If a Nash equilibrium
 653 exists, the algorithm also computes a (ultimately periodic) witness of it, that is, a
 654 computation π in G , that, in particular, satisfies the goals of players in W . At this
 655 point, using this information, we are able to define a strategy σ_i for each player $i \in N$
 656 in the game (*i.e.*, including those not in W), as follows: while no deviation occurs, play
 657 the action that contributes to generate π , and if a deviation of player j occurs, then
 658 play the (memoryless) strategy $\sigma_i^{\text{pun}j}$ that is defined in the game to punish player j in
 659 case j were to deviate. Notice, in addition, that because of Lemma 1 and Theorem 1,
 660 every strategy for player i in the game with parity goals is also a valid strategy for
 661 player i in the game with LTL goals, and that such a strategy, being bisimulation-
 662 invariant, is also a strategy for every possible bisimilar representation of player i . In
 663 this way, our technique can also solve the synthesis problem for every player, that
 664 is, can compute individual bisimulation-invariant strategies for every player (system
 665 component) in the original multi-player game (concurrent system).

666 For *verification*, one can use a reduction of the following two problems, called
 667 E-NASH and A-NASH in [15, 8, 7], to NON-EMPTINESS.

668 *Given:* Game \mathcal{G}_{LTL} , LTL formula φ .

669 E-NASH: Is it the case that $\pi(\vec{\sigma}) \models \varphi$, for some $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{LTL}})$?

670 A-NASH: Is it the case that $\pi(\vec{\sigma}) \models \varphi$, for all $\vec{\sigma} \in \text{NE}(\mathcal{G}_{\text{LTL}})$?

671 We write $(\mathcal{G}_{\text{LTL}}, \varphi) \in \text{E-NASH}$ to denote that $(\mathcal{G}_{\text{LTL}}, \varphi)$ is an instance of E-NASH, *i.e.*,

672 given a game \mathcal{G}_{LTL} and a LTL formula φ , the answer to E-NASH problem is a “yes”;
 673 and, similarly for A-NASH.

674 Because we are working on a bisimulation-invariant setting, we can ensure some-
 675 thing even stronger: that for any two games \mathcal{G}_{LTL} and $\mathcal{G}'_{\text{LTL}}$, whose underlying CGSs
 676 are \mathcal{M} and \mathcal{M}' , respectively, we know that if \mathcal{M} is bisimilar to \mathcal{M}' , then $(\mathcal{G}_{\text{LTL}}, \varphi) \in$
 677 E-NASH if and only if $(\mathcal{G}'_{\text{LTL}}, \varphi) \in$ E-NASH, for all LTL formulae φ ; and, similarly for
 678 A-NASH, as desired.

679 In order to solve E-NASH and A-NASH via NON-EMPTYNESS, one could use the
 680 following result, whose proof is a simple adaptation of the same result for iterated
 681 Boolean games [15] and for multi-player games with LTL goals modelled using SRML [7],
 682 which was first presented in [38].

683 **Lemma 2.** *Let G be a game and φ be an LTL formula. There is a game H of linear size*
 684 *in G , such that $\text{NE}(H) \neq \emptyset$ if and only if $\exists \vec{\sigma} \in \text{NE}(G). \pi(\vec{\sigma}) \models \varphi$.*

685 However, since we have Algorithm 1 at our disposal, an easier – and more direct
 686 – solution can be obtained. To solve E-NASH we can modify line 12 of Algorithm 1
 687 to include the restriction that such an algorithm, which now receives φ as a param-
 688 eter, returns “Yes” in line 13 if and only if φ is satisfied in some run in the set of
 689 Nash equilibrium witnesses. The new line 12 is “if $\mathcal{L}(\times_{i \in W} (\mathcal{S}_i) \times \mathcal{S}_\varphi) \neq \emptyset$ ”, where
 690 \mathcal{S}_φ is the DSW automaton representing φ . All complexities remain the same; the
 691 modified algorithm for E-NASH is denoted as Algorithm 1'. We can then use Algo-
 692 rithm 1' to solve A-NASH, also as described in [38]: essentially, we can check whether
 693 Algorithm 1'($\mathcal{G}_{\text{LTL}}, \neg\varphi$) returns “No” in line 16. If it does, then no Nash equilibrium
 694 of \mathcal{G}_{LTL} satisfies $\neg\varphi$, either because no Nash equilibrium exists at all (thus, A-NASH
 695 is vacuously true) or because all Nash equilibria of \mathcal{G}_{LTL} satisfy φ , then solving A-
 696 NASH positively. Note that in this case, since A-NASH is solved positively when the
 697 algorithm returns “No” in line 16, then no specific Nash equilibrium strategy profile
 698 is synthesised, as expected. However, if the algorithm returns “Yes”, that is, the case
 699 when the answer to A-NASH problem with $(\mathcal{G}_{\text{LTL}}, \varphi)$ instance is negative, then a strat-
 700 egy profile is synthesised from Algorithm 1' which corresponds to a counter-example
 701 for $(\mathcal{G}_{\text{LTL}}, \varphi) \in$ A-NASH. It should be easy to see that implementing E-NASH and A-

702 NASH is straightforward from Algorithm 1. Also, as already known, it is also easy to
703 see that Algorithm 1' solves NON-EMPTYNESS if and only if $(\mathcal{G}_{\text{LTL}}, \top) \in \text{E-NASH}$.

704 8. Implementation

705 We have implemented the decision procedures presented in this paper. Our im-
706 plementation uses SRML [10] as a modelling language. SRML is based on the Reac-
707 tive Modules language [24] which is used in a number of verification tools, including
708 PRISM [26] and MOCHA [25]. The tool that implements our algorithms is called EVE
709 (for *Equilibrium Verification Environment*) [23]. EVE is the *first and only tool* able
710 to analyse the linear temporal logic properties that hold in equilibrium in a concur-
711 rent, reactive, and multi-agent system within a bisimulation-invariant framework. It
712 is also the only tool that supports all of the following combined features: a high-level
713 description language using SRML, general-sum multi-player games with LTL goals,
714 bisimulation-invariant strategies, and perfect recall. It is also the only tool for Nash
715 equilibrium analysis that relies on a procedure based on the solution of parity games,
716 which has allowed us to solve the (rational) synthesis problem for individual play-
717 ers in the system using very powerful techniques originally developed to solve the
718 synthesis problem from (linear-time) temporal logic specifications.

719 To the best of our knowledge, there are only two other tools that can be used
720 to reason about temporal logic equilibrium properties of concurrent/multi-agent sys-
721 tems: PRALINE [39] and MCMAS [40, 41].

722 PRALINE allows one to compute a Nash equilibrium in a game played in a con-
723 current game structure [39]. The underlying technique uses alternating Büchi au-
724 tomata and relies on the solution of a two-player zero-sum game called the ‘suspect
725 game’ [36]. PRALINE can be used to analyse games with different kinds of players
726 goals (*e.g.*, reachability, safety, and others), but does not permit LTL goals, and does
727 not compute bisimulation-invariant strategies.

728 MCMAS is a model checking tool for multi-agent systems [42]. Since it can be
729 used to model check Strategy Logic (SL [12]) formulae [41], and SL can express the
730 existence of a Nash equilibrium, one can model a multi-agent system in MCMAS and

731 check for the existence of a Nash equilibrium in such a system using SL. However, MC-
732 MAS only supports SL with memoryless strategies (while our implementation does
733 not have this restriction) and, as PRALINE, does not compute bisimulation-invariant
734 strategies either.

735 From the many differences between PRALINE, MCMAS, and EVE (and their asso-
736 ciated underlying reasoning and verification techniques), one of the most important
737 ones is bisimulation-invariance, a feature needed to be able to do verification and syn-
738 thesis, *e.g.*, when using symbolic methods with OBDDs or some model-minimisation
739 techniques. Not being bisimulation-invariant also means that in some cases PRALINE,
740 MCMAS, and EVE would deliver completely different answers. For instance, unlike
741 EVE, with PRALINE and MCMAS it may be the case that for two bisimilar systems
742 PRALINE and MCMAS would compute a Nash equilibrium in one of them and none
743 in the other. A particular instance is the “motivating example” in [28]. Since the two
744 systems there are bisimilar, EVE is able to compute a bisimulation-invariant Nash
745 equilibrium in both systems, while PRALINE and MCMAS, both of which are not us-
746 ing bisimulation-invariant model of strategies, cannot. The experiment supporting
747 this claim is reported in Section 8.4 along with the performance results. Indeed, even
748 in cases where all tools are able to compute a Nash equilibrium, EVE outperforms the
749 other two tools as the size of the input system grows, despite the fact that the model
750 of strategies we use in our procedure is *richer* in the sense that it takes into account
751 more information of the underlying game.¹⁰

752 8.1. Tool Description

753 *Modelling Language.* Systems in EVE are specified with the *Simple Reactive Modules*
754 *Language* (SRML [10]), that can be used to model non-deterministic systems. Each
755 system component (agent/player) in SRML is represented as a *module*, which con-
756 sists of an *interface* that defines the name of the module and lists a non-empty set of
757 Boolean variables controlled by the module, and a set of *guarded commands*, which de-

¹⁰As mentioned before, not all games can be tested in all tools since, for instance, PRALINE does not support LTL objectives, but only goals expressed directly as Büchi conditions.

