The Complexity of Pure Strategy Relevant Equilibria in Concurrent Games

Purandar Bhaduri
Indian Institute of Technology Guwahati, India
pbhaduri@iitg.ac.in

We study rational synthesis problems for concurrent games with ω-regular objectives. Our model of rationality considers only pure strategy Nash equilibria that satisfy either a social welfare or Pareto optimality condition with respect to an ω-regular objective for each agent. This extends earlier work on equilibria in concurrent games, without consideration about their quality. Our results show that the existence of Nash equilibria satisfying social welfare conditions can be computed as efficiently as the constrained Nash equilibrium existence problem. On the other hand, the existence of Nash equilibria satisfying the Pareto optimality condition possibly involves a higher upper bound, except in the case of Büchi and Muller games, for which all three problems are in the classes P and PSPACE-complete, respectively.

This work has been supported by the MATRICS grant MTR/2021/000003 from SERB India, titled "Games and Controller Synthesis".

1 Introduction

Infinite games on finite graphs play a fundamental role in the automated synthesis of reactive systems from their specification [6, 2]. The goal of reactive synthesis is to design a system that meets its specification in all possible environments. This problem can be modelled as a zero-sum game between the system and the environment, where a winning strategy for the system yields the design of a correct-by-construction controller.

In many situations involving autonomous agents, such as robots, drones, and autonomous vehicles, a purely adversarial view of the environment of a system is not appropriate. Instead, each agent should be viewed as trying to satisfy her own objective rather than preventing the other agents from meeting theirs. This gives rise to the notion of *rational synthesis* [7, 11], where we restrict our attention to agent behaviours that arise from game-theoretic equilibria. The most widely studied equilibrium is the Nash equilibrium (NE) [12], a strategy profile for which no agent has the incentive to unilaterally deviate from her strategy.

In this paper we investigate the complexity of computing certain desirable Nash equilibria in concurrent games played on finite graphs. We call these equilibria *relevant* after [4]. Our games have multiple agents (or players), each with an ω-regular objective. The games are concurrent, where the agents make their moves simultaneously. Each agent tries to fulfil her own objective while being subject to the decisions of the other agents. Such games provide a common framework for modelling systems with multiple, distributed, and autonomous agents.

It is well known that for such games pure strategy Nash equilibria may not exist and may not be unique when they do. Some equilibria may not be optimal, where only a few or none of the agents meet their objectives. We focus on two measures of the quality of equilibria – social welfare and Pareto optimality.

Submitted to:
GANDALF 2025

Objective	Constrained NE Existence [3]	Social Welfare (SWDP)	Pareto Optimality (PODP)
Reachability	NP-c	NP-c	NP-h, P^{NP}
Safety	NP-c	NP-c	NP-h, P^{NP}
Büchi	P-c	P	P
coBüchi	NP-c	NP-c	NP-h, P^{NP}
Parity	$\mathsf{P}^{\mathsf{NP}}_{\parallel}$-c	NP-h, $\mathsf{P}^{\mathsf{NP}}_{\parallel}$	P^{NP}
Muller	PSPACE-c	PSPACE-c	PSPACE-c

Table 1: Summary of complexities for relevant equilibria

Intuitively, the social welfare criterion considers the number of agents who meet their objectives in a given equilibrium. Pareto optimal equilibria, on the other hand, consider equilibria which are maximal in the sense that no additional agent can meet her objective in any behaviour of the system, whether it is the result of an equilibrium or not.

The specific problems which we are interested in are (i) the Constrained NE Existence Problem, asking whether there is an equilibrium satisfying a given lower and upper threshold for the payoff of each agent, (ii) the Social Welfare Decision Problem (SWDP), asking whether there is an equilibrium where the number of agents meeting their objectives is above a given threshold, and (iii) the Pareto Optimal Decision Problem (PODP), where the problem is to decide whether there is an equilibrium such that no other strategy profile, whether an equilibrium or not, results in a strict superset of agents meeting their objectives. The complexity of the Constrained NE Existence Problem for different ω-regular objectives was extensively studied by Bouyer *et al.* [3]. We mention them here as a point of reference and also because many of our results depend on them.