758 fine the choices available to the module at each state. There are two kinds of guarded
759 commands: **init**, used for initialising the variables, and **update**, used for updating
760 variables subsequently.

761 A guarded command has two parts: a “condition” part (the “guard”) and an “ac-
762 tion” part. The “guard” determines whether a guarded command can be executed or
763 not given the current state, while the “action” part defines how to update the value
764 of (some of) the variables controlled by a corresponding module. Intuitively, $\varphi \rightsquigarrow \alpha$
765 can be read as “if the condition φ is satisfied, then *one* of the choices available to the
766 module is to execute α ”. Note that the value of φ being true does not guarantee the
767 execution of α , but only that it is *enabled* for execution, and thus *may be chosen*. If
768 no guarded command of a module is enabled in some state, then that module has no
769 choice and the values of the variables controlled by it remain unchanged in the next
770 state.

Formally, an SRML module m_i is defined as a triple $m_i = (\Phi_i, I_i, U_i)$, where
 $\Phi_i \subseteq \Phi$ is the finite set of Boolean variables controlled by m_i , I_i a finite set of **init**
guarded commands, such that for all $g \in I_i$, we have $ctr(g) \subseteq \Phi_i$, and U_i a finite
set of **update** guarded commands, such that for all $g \in U_i$, we have $ctr(g) \subseteq \Phi_i$. A
guarded command g over a set of variables Φ is an expression

$$g : \quad \varphi \rightsquigarrow x'_1 := \psi_1; \dots; x'_k := \psi_k$$

where the guard φ is a propositional logic formula over Φ , each x_i is a member of
 Φ and ψ_i is a propositional logic formula over Φ . Let $guard(g)$ denote the guard of
 g , thus, in the above rule, we have $guard(g) = \varphi$. It is required that no variable x_i
appears on the left hand side of more than one assignment statements in the same
guarded command, hence no issue on the (potentially) conflicting updates arises. The
variables x_1, \dots, x_k are controlled variables in $g \in U_i$ and we denote this set by
 $ctr(g)$. If no guarded command of a module is enabled, then the values of all variables
in $ctr(g)$ are unchanged. A set of guarded commands is said to be *disjoint* if their
controlled variables are mutually disjoint. To make it clearer, here is an example of a

```

module toggle controls x
  init
  ::  $\top \rightsquigarrow x' := \top$ ;
  ::  $\top \rightsquigarrow x' := \perp$ ;
  update
  ::  $\neg x \rightsquigarrow x' := \top$ ;
  ::  $x \rightsquigarrow x' := \perp$ ;

```

Figure 5: Example of module toggle in SRML.

guarded command:

$$\underbrace{(p \wedge q)}_{\text{guard}} \rightsquigarrow \underbrace{p' := \top; q' := \perp}_{\text{action}}$$

771 The guard is the propositional logic formula $(p \wedge q)$, so this guarded command will be
 772 enabled if both p and q are true. If the guarded command is chosen (to be executed),
 773 then in the next time-step, variable p will be assigned true and variable q will be
 774 assigned false.

775 Figure 5 shows a module named *toggle* that controls a Boolean variable named
 776 x . There are two **init** guarded commands and two **update** guarded commands. The
 777 **init** guarded commands define two choices for the initialisation of variable x : true or
 778 false. The first **update** guarded command says that if x has the value of true, then
 779 the corresponding choice is to assign it to false, while the second command says that
 780 if x has the value of false, then it can be assigned to true. Intuitively, the module
 781 would choose (in a non-deterministic manner) an initial value for x , and then on
 782 subsequent rounds toggles this value. In this particular example, the **init** commands
 783 are non-deterministic, while the **update** commands are deterministic. We refer to [7]
 784 for further details on the semantics of SRML. In particular, in Figure 12 of [7], we detail
 785 how to build a Kripke structure that models the behaviour of an SRML system. In
 786 addition, we associate each module with a goal, which is specified as an LTL formula.

787 At this point, readers might notice that the way SRML modules are defined leads
 788 to the possibility of having multiple initial states – which appears to contradict the




Figure 6: High-level workflow of EVE.

789 definition of CMGS. However, this is not a problem, since we can always add an extra
 790 “pre”-initial state whose outgoing edges are labelled according to init guarded com-
 791 mands, and use it as the “real” initial state.

792 **Automated Temporal Equilibrium Analysis.** Once a multi-agent system is mod-
 793 elled in SRML, it can be seen as a multi-player game in which players (the modules)
 794 use strategies to resolve the non-deterministic choices in the system. EVE uses Algo-
 795 rithm 1 to solve NON-EMPTYNESS. The main idea behind this algorithm is illustrated
 796 in Figure 6. The general flow of the implementation is as follows. Let \mathcal{G}_{LTL} be a game,
 797 modelled using SRML, with a set of players/modules $N = \{1, \dots, n\}$ and LTL goals
 798 $\Gamma = \{\gamma_1, \dots, \gamma_n\}$, one for each player. Using \mathcal{G}_{LTL} we construct an associated con-
 799 current game with parity goals \mathcal{G}_{PAR} in order to shift reasoning on the set of Nash
 800 equilibria of \mathcal{G}_{LTL} into the set of Nash equilibria of \mathcal{G}_{PAR} . The basic idea of this con-
 801 struction is, firstly, to transform all LTL goals in \mathcal{G}_{LTL} into deterministic parity word
 802 (DPW) automata. To do this, we use LTL2BA tool [43, 44] to transform the formulae
 803 into nondeterministic Büchi word (NBW) automata. From NBWs, we construct the
 804 associated deterministic parity word (DPW) automata via construction described in
 805 [33]. Secondly, to perform a product construction of the Kripke structure that repre-
 806 sents \mathcal{G}_{LTL} with the collection of DPWs in which the set of Nash equilibria of the input
 807 game is preserved. With \mathcal{G}_{PAR} in our hands, we can then reason about Nash equilibria
 808 by solving a collection of parity games. To solve these parity games, we use PGSolver

809 tool [45, 46]. EVE then iterates through all possible set of “winners” $W \subseteq N$ (Algo-
810 rithm 1 line 4) and computes a punishment region $\text{Pun}_j(\mathcal{G}_{\text{PAR}})$ for each $j \in L = N \setminus W$,
811 with which a reduced parity game $\mathcal{G}_{\text{PAR}}^{-L} = \bigcap_{j \in L} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$ is built. Notice that for
812 each player j , $\text{Pun}_j(\mathcal{G}_{\text{PAR}})$ need only computed once and can be stored, thus result-
813 ing in a more efficient running time. Lastly, EVE checks whether there exists a path
814 ρ in $\mathcal{G}_{\text{PAR}}^{-L}$ that satisfies the goals of each $i \in W$. To do this, we translate $\mathcal{G}_{\text{PAR}}^{-L}$ into a
815 deterministic Streett automata, whose language is empty if and only if so is the set
816 of Nash equilibria of \mathcal{G}_{PAR} . For E-NASH problem, we simply need to find a run in the
817 witness returned when we check for NON-EMPTYNESS; this can be done via automata
818 intersection¹¹.

819 EVE was developed in Python and available online from [9]. EVE takes as input
820 a concurrent and multi-agent system described in SRML code, with player goals and
821 a property φ to be checked specified in LTL. For NON-EMPTYNESS, EVE returns “YES”
822 (along with a set of winning players W) if the set of Nash equilibria in the system is
823 not empty, and returns “NO” otherwise. For E-NASH (A-NASH), EVE returns “YES” if
824 φ holds on *some (every)* Nash equilibrium of the system, and “NO” otherwise.

825 In the next subsection, we present some case studies to evaluate the performance
826 of EVE. The case studies are based on distributed and concurrent systems that can nat-
827 urally be modelled as multi-agent systems. We note, however, that such case studies
828 bear no special relevance to multi-agent systems research. Instead, our only purpose
829 is to use such case studies and multi-agent systems to evaluate EVE’s *performance*,
830 rather than to solve problems of particular relevance in the AI or multi-agent sys-
831 tems literatures. Nevertheless, one could easily see that the case studies are based on
832 systems that one can imagine to be found in many AI systems nowadays.

833 8.2. Case Studies

834 In this section, we present two examples from the literature of concurrent and
835 distributed systems to illustrate the practical usage of EVE. Among other things, these
836 two examples differ in the way they are modelled as a concurrent game. While the

¹¹For A-NASH is straightforward, since it is the dual of E-NASH.




Figure 7: Gossip framework structure.

```

module RM1 controls s1
init
:: true ~> s1':=true;
update
:: s1 ~> s1':=false;
:: s1 ~> s1':=true;
:: !s1 and (!s2 or ... or !sn)
  ~> s1':=true;
goal
:: G F (!s1);

```

Figure 8: SRML machine readable code for module RM_1 as written in EVE's input code.

837 first one is played in an arena implicitly given by the specification of the players in the
 838 game (as done in [7]), the second one is played on a graph, *e.g.*, as done in [47] with
 839 the use of concurrent game structures. Both of these models of games (modelling
 840 approaches) can be used within our tool. We will also use these two examples to
 841 evaluate EVE's practical performance and compare it against MCMAS and PRALINE
 842 in Section 8.3. Furthermore, since PRALINE and MCMAS use different modelling
 843 languages – ISPL in the case of MCMAS – we need to translate the examples modelled
 844 in SRML into PRALINE's input language and ISPL. Given the high-level nature of
 845 SRML, the translation might introduce exponential blowup. However, we argue that
 846 this is not a problem from the comparison point of view, since the exponential blowup
 847 is also unavoidable when building Kripke structures from SRML games.

848 *Gossip protocols.* These are a class of networking and communication protocols that
 849 mimic the way social networks disseminate information. They have been used to
 850 solve problems in many large-scale distributed systems, such as *peer-to-peer* and *cloud*
 851 computing systems. Ladin *et al.* [48] developed a framework to provide high avail-
 852 ability services via replication which is based on the gossip approach first introduced
 853 in [49, 50]. The main feature of this framework is the use of *replica managers* (RMs)
 854 which exchange “gossip” messages periodically in order to keep the data updated. The
 855 architecture of such an approach is shown in Figure 7.

856 We can model each RM as a module in SRML as follows: (1) When in *servicing*
 857 *mode*, an RM can choose either to keep in servicing mode or to switch to gossiping
 858 *mode*; (2) If it is in gossiping mode and there is at least another RM also in gossiping

859 mode¹², since the information during gossip exchange is of (small) bounded size, it
860 goes back to servicing mode in the subsequent step. We then set the goal of each RM
861 to be able to gossip infinitely often. As shown in Figure 8, the module RM1 controls
862 a variable: $s1$. Its value being true signifies that RM1 is in servicing mode; other-
863 wise, it is in gossiping mode. Behaviour (1) is reflected in the first and second update
864 commands, while behaviour (2) is reflected in the third update command. The goal of
865 RM1 is specified with the LTL formula $\mathbf{GF} \neg s1$, which expresses that RM1’s goal is
866 to gossip infinitely often: “always” (**G**) “eventually” (**F**) gossip ($\neg s1$).

867 Observe that with all RMs rationally pursuing their goals, they will adopt any
868 strategy which induces a run where each RM can gossip (with at least one other RM)
869 infinitely often. In fact, this kind of game-like modelling gives rise to a powerful
870 characteristic: on *all* runs that are sustained by a Nash equilibrium, the distributed
871 system is guaranteed to have two crucial *non-starvation/liveness* properties: RMs can
872 gossip infinitely often and clients can be served infinitely often. Indeed, these prop-
873 erties are verified in the experiments; with E-NASH: no Nash equilibrium sustains “all
874 RMs forever gossiping”; and with A-NASH: in all Nash equilibria at least one of the
875 RM is in servicing mode infinitely often. We also notice that each RM is modelled as
876 a non-deterministic open system: non-determinism is used in the first two updated
877 commands, as they have the same guard $s1$ and therefore will be both enabled at the
878 same time; and the system is open since each module’s state space and choices depend
879 on the states of other modules, as reflected by the third updated command.

880 *Replica Control Protocol.* Consensus is a key issue in distributed computing and multi-
881 agent systems. An important application domain is in maintaining data consistency.
882 Gifford [51] proposed a quorum-based voting protocol to ensure data consistency by
883 not allowing more than one processes to read/write a data item concurrently. To do
884 this, each copy of a replicated item is assigned a vote.

885 We can model a (modified version of) Gifford’s protocol as a game as follows. The
886 set of players $N = \{1, \dots, n\}$ in the game is arranged in a request queue represented

¹²The core of the protocol involves (at least) pairwise interactions periodically.




Figure 9: Gifford's protocol modelled as a game.

887 by the sequence of states q_1, \dots, q_n , where q_i means that player i is requesting to
 888 read/write the data item. At state q_i , other players in $N \setminus \{i\}$ then can vote whether
 889 to allow player i to read/write. If the majority of players in N vote “yes”, then the
 890 transition goes to q_0 , *i.e.*, player i is allowed to read/write, and otherwise it goes to
 891 q_{i+1} ¹³. The voting process then restarts from q_1 . The protocol's structure is shown in
 892 Figure 9. Notice that at the last state, q_n , there is only one outgoing arrow to q_0 . As
 893 in the previous example, the goal of each player i is to visit q_0 right after q_i infinitely
 894 often, so that the desired behaviour of the system is sustained on all Nash equilibria of
 895 the system: a data item is not concurrently accessed by two different processes and the
 896 data is updated in *every* round. The associated temporal properties are automatically
 897 verified in the experiments in Section 8.3. Specifically, the temporal properties we
 898 check are as follows. With E-NASH: there is no Nash equilibrium in which the data is
 899 never updated; and, with A-NASH: on all Nash equilibria, for each player, its request
 900 will be granted infinitely often. Also, in this example, we define a module, called
 901 “Environment”, which is used to represent the underlying concurrent game structure,
 902 shown in Figure 9, where the game is played.

903 8.3. Experiment I

904 In order to evaluate the practical performance of our tool and approach (against
 905 MCMAS and PRALINE), we present results on the temporal equilibrium analysis for
 906 the examples in Section 8.2. We ran the tools on the two examples with different

¹³We assume arithmetic modulo $(|N| + 1)$ in this example.

Table 1: Gossip Protocol experiment results.

P	S	E	EVE			PRALINE			MCMAS		
			ν (s)	ϵ (s)	α (s)	ν (s)	ϵ (s)	α (s)	ν (s)	ϵ (s)	α (s)
2	4	9	0.02	0.24	0.08	0.02	1.71	1.73	0.01	0.01	0.01
3	8	27	0.09	0.43	0.26	0.33	26.74	27.85	0.02	0.06	0.06
4	16	81	0.42	3.51	1.41	0.76	547.97	548.82	760.65	3257.56	3272.57
5	32	243	2.30	35.80	25.77	10.06	TO	TO	TO	TO	TO
6	64	729	16.63	633.68	336.42	255.02	TO	TO	TO	TO	TO
7	128	2187	203.05	TO	TO	5156.48	TO	TO	TO	TO	TO
8	256	6561	4697.49	TO	TO	TO	TO	TO	TO	TO	TO

Table 2: Replica control experiment results.

P	S	E	EVE			PRALINE			MCMAS		
			ν (s)	ϵ (s)	α (s)	ν (s)	ϵ (s)	α (s)	ν (s)	ϵ (s)	α (s)
2	3	8	0.04	0.11	0.10	0.05	0.64	0.74	0.01	0.01	0.02
3	4	20	0.11	1.53	0.22	0.12	4.96	5.46	0.02	0.06	0.11
4	5	48	0.34	1.73	0.68	0.56	65.50	67.45	1.99	4.15	11.28
5	6	112	1.43	2.66	2.91	6.86	1546.90	1554.80	1728.73	6590.53	TO
6	7	256	5.87	13.69	16.03	94.39	TO	TO	TO	TO	TO
7	8	576	32.84	76.50	102.12	2159.88	TO	TO	TO	TO	TO
8	9	1280	166.60	485.99	746.55	TO	TO	TO	TO	TO	TO

907 numbers of players (“P”), states (“S”), and edges (“E”). The experiments were obtained
908 on a PC with Intel i5-4690S CPU 3.20 GHz machine with 8 GB of RAM running Linux
909 kernel version 4.12.14-300.fc26.x86_64. We report the running time¹⁴ for solving NON-

¹⁴To carry out a fairer comparison (since PRALINE does not accept LTL goals), we added to PRALINE’s running time the time needed to convert LTL games into its input.

910 EMPTINESS (“ ν ”), E-NASH (“ ϵ ”), and A-NASH (“ α ”). For the last two problems, since
 911 there is no direct support in PRALINE and MCMAS, we used the reduction of E/A-
 912 NASH to NON-EMPTINESS presented in [38]. Intuitively, the reduction is as follows:
 913 given a game G and formula φ , we construct a new game H with two additional agents,
 914 say $n + 1$ and $n + 2$, with goals $\gamma_{n+1} = \varphi \vee (p \leftrightarrow q)$ and $\gamma_{n+2} = \varphi \vee \neg(p \leftrightarrow q)$,
 915 where $\Phi_{n+1} = \{p\}$ and $\Phi_{n+2} = \{q\}$, p and q are fresh Boolean variables. This means
 916 that it is the case $\text{NE}(H) \neq \emptyset$ if and only if there exists a Nash equilibrium run in G
 917 satisfying φ .