A summary of our results for SWDP and PODP is shown in Table 1. The results show that the existence of NE satisfying a social welfare condition can be decided as efficiently as the constrained NE existence problem. However, a Pareto optimality condition seems to entail a higher upper bound, except for Büchi and Muller objectives, for which all three problems are in the classes P and PSPACE-complete, respectively.

1.1 Related Work

Concurrent game structures were introduced by Alur *et al.* [1] for modelling the behaviour of open systems containing both system and environment components, called players or agents. Every state transition in such a game is determined by a choice of move by each player and models their synchronous composition.

The seminal work related to problems on Nash equilibria for concurrent games is by Bouyer *et al.* [3]. This paper studies Nash equilibria for pure-strategy multiplayer concurrent deterministic games for various ω-regular objectives and their generalisations. It gives comprehensive results for the Value Problem, the NE Existence Problem and the Constrained NE Existence Problem for all these objectives. The present work is an extension where we also study the existence of desirable equilibria, namely the existence of NE with lower bounds on social welfare and those satisfying Pareto optimality.

The term 'relevant equilibria' was introduced by Brihaye *et al.* [4] to refer to desirable equilibria such as those with high social welfare and those that are Pareto optimal. This paper was in the context of turn-based multiplayer games with quantitative reachability objectives. In contrast, our work is about concurrent qualitative games with the ω-regular objectives specified above.

Another line of work that deals with infinite concurrent games on finite graphs with both qualitative and quantitative objectives is by Gutierrez *et al.* [9]. In contrast to our work, their focus is on LTL and GR(1) objectives, a setting where the input sizes can be exponentially more succinct and the expressive power somewhat restricted.

Condurache *et al.* [5] have also investigated rational synthesis problems for concurrent games for the same objectives that we consider in this paper. However, their interest lies in *non-cooperative* rational synthesis, where the aim is to synthesise a strategy for agent 0 such that with respect to *all* NE involving the rest of the agents, agent 0's payoff is always 1. Computationally, this is a harder problem than the *cooperative* rational synthesis problem in our work, where we are interested in the existence of *at least one* NE satisfying some constraints. Moreover, [5] does not consider the quality of NEs, in contrast to the present work.

2 Preliminaries

2.1 Notation

$\mathbb{B}$ and $\mathbb{N}$ denote the sets $\{0,1\}$ of Boolean values and of natural numbers, respectively. A *word* over a finite alphabet Σ is a finite sequence of symbols from Σ. Σ^* denotes the set of all words, including the empty word ε. Σ^ω denotes the set of all *infinite words*, *i.e.*, infinite sequences over Σ. For $m,n \in \mathbb{N}$ with $m \leq n$ and a (finite or infinite) word α, we denote by $\alpha[j]$ the $j+1$-th letter of α and by $\alpha[m,n]$ the finite word $\alpha[m]\alpha[m+1]\dots\alpha[n]$.

2.2 Concurrent Game Structures

Definition 2.1. *A* concurrent game structure *(CGS) is a tuple* $\mathcal{G} = (\mathsf{St},\mathsf{Agt},\mathsf{Act},\mathsf{avb},\mathsf{tr})$ *where* St *is a finite non-empty set of states,* $\mathsf{Agt} = \{1,\dots,n\}$ *is the set of agents (or players),* Act *is a finite set of actions, the map* $\mathsf{avb} : \mathsf{St}\times\mathsf{Agt} \to 2^{\mathsf{Act}} \setminus \{\emptyset\}$ *indicates the actions available in a given state for a given agent, and* $\mathsf{tr} : \mathsf{St}\times\mathsf{Act}^n \to \mathsf{St}$ *is the transition function.*