918 From the experiment results shown in Table 1 and 2, we observe that, in general,
 919 EVE has the best performance, followed by PRALINE and MCMAS. Although PRA-
 920 LINE performed better than MCMAS, both struggled (timed-out¹⁵) with inputs with
 921 more than 100 edges, while EVE could handle up to 6000 edges (for NON-EMPTINESS).

922 8.4. Experiment II

923 This experiment is taken from the motivating examples in [28]. Suppose the sys-
 924 tems shown in Figure 10 and 11 represents a 3-player game, where each transition
 925 is labelled by the actions x, y, z of player 1, 2, and 3, respectively, an asterisk * be-
 926 ing a wildcard. The goals of the players can be represented by the LTL formulae
 927 $\gamma_1 = \mathbf{F}p$, $\gamma_2 = \mathbf{F}q$, and $\gamma_3 = \mathbf{G}\neg(p \vee q)$. The system in Figure 10 has a Nash equi-
 928 librium, whereas no (non-bisimulation-invariant strategies) Nash equilibria exists in
 929 the (bisimilar) system in Figure 11.

930 In this experiment, we extended the number of states by adding more layers to
 931 the game structures used there in order to test the practical performance of EVE,
 932 MCMAS, and PRALINE. The experiments were performed on a PC with Intel i7-
 933 4702MQ CPU 2.20GHz machine with 12GB of RAM running Linux kernel version
 934 4.14.16-300.fc26.x86_64. We divided the test cases based on the number of Kripke
 935 states and edges; then, for each case, we report (i) the total running time¹⁶ (“time”)
 936 and (ii) whether the tools find any Nash equilibria (“NE”).

¹⁵Time-out was fixed to be 7200 seconds.

¹⁶Similarly to Experiment I (Section 8.3), we added to PRALINE’s running time the time needed to convert LTL games into its input to carry out a fairer comparison.

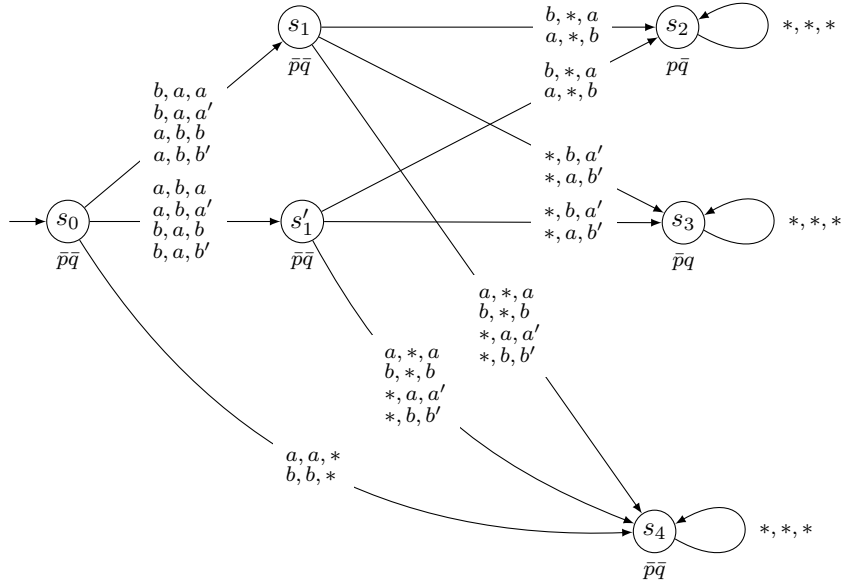


Figure 10: A 3-player game with Nash equilibrium.

937 Table 3 shows the results of the experiments on the example in which the model
 938 of strategies that depends only on the run (sequence of states) of the game (run-based
 939 strategies [28]) cannot sustain any Nash equilibria, a model of strategies that is not
 940 invariant under bisimilarity. Indeed, since MCMAS and PRALINE use this model of
 941 strategies, both did not find any Nash equilibria in the game, as shown in Table 3.
 942 EVE, which uses a model of strategies that not only depends on the run of the game
 943 but also on the actions of players (computation-based [28]), found a Nash equilibrium
 944 in the game. We can also see that EVE outperformed MCMAS on games with 14 or
 945 more states. In fact, MCMAS timed-out¹⁷ on games with 17 states or more, while EVE
 946 kept working efficiently for games of bigger size. We can also observe that PRALINE
 947 performed almost as efficiently as EVE in this experiment, although EVE performed
 948 better in both small and large instances of these games.

¹⁷We fixed the time-out value to be 3600 seconds (1 hour).




Figure 11: A 3-player game without (non-bisimulation-invariant strategies) Nash equilibria.

949 In Table 4, we used the example in which Nash equilibria is sustained in run-
 950 based strategies. As shown in the table, MCMAS found Nash equilibria in games with
 951 6 and 9 states. However, since MCMAS uses imperfect recall, when the third layer
 952 was added (case with 12 states in Table 4) to the game, it could not find any Nash
 953 equilibria. Regarding running times, EVE outperformed MCMAS from the game with
 954 12 states and beyond, where MCMAS timed-out on games with 15 or more states.
 955 As for PRALINE, it performed comparably to EVE in this experiment, but again, EVE
 956 performed better in all instances.

957 *8.5. Experiment III*

This experiment is based on the example previously presented in Section 2. For this particular experiment, we assume that initially the agents are located at opposing corners of the grid; specifically, agent 1 is located at the top-left corner (coordinate $(0, 0)$) and agent 2 at the bottom-right corner $(n - 1, n - 1)$. A number of obstacles are also placed (uniformly) randomly on the grid. We use a binary encoding to represent

Table 3: Example with no Nash equilibrium.

states	edges	MCMAS		EVE		PRALINE	
		time (s)	NE	time (s)	NE	time (s)	NE
5	80	0.04	No	0.75	Yes	0.77	No
8	128	0.24	No	2.99	Yes	2.06	No
11	176	6.28	No	3.86	Yes	4.42	No
14	224	273.14	No	7.46	Yes	8.53	No
17	272	TO	-	13.31	Yes	15.33	No
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
50	800	TO	-	655.80	Yes	789.77	No

Table 4: Example with Nash equilibria

states	edges	MCMAS		EVE		PRALINE	
		time (s)	NE	time (s)	NE	time (s)	NE
6	96	0.02	Yes	1.09	Yes	1.19	Yes
9	144	0.77	Yes	3.36	Yes	3.76	Yes
12	192	65.31	No	7.45	Yes	8.89	Yes
15	240	TO	-	15.52	Yes	17.72	Yes
18	288	TO	-	30.06	Yes	30.53	Yes
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
51	816	TO	-	1314.47	Yes	1563.79	Yes

the spatial information of the grid world which includes the grid coordinates, as well as the obstacles and the agents locations. For instance, to encode a position of an agent 1 in 4×4 grid, we need 4 Boolean variables arranged as a tuple $pos_1 = \langle x_0^1, x_1^1, y_0^1, y_1^1 \rangle$. An instance of such a tuple $pos_1 = \langle 0, 1, 1, 0 \rangle$ means that agent 1 is at (2, 1). For each time step and $i \in \{1, 2\}$, the update guarded command set U_i is such a way that agent i can only move horizontally and vertically, 1 step at a time. Furthermore,




Figure 12: Plots from Table 5. Y-axis is in logarithmic scale.

the commands in U_i respect the legality of movement, *i.e.*, agent i cannot move out of bound or into an obstacle. The goal of each agent can be expressed by the LTL formulae

$$\gamma_1 = \mathbf{F} \left(\bigwedge_{i \in \{0, \dots, n-1\}} x_i^1 \wedge \bigwedge_{i \in \{0, \dots, n-1\}} y_i^1 \right)$$

and

$$\gamma_2 = \mathbf{F} \left(\bigwedge_{i \in \{0, \dots, n-1\}} \neg x_i^2 \wedge \bigwedge_{i \in \{0, \dots, n-1\}} \neg y_i^2 \right).$$

A safety specification (no more than one agent occupying the same position at the same time) can be expressed by the following LTL formula:

$$\varphi = \mathbf{G} \neg \left(\bigwedge_{i \in \{0, \dots, n-1\}} (x_i^1 \leftrightarrow x_i^2) \wedge \bigwedge_{i \in \{0, \dots, n-1\}} (y_i^1 \leftrightarrow y_i^2) \right).$$

958 The experiment was obtained on a PC with Intel i5-4690S CPU 3.20 GHz machine
 959 with 8 GB of RAM running Linux kernel version 4.12.14-300.fc26.x86_64. We varied
 960 the size of the grid world (“size”) from 3×3 to 10×10 , each with a fixed number
 961 of obstacles (“# Obs”), randomly distributed on the grid. We report the number of
 962 Kripke states (“KS”), Kripke edges (“KE”), \mathcal{G}_{PAR} states (“GS”), \mathcal{G}_{PAR} edges (“GE”),
 963 NON-EMPTYNESS execution time (“ ν ”), and E-NASH execution time (“ ϵ ”). We ran the
 964 experiment for five replications, and report the average (*ave*), minimum (*min*), and

Table 5: Grid world experiment results.

Size	# Obs	KS	KE	GS
3	3	15(13, 18)	44(32, 72)	60(53, 73)
4	6	40(32, 52)	150(98, 200)	156(121, 209)
5	10	94(61, 125)	398(242, 512)	376(453, 741)
6	15	155(113, 185)	655(450, 800)	619(453, 741)
7	21	228(181, 290)	994(800, 1250)	909(725, 1161)
8	28	491(394, 666)	2297(1922, 2888)	1963(1577, 2665)
9	36	564(269, 765)	2687(1352, 3698)	2256(1077, 3061)
10	45	916(730, 1258)	4780(3528, 6498)	3657(2921, 5033)

Size	GE	ν (s)	ϵ (s)
3	173(129, 289)	0.44(0.19, 1.14)	1.21(0.5, 2.63)
4	595(379, 801)	0.98(0.63, 1.16)	1.57(1.01, 2.24)
5	1591(969, 2049)	4.73(2.62, 6.22)	22.51(18.22, 26.25)
6	2622(1801, 3201)	9.53(7.13, 11.49)	32.32(26.05, 37.35)
7	3969(3161, 5001)	17.69(13.81, 21.58)	48.90(39.70, 59.50)
8	9190(7689, 11553)	50.91(38.38, 72.49)	121.33(95.03, 167.25)
9	10748(5409, 14793)	100.94(45.81, 137.91)	6002.80(5477.63, 6374.26)
10	19102(14113, 25993)	211.30(152.74, 311.43)	6871.16(6340.64, 7650.87)

965 maximum (*max*) times from the replications. The results are reported in Table 5, with
966 the following format: *ave(min, max)*.

967 From the experiment results, we see that EVE works well for NON-EMPTYNESS up
968 until size 10. From the plots in Figure 12, we can clearly see that the values of each
969 variable, except for ϵ , grow exponentially. For ϵ (E-NASH), however, it seems to grow
970 faster than the rest. Specifically, it is clearly visible in transitions between numbers

971 that have different size of bit representation, *i.e.*, 4 to 5 and 8 to 9¹⁸. These jumps
972 correspond to the time used to build deterministic parity automata on words from
973 LTL properties to be checked in E-NASH, which is essentially, bit-for-bit comparisons
974 between the position of agent 1 and 2.

975 From the experiments shown in this section it is also clear that the bottleneck in
976 the performance is the translation of LTL goals and the high-level description of the
977 game into the underlying parity game. Once an explicit parity game is constructed,
978 then the performance improves radically. This result is perfectly consistent with what
979 the theoretical complexity of the decision procedure predicts: our algorithm works
980 in doubly-exponential time in the size of the goals of the players, while it is only
981 singly-exponential in the size of the SRML specification. These two exponential-time
982 reductions are in fact optimal, so there is no hope that they can be improved, at least
983 in theory. On the other hand, the actual subroutine that finds a Nash equilibrium and
984 computes players' strategies from the parity games representation of the problem is
985 rather efficient in theory – but still not known to be in polynomial time using the best
986 algorithms to solve parity games. Then, it is clear that a natural way to make rational
987 verification a feasible problem, in theory, is to look at cases where goals and/or game
988 representations are simpler. Such study is conducted in [52], where several positive
989 results on the complexity of solving the rational verification problem are obtained.

990 **9. Concluding Remarks and Related Work**

991 This paper contains a complete study, from theory to implementation, of the tem-
992 poral equilibrium analysis of multi-agent AI systems formally modelled as multi-
993 player games. The two main contributions of the paper are: (1) a novel and optimal
994 decision procedure, based on the solution of parity games, that can be used to solve
995 both the rational verification and the automated synthesis problems for multi-player
996 games; and (2) a complete implementation of the general game-theoretic modelling
997 and reasoning framework – with full support of goals expressed as LTL formulae and

¹⁸Since the grid coordinate index starts at 0, the “actual” transitions are 3 to 4 and 7 to 8.

998 high-level game descriptions in SRML – which is available online. Our work builds
999 on several previous results in the computer science (synthesis and verification) and AI
1000 literatures (multi-agent systems). Relevant related literature will be discussed next.

1001 **Equilibrium Analysis in Multi-Agent Systems.** Rational verification was pro-
1002 posed as an complementary verification methodology to conventional methods, such
1003 as model checking. A legitimate question is, then, when is rational verification an ap-
1004 propriate verification approach? A possible answer is given next. The verification
1005 problem [1], as conventionally formulated, is concerned with checking that some
1006 property, usually defined using a modal or a temporal logic [53], holds on some or
1007 on every computation run of a system. In a game-theoretic setting, this can be a
1008 very strong requirement – and in some cases even inappropriate – since only some
1009 computations of the system will arise (be sustained) as the result of agents in the sys-
1010 tem choosing strategies in equilibrium, that is, due to strategic and rational play. It
1011 was precisely this concern that motivated the rational verification approach [7, 8].
1012 In rational verification, we ask if a given temporal property holds on some or every
1013 computation run that can be sustained by agents choosing Nash equilibrium strate-
1014 gies. Rational verification can be reduced to the NON-EMPTYNESS problem, as stated
1015 in this paper; cf., [38]. As a consequence, along with the polynomial transformations
1016 in [38], our results provide a complete framework (theory, algorithms, and imple-
1017 mentation) for automated temporal equilibrium analysis, specifically, to do rational
1018 synthesis and formal verification of logic-based multi-agent systems. The framework,
1019 in particular, provides a concrete and algorithmic solution to the rational synthesis
1020 problem as studied in [14], where the Boolean case (iterated games where players
1021 control Boolean variables, whose valuations define sequences of states in the game,
1022 *i.e.*, the plays in the game) was given an interesting automata-theoretic solution via
1023 (an extension of) Strategy Logic [16].

1024 **Automata and logic.** In computer science, a common technique to reason about
1025 Nash equilibria in multi-player games is using alternating parity *automata on infinite*
1026 *trees* (APTs [18]). This approach is used to do rational synthesis [14, 54]; equilibrium
1027 checking and rational verification [8, 15, 7]; and model checking of logics for strategic

1028 reasoning capable to specify the existence of a Nash equilibrium in concurrent game
1029 structures [47], both in two-player games [16, 55] and in multi-player games [56, 12].
1030 In cases where players’ goals are simpler than general LTL formulae, *e.g.*, for reacha-
1031 bility or safety goals, alternating Büchi automata can be used instead [36]. *Our tech-*
1032 *nique is different from all these automata-based approaches, and in some cases more*
1033 *general*, as it can be used to handle either a more complex model of strategies or a
1034 more complex type of goals, and delivers an immediate procedure to synthesise indi-
1035 vidual strategies for players in the game, while being amenable to implementation.

1036 **Tools and algorithms.** In theory, the kind of equilibrium analysis that can be done
1037 using MCMAS [40, 57, 58] and PRALINE [39, 36] rely on the automata-based approach.
1038 However, the algorithms that are actually implemented have a different flavour. MC-
1039 MAS uses a procedure for SL which works as a *labelling algorithm* since it only consid-
1040 ers memoryless strategies [58]. On the other hand, PRALINE, which works for Büchi
1041 definable objectives, uses a procedure based on the “*suspect game*” [36]. Despite some
1042 similarities between our construction and the suspect game, introduced in [36], the
1043 two procedures are substantially different. Unlike our procedure, the suspect game is
1044 a standard two-player zero-sum turn-based game $\mathcal{H}(\mathcal{G}, \pi)$, constructed from a game
1045 \mathcal{G} and a possible path π , in which one of the players (“Eve”) has a winning strategy
1046 if, and only if, π can be sustained by a Nash equilibrium in \mathcal{G} . The overall procedure
1047 in [36] relies on the construction of such a game, whose size (space complexity) is
1048 exponential in the number of agents [36, Section 4.3]. Instead, our procedure solves,
1049 independently, a collection of parity games that avoids an exponential use of space but
1050 may require to be executed exponentially many times. Key to the correctness of our
1051 approach is that we deal with parity conditions, which are prefix-independent, ensur-
1052 ing that punishment strategies do not depend on the history of the game. Regarding
1053 similarities, our procedure also checks for the existence of a path sustained by a Nash
1054 Equilibrium, but our algorithm does this for every subset $W \subseteq \mathbb{N}$ of agents, if needed.
1055 Doing this (*i.e.*, trading exponential space for exponential time), at every call of this
1056 subroutine, our algorithm avoids building an exponentially sized game, like \mathcal{H} . On the
1057 other hand, from a practical point of view, avoiding the construction of such an expo-