An *action profile* or *move* $\bar{a} = (a_1,a_2,\dots,a_n) \in \mathsf{Act}^n$ is just an n-tuple of actions. Here, a_i is the action taken by agent i. We often write $\bar{a}(i)$ for a_i. We say $\bar{a}$ is *legal* at a state s if $a_i \in \mathsf{avb}(s,i)$ for all $i \in \mathsf{Agt}$. We call a CGS *turn-based* if, for each state, the set of available moves is a singleton for all but at most one player; such a player is said to *own* the state.[1]

Definition 2.2. *A* play *in the CGS* $\mathcal{G}$ *is a sequence of states* $\rho = s_0s_1s_2\dots$ *such that* $s_{j+1} = \mathsf{tr}(s_j,\bar{a}_j)$ *for some legal move* $\bar{a}_j$ *in state* s_j*, for all* $j \geq 0$.

Let $\mathsf{Plays}(\mathcal{G})$ be the set of all plays in the CGS $\mathcal{G}$. A *history* h is any finite prefix of a play. Let $\mathsf{Hist}(\mathcal{G})$ be the set of all histories in $\mathcal{G}$. We often drop $\mathcal{G}$ when referring to plays and histories when it is clear from the context. The last element of a history h is denoted by $\mathsf{last}(h)$.

Definition 2.3. *A strategy for agent i is a map* $\sigma_i : \mathsf{Hist} \to \mathsf{Act}$ *such that* $\sigma_i(h) \in \mathsf{avb}(\mathsf{last}(h),i)$*. A* strategy profile *is a tuple* $\bar{\sigma} = \langle\sigma_1,\dots,\sigma_n\rangle$ *of strategies, one for each agent.*

By convention, $\bar{\sigma}_{-i}$ is the tuple of strategies excluding that of agent i and $\langle\bar{\sigma}_{-i},\sigma'_i\rangle$ is obtained from the profile $\bar{\sigma}$ by substituting agent i's strategy σ_i by σ'_i. Note that our strategies are *pure*, *i.e.*, they do not involve any randomisation.

[1] We prefer the term player to agent when referring to turn-based games.

We say that a play ρ is *compatible* with a strategy σ_i of agent i if for every prefix $\rho[0,k]$ of ρ with $k \geq 0$ we have $\sigma_i(\rho[0,k]) = a_{i_k}$ and $\mathsf{tr}(\rho[k], \langle a_{-i_k}, a_{i_k}\rangle) = \rho[k+1]$, for some action profile a_{-i_k} of the other agents making $\langle a_{-i_k}, a_{i_k}\rangle$ legal in $\rho[k]$. We can define compatibility between a history and an agent i strategy in a similar way. For a coalition $P \subseteq \mathsf{Agt}$, and a tuple σ_P of strategies for the agents in P, we write $\mathsf{Out}_{\mathcal{G}}(\sigma_P)$ for the set of plays (called *outcomes*) in $\mathcal{G}$ that are compatible with strategy σ_i for every $i \in P$. Note that a strategy profile $\bar{\sigma}$ and an initial state s uniquely define a play $\mathsf{Out}(s, \bar{\sigma})$, referred to as its *outcome*.

Remark. *We assume that the transition function* tr *is represented explicitly as a table when* $\mathcal{G}$ *is an input to an algorithm. Its size* $|\mathsf{tr}|$ *is* $\sum_{s\in\mathsf{St}} \prod_{i\in\mathsf{Agt}} \mathsf{avb}(s,i) \cdot \lceil \log(|\mathsf{St}|) \rceil$ *and this can be exponential in the number* n *of agents.*