1058 nential sized game leads to better performance (running times), even in cases where
 1059 no Nash equilibrium exists, when our subroutine is necessarily called exponentially
 1060 many times. In addition to all of the above, neither the algorithm used for MCMAS nor
 1061 the one used for PRALINE computes pure Nash equilibria in a bisimulation-invariant
 1062 framework, as our procedure does. While MCMAS and PRALINE are the two closest
 1063 tools to EVE, they are not the only available options to reason about games. For in-
 1064 stance, PRISM-games [59], EAGLE [60], and UPPAAL [61] are other interesting tools
 1065 to reason about games. PRISM-games allows one to do strategy synthesis for turn-
 1066 based stochastic games as well as model checking for long-run, average, and ratio
 1067 rewards properties. Only until very recently, PRISM-games had no support of equi-
 1068 librium reasoning, but see [62]. EAGLE is a tool specifically designed to reason about
 1069 pure Nash equilibria in multi-player games. EAGLE considers games where goals
 1070 are given as CTL formulae and allows one to check if a given strategy profile is a
 1071 Nash equilibrium of a given multi-agent system. This decision problem, called MEM-
 1072 BERSHIP within the rational verification framework [8], is, theoretically, simpler than
 1073 NON-EMPTYNESS: while the former can be solved in EXPTIME (for branching-time
 1074 goals expressed using CTL formulae [13]), the latter is 2EXPTIME-complete for LTL
 1075 goals, and even 2EXPTIME-hard for CTL goals and nondeterministic strategies [13].
 1076 UPPAAL is another tool that can be used to analyse equilibrium behaviour in a sys-
 1077 tem [63, 64]. However, UPPAAL differs from EVE in various critical ways: *e.g.*, it works
 1078 in a quantitative setting, uses statistical model checking, and most importantly, com-
 1079 putes approximate Nash equilibria of a game.

1080 **The Role of Bisimilarity.** One crucial aspect of our approach to rational verification
 1081 and synthesis is the role of *bisimilarity* [65, 31, 66, 67]. Bisimulation is the most impor-
 1082 tant type of behavioural equivalence relation considered in computer science, and in
 1083 particular two bisimilar systems will satisfy the same temporal logic properties. In our
 1084 setting, it is highly desirable that properties which hold in equilibrium are sustained
 1085 across all bisimilar systems to P_1, \dots, P_n . That is, that for every (temporal logic)
 1086 property φ and every system component P'_i modelled as an agent in a multi-player
 1087 game, if P'_i is bisimilar to $P_i \in \{P_1, \dots, P_n\}$, then φ is satisfied in equilibrium – that

1088 is, on a run induced by some Nash equilibrium of the game – by $P_1, \dots, P_i, \dots, P_n$ if
1089 and only if is also satisfied in equilibrium by $P_1, \dots, P'_i, \dots, P_n$, the system in which
1090 P_i is replaced by P'_i , that is, across all bisimilar systems to P_1, \dots, P_n . This property
1091 is called *invariance under bisimilarity*. Unfortunately, as shown in [34, 28], the satis-
1092 faction of temporal logic properties in equilibrium is not invariant under bisimilarity,
1093 thus posing a challenge for the modular and compositional reasoning of concurrent
1094 systems, since individual system components in a concurrent system cannot be re-
1095 placed by (behaviourally equivalent) bisimilar ones, while preserving the temporal
1096 logic properties that the overall multi-agent system satisfies in equilibrium. This is
1097 also a problem from a synthesis point of view. Indeed, a strategy for a system com-
1098 ponent P_i may not be a valid strategy for a bisimilar system component P'_i . As a
1099 consequence, the problem of building strategies for individual processes in the con-
1100 current system $P_1, \dots, P_i, \dots, P_n$ may not, in general, be the same as building strate-
1101 gies for a bisimilar system $P_1, \dots, P'_i, \dots, P_n$, again, deterring any hope of being able
1102 to do modular reasoning on concurrent and multi-agent systems. These problems
1103 were first identified in [34] and further studied in [28]. However, no algorithmic so-
1104 lutions to these two problems were presented in either [34] or [28]. Specifically, in
1105 this paper, bisimilarity was exploited in two ways. Firstly, our construction of punish-
1106 ment strategies (used in the characterisation of Nash equilibrium given by Theorem 3)
1107 assumes that players have access to the history of choices that other players in the
1108 game have made. As shown in [28, 29], with a model of strategies where this is not
1109 the case, the preservation of Nash equilibria in the game, as well as of temporal logic
1110 properties in equilibrium, may not be guaranteed. Secondly, our implementation in
1111 EVE guarantees that any two games whose underlying CGSs are bisimilar, and there-
1112 fore should be regarded as observationally equivalent from a concurrency point of
1113 view, will produce the same answers to the rational verification and automated syn-
1114 thesis problems. It is also worth noting that even though bisimilarity is probably the
1115 most widely used behavioural equivalence in concurrency, in the context of multi-
1116 agent systems other relations may be preferred, for instance, equivalence relations
1117 that take a detailed account of the independent interactions and behaviour of indi-
1118 vidual components in a multi-agent system. In such a setting, “alternating” relations

1119 with natural ATL* characterisations have been studied [68]. Alternating bisimulation
1120 is very similar to bisimilarity on labelled transition systems [65, 31], only that when
1121 defined on CGSs, instead of action profiles (directions) taken as possible transitions,
1122 one allows individual player’s actions, which must be matched in the bisimulation
1123 game. Because of this, it immediately follows that any alternating bisimulation as de-
1124 fined in [68] is also a bisimilarity as defined here. Despite having a different formal
1125 definition, a simple observation can be made: Nash equilibria are not preserved by
1126 the alternating (bisimulation) equivalence relations in [68] either, which discourages
1127 the use of these even stronger equivalence relations for multi-agent systems. In fact,
1128 as discussed in [69], the “right” notion of equivalence for games (which can be indi-
1129 rectly used as an observationally equivalence between multi-agent systems) and their
1130 game theoretic solution concepts is, undoubtedly, an important and interesting topic
1131 of debate, which deserves to be investigated further.

1132 **Some features of our framework.** Unlike other approaches to rational synthesis
1133 and temporal equilibrium analysis, *e.g.* [58, 36, 14, 7], we employ parity games [19],
1134 which are an intuitively *simple verification model* with an abundant associated set of
1135 algorithmic solutions [70]. In particular, strategies in our framework, as in [7], can
1136 depend on players’ actions, leading to a much richer game-theoretic setting where
1137 Nash equilibrium is invariant under bisimilarity [28, 29], a desirable property for con-
1138 current and reactive systems [65, 31, 66, 67]. Our reasoning and verification approach
1139 applies to multi-player games that are concurrent and synchronous, with perfect re-
1140 call and perfect information, and which can be represented in a high-level, succinct
1141 manner using SRML [10]. In addition, the technique developed in this paper, and its
1142 associated implementation, considers games with LTL goals, deterministic and pure
1143 strategies, and dichotomous preferences. In particular, strategies in these games are
1144 assumed to be able to see all past players’ actions. We do not consider mixed or non-
1145 deterministic strategies, or goals given by branching-time formulae. We also do not
1146 allow for quantitative or probabilistic systems, *e.g.*, such as stochastic games or similar
1147 game models. We note, however, that some of these aspects of our reasoning frame-
1148 work have been placed to avoid undesirable computational properties. For instance, it

1149 is known that checking for the existence of a Nash equilibrium in multi-player games
1150 like the ones we consider is an undecidable problem if either imperfect information
1151 or (various kinds of) quantitative/probabilistic information is allowed [17, 71].

1152 **Future Work.** This paper gives a solution to the temporal equilibrium problem (both
1153 automated synthesis and formal verification) in a noncooperative setting. In future
1154 work, we plan to investigate the cooperative games setting [72]. The paper also solves
1155 the problem in practice for perfect information games. We also plan to investigate if
1156 our main algorithms can be extended to decidable classes of imperfect information
1157 games, for instance, as those studied to model the behaviour of multi-agent systems
1158 in [17, 73, 74, 75]. Whenever possible, such studies will be complemented with prac-
1159 tical implementations in EVE. Finally, extensions to epistemic systems and quantita-
1160 tive information in the context of multi-agent systems may be another avenue for
1161 further applications [76, 77], as well as settings with more complex preference rela-
1162 tions [13, 14, 78, 79], which would provide a strictly stronger modelling power.

1163 *Acknowledgements.* The authors gratefully acknowledge the financial support of the
1164 ERC Advanced Investigator Grant 291528 (“RACE”) at Oxford. Giuseppe Perelli con-
1165 ducted this research partially while being member of the University of Oxford, work-
1166 ing on the aforementioned grant, and now supported in part by European Research
1167 Council under the European Union’s Horizon 2020 Programme through the ERC Ad-
1168 vanced Investigator Grant 834228 (“WhiteMech”). Muhammad Najib was supported
1169 by the Indonesia Endowment Fund for Education (LPDP) while working on this re-
1170 search at the University of Oxford, and now by the European Research Council (ERC)
1171 under the European Union’s Horizon 2020 research and innovation programme (grant
1172 agreement no 759969). Part of this paper, focussing on the EVE system, has been pre-
1173 sented at ATVA’18 [23].

1174 [1] E. M. Clarke, O. Grumberg, D. A. Peled, Model Checking, MIT Press, Cambridge,
1175 MA, USA, 2002.

1176 [2] M. Wooldridge, Introduction to Multiagent Systems, Wiley, Chichester, UK, 2001.

- 1177 [3] Y. Shoham, K. Leyton-Brown, Multiagent Systems: Algorithmic, Game-
1178 Theoretic, and Logical Foundations, CUP, 2008.
- 1179 [4] M. J. Osborne, A. Rubinstein, A Course in Game Theory, MIT Press, 1994.
- 1180 [5] J. Y. Halpern, Beyond nash equilibrium: Solution concepts for the 21st century,
1181 in: KR, AAAI Press, 2008, pp. 6–15.
- 1182 [6] I. Abraham, L. Alvisi, J. Y. Halpern, Distributed computing meets game theory:
1183 combining insights from two fields, SIGACT News 42 (2) (2011) 69–76.
- 1184 [7] J. Gutierrez, P. Harrenstein, M. Wooldridge, From model checking to equilibrium
1185 checking: Reactive modules for rational verification, Artificial Intelligence 248
1186 (2017) 123–157.
- 1187 [8] M. Wooldridge, J. Gutierrez, P. Harrenstein, E. Marchioni, G. Perelli, A. Toumi,
1188 Rational verification: From model checking to equilibrium checking, in: AAAI,
1189 2016, pp. 4184–4191.
- 1190 [9] EVE: A tool for temporal equilibrium analysis, [https://github.com/
1191 eve-mas/eve-parity](https://github.com/eve-mas/eve-parity), accessed: 09-09-2019.
- 1192 [10] W. van der Hoek, A. Lomuscio, M. Wooldridge, On the complexity of practical
1193 ATL model checking, in: AAMAS, ACM, 2006, pp. 201–208.
- 1194 [11] A. Pnueli, The temporal logic of programs, in: FOCS, IEEE, 1977, pp. 46–57.
- 1195 [12] F. Mogavero, A. Murano, G. Perelli, M. Y. Vardi, Reasoning about strategies: On
1196 the model-checking problem, ACM Transactions on Computational Logic 15 (4)
1197 (2014) 34:1–34:47.
- 1198 [13] J. Gutierrez, P. Harrenstein, M. Wooldridge, Reasoning about equilibria in game-
1199 like concurrent systems, Annals of Pure Applied Logic 168 (2) (2017) 373–403.
- 1200 [14] D. Fisman, O. Kupferman, Y. Lustig, Rational synthesis, in: TACAS, Vol. 6015 of
1201 LNCS, Springer, 2010, pp. 190–204.

- 1202 [15] J. Gutierrez, P. Harrenstein, M. Wooldridge, Iterated Boolean games, *Information*
1203 *and Computation* 242 (2015) 53–79.
- 1204 [16] K. Chatterjee, T. A. Henzinger, N. Piterman, Strategy logic, *Information and*
1205 *Computation* 208 (6) (2010) 677–693.
- 1206 [17] J. Gutierrez, G. Perelli, M. Wooldridge, Imperfect information in reactive mod-
1207 ules games, *Information and Computation* 261 (Part) (2018) 650–675.
- 1208 [18] C. Löding, Basics on tree automata, in: *Modern Applications of Automata The-*
1209 *ory*, Indian Institute of Science, Bangalore, India, 2012, pp. 79–110.
- 1210 [19] E. A. Emerson, C. S. Jutla, Tree automata, mu-calculus and determinacy, in:
1211 *FOCS, IEEE*, 1991, pp. 368–377.
- 1212 [20] M. Jurdzinski, Deciding the winner in parity games is in $UP \cap co-up$, *Information*
1213 *Processing Letters* 68 (3) (1998) 119–124.
- 1214 [21] C. S. Calude, S. Jain, B. Khossainov, W. Li, F. Stephan, Deciding parity games in
1215 quasipolynomial time, in: *STOC, ACM*, 2017, pp. 252–263.
- 1216 [22] O. Kupferman, Automata theory and model checking, in: *Handbook of Model*
1217 *Checking*, Springer International Publishing, 2018, pp. 107–151. doi:10.
1218 1007/978-3-319-10575-8\4.
- 1219 [23] J. Gutierrez, M. Najib, G. Perelli, M. Wooldridge, Eve: A tool for temporal equilib-
1220 rium analysis, in: *ATVA, Vol 11138 of LNCS*, Springer, Cham, 2018, pp. 551–557.
- 1221 [24] R. Alur, T. A. Henzinger, Reactive modules, *Formal Methods in System Design*
1222 15 (1) (1999) 7–48.
- 1223 [25] R. Alur, T. A. Henzinger, F. Y. C. Mang, S. Qadeer, S. K. Rajamani, S. Tasiran,
1224 *MOCHA: modularity in model checking*, in: *CAV, Vol. 1427 of LNCS*, Springer,
1225 1998, pp. 521–525.
- 1226 [26] M. Z. Kwiatkowska, G. Norman, D. Parker, PRISM: probabilistic model checking
1227 for performance and reliability analysis, *SIGMETRICS Performance Evaluation*
1228 *Review* 36 (4) (2009) 40–45.

- 1229 [27] R. I. Brafman, C. Domshlak, From one to many: Planning for loosely coupled
1230 multi-agent systems, in: Proceedings of the Eighteenth International Confer-
1231 ence on Automated Planning and Scheduling, ICAPS 2008, Sydney, Australia,
1232 September 14-18, 2008, 2008, pp. 28–35.
- 1233 [28] J. Gutierrez, P. Harrenstein, G. Perelli, M. Wooldridge, Nash equilibrium and
1234 bisimulation invariance, in: CONCUR, Vol. 85 of LIPIcs, Schloss Dagstuhl, 2017,
1235 pp. 17:1–17:16.
- 1236 [29] J. Gutierrez, P. Harrenstein, G. Perelli, M. J. Wooldridge, Nash equilibrium and
1237 bisimulation invariance, Logical Methods in Computer Science 15 (3).
- 1238 [30] R. Milner, Communication and Concurrency, Prentice Hall, 1989.
- 1239 [31] M. Hennessy, R. Milner, Algebraic laws for nondeterminism and concurrency,
1240 Journal of the ACM 32 (1) (1985) 137–161.
- 1241 [32] A. P. Sistla, M. Y. Vardi, P. Wolper, The complementation problem for Büchi au-
1242 tomata with applications to temporal logic, Theoretical Computer Science 49
1243 (1987) 217–237.
- 1244 [33] N. Piterman, From nondeterministic Büchi and Streett automata to deterministic
1245 parity automata, Logical Methods in Computer Science 3 (3) (2007) 1–21.
- 1246 [34] J. Gutierrez, P. Harrenstein, M. Wooldridge, Expressiveness and complexity re-
1247 sults for strategic reasoning, in: CONCUR, Vol. 42 of LIPIcs, Schloss Dagstuhl,
1248 2015, pp. 268–282.
- 1249 [35] D. Perrin, J. Pin, Infinite Words., Pure and Applied Mathematics., Elsevier, 2004.
- 1250 [36] P. Bouyer, R. Brenguier, N. Markey, M. Ummels, Pure Nash equilibria in con-
1251 current deterministic games, Logical Methods in Computer Science 11 (2) (2015)
1252 1–72.
- 1253 [37] W. Zielonka, Infinite games on finitely coloured graphs with applications to au-
1254 tomata on infinite trees, Theoretical Computer Science 200 (1-2) (1998) 135–183.