2.3 Concurrent Games and Solution Concepts

Omega-regular Games A concurrent (or multiplayer) game is a pair $\langle \mathcal{G}, (\mathsf{Obj}_i)_{i\in\mathsf{Agt}}\rangle$ where $\mathcal{G}$ is a CGS and $\mathsf{Obj}_i \subseteq \mathsf{St}^\omega$ is the *objective* for agent i. Thus, an objective is a set of infinite sequences of states in $\mathcal{G}$. For us, an objective can be any one of safety, reachability, Büchi, coBüchi, parity or Muller, defined below. For a play ρ in a CGS $\mathcal{G}$ we write $occ(\rho)$ for the states that occur in ρ and $inf(\rho)$ for the states that occur infinitely often in ρ, *i.e.*, $occ(\rho) = \{s \in \mathsf{St} \mid \exists j \geq 0.\ s = \rho[j]\}$ and $inf(\rho) = \{s \in \mathsf{St} \mid \forall j \geq 0.\ \exists k \geq j.\ s = \rho[k]\}$. Then, we consider the following objectives:

1. *Reachability:* Given a set $F \subseteq \mathsf{St}$ of *target states*, $\textsc{Reach}(F) = \{\rho \in \mathsf{St}^\omega \mid occ(\rho) \cap F \neq \emptyset\}$;
2. *Safety:* Given a set $F \subseteq \mathsf{St}$ of *unsafe states*, $\textsc{Safe}(F) = \{\rho \in \mathsf{St}^\omega \mid occ(\rho) \cap F = \emptyset\}$;
3. *Büchi:* Given a set $F \subseteq \mathsf{St}$ of *accept states*, $\textsc{Büchi}(F) = \{\rho \in \mathsf{St}^\omega \mid inf(\rho) \cap F \neq \emptyset\}$;
4. *coBüchi:* Given a set $F \subseteq \mathsf{St}$ of *reject states*, $\textsc{coBüchi}(F) = \{\rho \in \mathsf{St}^\omega \mid inf(\rho) \cap F = \emptyset\}$;
5. *Parity:* For a given *priority function* $p : \mathsf{St} \to \mathbb{N}$, $\textsc{Parity}(p) = \{\rho \in \mathsf{St}^\omega \mid \min\{p(s) \mid s \in inf(\rho)\}$ is even$\}$;
6. *Muller:* For a given finite set C of *colours*, a colouring function $c : \mathsf{St} \to C$ and a set $\mathcal{F} \subseteq 2^C$, $\textsc{Muller}(\varphi) = \{\rho \in \mathsf{St}^\omega \mid inf(c(\rho)) \in \mathcal{F}\}$. Here $c(\rho)$ is the infinite sequence of colours of the states in the sequence ρ and $inf(c(\rho))$ is the set of colours appearing infinitely often in the sequence $c(\rho)$.

For a given play ρ, the *payoff* of ρ, denoted $\mathsf{Payoff}(\rho)$, is given by the tuple $\langle \mathsf{Payoff}_i(\rho)\rangle_{i\in\mathsf{Agt}}$. Here, $\mathsf{Payoff}_i(\rho) \in \{0,1\}$, the payoff of agent i, is defined by $\mathsf{Payoff}_i(\rho) = 1 \Leftrightarrow \rho \in \mathsf{Obj}_i$. We say agent i wins the play ρ if her payoff is 1 and she loses ρ otherwise. For a given state s in $\mathcal{G}$, we write $\mathsf{Payoff}(s, \bar{\sigma})$ (respectively, $\mathsf{Payoff}_i(s, \bar{\sigma})$) for the payoff (respectively, payoff of agent i) for the unique play ρ that is the outcome of the strategy profile $\bar{\sigma}$ starting from state s.

A two-player concurrent game is *zero-sum* if for any play ρ, Player 1 wins ρ if, and only if, Player 2 loses it. This is a purely adversarial setting, where one player loses if the other wins and there are no ties. In a zero-sum game, we are interested in finding a winning strategy for a player from a given state if it exists. Such a strategy allows the player to win no matter how the other player moves. In non-zero-sum games, where each player has her own objective and is not necessarily trying to play spoilsport, winning strategies are too restrictive and seldom exist. Instead, the notion of an equilibrium, a strategy profile that is best possible for each player given the strategies of the other players in the profile, is the key concept. The most celebrated equilibrium in the literature is that of Nash equilibrium, defined below.