- 1255 [38] T. Gao, J. Gutierrez, M. Wooldridge, Iterated Boolean games for rational verifi-
1256 cation, in: AAMAS, ACM, 2017, pp. 705–713.
- 1257 [39] R. Brenguier, PRALINE: A tool for computing Nash equilibria in concurrent
1258 games, in: CAV, Vol. 8044 of LNCS, Springer, 2013, pp. 890–895.
- 1259 [40] P. Cermák, A. Lomuscio, F. Mogavero, A. Murano, MCMAS-SLK: A model
1260 checker for the verification of strategy logic specifications, in: CAV, Vol. 8559
1261 of LNCS, Springer, 2014, pp. 525–532.
- 1262 [41] P. Cermák, A. Lomuscio, F. Mogavero, A. Murano, Practical Verification of
1263 Multi-Agent Systems Against SLK Specifications, *Information and Computation*
1264 261 (Part) (2018) 588–614.
- 1265 [42] A. Lomuscio, H. Qu, F. Raimondi, MCMAS: an open-source model checker for
1266 the verification of multi-agent systems, *STTT* 19 (1) (2017) 9–30.
- 1267 [43] P. Gastin, D. Oddoux, Fast LTL to Büchi automata translation, in: CAV, Springer,
1268 2001, pp. 53–65.
- 1269 [44] LTL 2 BA: fast translation from LTL formulae to Büchi automata, [http://](http://www.lsv.fr/~gastin/ltl2ba/)
1270 www.lsv.fr/~gastin/ltl2ba/, accessed: 09-09-2019.
- 1271 [45] O. Friedmann, M. Lange, The pgsolver collection of parity game solvers – version
1272 3 (2010).
- 1273 [46] PGSolver, <https://github.com/tcsprojects/pgsolver>, ac-
1274 cessed: 09-09-2019.
- 1275 [47] R. Alur, T. A. Henzinger, O. Kupferman, Alternating-time temporal logic, *Journal*
1276 *of the ACM* 49 (5) (2002) 672–713.
- 1277 [48] R. Ladin, B. Liskov, L. Shriru, S. Ghemawat, Providing high availability using lazy
1278 replication, *ACM Transactions on Computer Systems* 10 (4) (1992) 360–391.
- 1279 [49] M. J. Fischer, A. Michael, Sacrificing serializability to attain high availability of
1280 data in an unreliable network, in: PODS, ACM, New York, NY, USA, 1982, pp.
1281 70–75.

- 1282 [50] G. T. Wuu, A. J. Bernstein, Efficient solutions to the replicated log and dictionary
1283 problems, in: PODC, ACM, New York, NY, USA, 1984, pp. 233–242.
- 1284 [51] D. K. Gifford, Weighted voting for replicated data, in: SOSp, ACM, New York,
1285 NY, USA, 1979, pp. 150–162.
- 1286 [52] J. Gutierrez, M. Najib, G. Perelli, M. J. Wooldridge, On computational tractability
1287 for rational verification, in: IJCAI, ijcai.org, 2019, pp. 329–335.
- 1288 [53] E. A. Emerson, Temporal and modal logic, in: Handbook of Theoretical Com-
1289 puter Science, Volume B: Formal Models and Semantics (B), MIT Press, Cam-
1290 bridge, MA, USA, 1990, pp. 995–1072.
- 1291 [54] O. Kupferman, G. Perelli, M. Y. Vardi, Synthesis with rational environments, An-
1292 nals of Mathematics and Artificial Intelligence 78 (1) (2016) 3–20.
- 1293 [55] B. Finkbeiner, S. Schewe, Coordination logic, in: CSL, Vol. 6247 of LNCS,
1294 Springer, 2010, pp. 305–319.
- 1295 [56] F. Laroussinie, N. Markey, Augmenting ATL with strategy contexts, Information
1296 and Computation 245 (2015) 98–123.
- 1297 [57] P. Cermák, A. Lomuscio, A. Murano, Verifying and synthesising multi-agent sys-
1298 tems against one-goal strategy logic specifications, in: AAI, AAAI Press, 2015,
1299 pp. 2038–2044.
- 1300 [58] P. Cermák, A. Lomuscio, F. Mogavero, A. Murano, Practical verification of
1301 multi-agent systems against SLK specifications, Information and Computation
1302 261 (Part) (2018) 588–614.
- 1303 [59] M. Kwiatkowska, D. Parker, C. Wiltsche, Prism-games 2.0: A tool for multi-
1304 objective strategy synthesis for stochastic games, in: TACAS, Vol. 9636 of LNCS,
1305 Springer, 2016, pp. 560–566.
- 1306 [60] A. Toumi, J. Gutierrez, M. Wooldridge, A tool for the automated verification of
1307 Nash equilibria in concurrent games, in: ICTAC, Vol. 9399 of LNCS, Springer,
1308 2015, pp. 583–594.

- 1309 [61] A. David, K. G. Larsen, A. Legay, M. Mikucionis, D. B. Poulsen, Uppaal SMC
1310 tutorial, *STTT* 17 (4) (2015) 397–415.
- 1311 [62] M. Kwiatkowska, G. Norman, D. Parker, G. Santos, Equilibria-based probabilistic
1312 model checking for concurrent stochastic games, in: *FM*, Vol. 11800 of LNCS,
1313 Springer, 2019, pp. 298–315.
- 1314 [63] A. David, P. G. Jensen, K. G. Larsen, M. Mikucionis, J. H. Taankvist, Uppaal strat-
1315 ego, in: *TACAS*, Vol. 9035 of LNCS, Springer, 2015, pp. 206–211.
- 1316 [64] P. E. Bulychev, A. David, K. G. Larsen, A. Legay, M. Mikucionis, Computing Nash
1317 equilibrium in wireless ad hoc networks: A simulation-based approach, in: *Pro-
1318 ceedings Second International Workshop on Interactions, Games and Protocols,
1319 IWIGP*, Vol. 78 of EPTCS, 2012, pp. 1–14.
- 1320 [65] R. Milner, *A Calculus of Communicating Systems*, Vol. 92 of LNCS, Springer,
1321 1980.
- 1322 [66] R. De Nicola, F. W. Vaandrager, Three logics for branching bisimulation, *Journal
1323 of the ACM* 42 (2) (1995) 458–487.
- 1324 [67] R. J. van Glabbeek, W. P. Weijland, Branching time and abstraction in bisimula-
1325 tion semantics, *Journal of the ACM* 43 (3) (1996) 555–600.
- 1326 [68] R. Alur, T. A. Henzinger, O. Kupferman, M. Y. Vardi, Alternating refinement re-
1327 lations, in: *CONCUR*, Vol. 1466 of LNCS, Springer, 1998, pp. 163–178.
- 1328 [69] J. van Benthem, Extensive games as process models, *Journal of Logic, Language
1329 and Information* 11 (3) (2002) 289–313.
- 1330 [70] O. Friedmann, M. Lange, Solving parity games in practice, in: *ATVA*, Vol. 5799
1331 of LNCS, Springer, 2009, pp. 182–196.
- 1332 [71] M. Ummels, D. Wojtczak, The complexity of Nash equilibria in stochastic multi-
1333 player games, *Logical Methods in Computer Science* 7 (3) (2011) 1–45.

- 1334 [72] T. Ågotnes, W. van der Hoek, M. Wooldridge, Reasoning about coalitional games,
1335 *Artificial Intelligence* 173 (1) (2009) 45–79.
- 1336 [73] F. Belardinelli, A. Lomuscio, A. Murano, S. Rubin, Verification of multi-agent
1337 systems with imperfect information and public actions, in: *AAMAS*, ACM, 2017,
1338 pp. 1268–1276.
- 1339 [74] B. Aminof, F. Mogavero, A. Murano, Synthesis of hierarchical systems, *Science*
1340 *of Computer Programming* 83 (2014) 56–79.
- 1341 [75] R. Berthon, B. Maubert, A. Murano, Decidability results for ATL^* with imperfect
1342 information and perfect recall, in: *AAMAS*, ACM, 2017, pp. 1250–1258.
- 1343 [76] A. Herzig, E. Lorini, F. Maffre, F. Schwarzenrüber, Epistemic Boolean games
1344 based on a logic of visibility and control, in: *IJCAI*, IJCAI/AAAI Press, 2016, pp.
1345 1116–1122.
- 1346 [77] F. Belardinelli, A. Lomuscio, Quantified epistemic logics for reasoning about
1347 knowledge in multi-agent systems, *Artificial Intelligence* 173 (9-10) (2009) 982–
1348 1013.
- 1349 [78] J. Gutierrez, A. Murano, G. Perelli, S. Rubin, M. J. Wooldridge, Nash equilibria in
1350 concurrent games with lexicographic preferences, in: *IJCAI*, ijcai.org, 2017, pp.
1351 1067–1073.
- 1352 [79] S. Almagor, O. Kupferman, G. Perelli, Synthesis of controllable nash equilibria
1353 in quantitative objective game, in: *Proceedings of the Twenty-Seventh Interna-*
1354 *tional Joint Conference on Artificial Intelligence, IJCAI 2018*, July 13-19, 2018,
1355 Stockholm, Sweden, 2018, pp. 35–41.