Nash Equilibrium and Relevant Equilibria The solution concept for games we consider in this paper is the Nash equilibrium [12], a strategy profile in which no agent has the incentive to unilaterally change her strategy.

Definition 2.4. *Let $\mathcal{G}$ be a concurrent game and let s be a state of $\mathcal{G}$. A strategy profile $\bar{\sigma}$ is a* Nash equilibrium *(NE) of $\mathcal{G}$ from s if for every agent i and every strategy σ'_i of agent i, it is the case that* $\mathsf{Payoff}_i(s,\bar{\sigma}) \geq \mathsf{Payoff}_i(s,\langle\bar{\sigma}_{-i},\sigma'_i\rangle)$.

In the kind of games we consider, the payoffs of agents for a given play ρ only depend on the set of states that are visited and the set of states that are visited infinitely often in ρ. For such games, the outcomes of Nash equilibria can be taken to be ultimately periodic sequences as shown in [3] (Proposition 3.1). We restate the result here.

Proposition 2.1. *Suppose $\mathcal{G}$ is a concurrent game where for any pair ρ,ρ' of plays, $occ(\rho) = occ(\rho')$ and* $\inf(\rho) = \inf(\rho')$ *imply* $\mathsf{Payoff}(\rho) = \mathsf{Payoff}(\rho')$*. If ρ is an outcome of an NE in $\mathcal{G}$ then there is an NE with outcome ρ' of the form $\alpha_1.\alpha_2^\omega$ such that* $\mathsf{Payoff}(\rho) = \mathsf{Payoff}(\rho')$*, where the lengths $|\alpha_1|$ and $|\alpha_2|$ have an upper bound of* $|\mathsf{St}|^2$.

In concurrent games, pure strategy Nash equilibria may not exist, or multiple equilibria may exist. We identify the equilibria that are desirable and call them *relevant* after [4]. For example, an equilibrium in which the number of players who meet their objectives is maximum among all equilibria (*i.e.*, one that maximises the *social welfare* defined below) would be considered a relevant one.

The relevant equilibria we consider are based on social welfare and Pareto optimality. Given a play ρ in a concurrent game $\langle\mathcal{G},(\mathsf{Obj})_i\rangle$, we denote by $\mathsf{Win}(\rho)$ the set of agents who meet their objective along ρ, *i.e.*, $\mathsf{Win}(\rho) = \{i \in \mathsf{Agt} \mid \rho \in \mathsf{Obj}_i\}$. The *social welfare* $\mathsf{sw}(\rho)$ of ρ is $|\mathsf{Win}(\rho)|$, *i.e.*, the number of agents who meet their objectives in ρ. For a given state s in $\mathcal{G}$ and a strategy profile $\bar{\sigma}$, we write $\mathsf{sw}(s,\bar{\sigma})$ for $\mathsf{sw}(\rho)$ where $\rho = \mathsf{Out}(s,\bar{\sigma})$.

To define Pareto optimality in a game, let $P = \{\langle w_i(\rho)\rangle_{i\in\mathsf{Agt}} \mid \rho \in \mathsf{Plays}(\mathcal{G})\}$, where $w_i(\rho) = 1$ if agent i wins ρ and 0 otherwise, *i.e.*, the set P is the set of winner profiles for all plays in $\mathcal{G}$. Then a winner profile $p \in P$ is *Pareto optimal* if it is maximal in P with respect to the componentwise ordering $\leq_P$ on P where $0 < 1$ in each component. This means no other agent can be added to the set of winners in p, *i.e.*, p represents a maximal set of winners along any play.

2.4 Rational Synthesis Problems

The rational synthesis problem generalises the synthesis problem to multiagent systems. The aim is to synthesise a game-theoretic equilibrium, a Nash equilibrium in our case, that satisfies additional desirable properties. The problems we study are detailed below.

For completeness, we start by referring to the *Constrained NE Problem* solved by Bouyer *et al.* [3]. This problem asks whether there is an NE in a given game whose payoff profile satisfies both an upper and a lower threshold.[2]

Problem 1 (Constrained NE Existence Problem). *Given a game $\mathcal{G}$, a state s in $\mathcal{G}$ and threshold tuples $\mathbf{v},\mathbf{u} \in \mathbb{B}^n$, decide whether there exists an NE $\bar{\sigma}$ such that $\mathbf{v} \leq (\mathsf{Payoff}(s,\bar{\sigma})) \leq \mathbf{u}$, where the ordering is defined componentwise.*

[2] The original formulation in [3] was stated in terms of a preference relation $\lesssim_A$ for each agent A as follows: Given two plays ρ_A^ℓ and ρ_A^u for each agent A, is there a Nash equilibrium $\bar{\sigma}$ which satisfies $\rho_A^\ell \lesssim_A \mathsf{Out}(\bar{\sigma}) \lesssim_A \rho_A^u$ for all $A \in \mathsf{Agt}$? Since for ω-regular objectives a threshold play ρ is encoded by the pair $(occ(\rho), inf(\rho))$ in [3], the way we state the problem is polynomially equivalent.

Bouyer *et al.* showed that the problem is P-complete for Büchi objectives, NP-complete for safety, reachability and co-Büchi objectives, and PSPACE-complete for Muller objectives; the complexity class for parity objectives is $\mathsf{P}^{\mathsf{NP}}_{\parallel}$-complete, where $\mathsf{P}^{\mathsf{NP}}_{\parallel}$ is the class of problems that can be solved by a deterministic Turing machine in polynomial time with an access to an oracle for solving NP problems and such that the oracle can be queried only once with a set of queries.

We now come to the relevant equilibrium problems that we explore in this paper. First, we define the social welfare problem as placing a lower bound on the social welfare, *i.e.*, the number of winning players, of an equilibrium.

Problem 2 (Social Welfare Decision Problem (SWDP)). *Given a game $\mathcal{G}$, a state s in $\mathcal{G}$ and a threshold value $v \in \mathbb{N}$, decide whether* there exists *an NE $\bar{\sigma}$ such that* $\mathsf{sw}(s,\bar{\sigma}) \geq v$.

The second problem we consider is the Pareto optimality decision problem for rational synthesis.

Problem 3 (Pareto Optimal Decision Problem (PODP)). *Given a game $\mathcal{G}$ and a state s in $\mathcal{G}$, decide whether* there exists *an NE $\bar{\sigma}$ such that* $\mathsf{Payoff}(s,\bar{\sigma})$ *is Pareto optimal.*

Notice that the problems above are about the *existence* of NE satisfying some conditions. Hence, it is a *cooperative* setting in which we assume that all agents will cooperate when presented with an NE.

2.5 The Suspect Game

Many results from Bouyer *et al.* [3] that we use below rely on a key construction. The idea is based on a correspondence between Nash equilibria in a concurrent game $\mathcal{G}$ and winning strategies in a two-player zero-sum game $\mathcal{H}$ derived from $\mathcal{G}$, called the suspect game. The game $\mathcal{H}$ is played between Eve and Adam. Intuitively, Eve's goal is to prove that the sequence of moves proposed by her results from a Nash equilibrium in $\mathcal{G}$, while Adam's task is to foil her attempt by exhibiting that some agent has a profitable deviation from the strategy suggested by Eve. Here we recall the basic definitions and results from [3].

Given two states s and s' and a move $\bar{a}$ in a concurrent game $\mathcal{G}$, the set of *suspect agents* for (s,s') and $\bar{a}$ is the set

$$\mathsf{Susp}((s,s'),\bar{a}) = \{i \in \mathsf{Agt} \mid \exists a' \in \mathsf{avb}(s,i).\ \mathsf{tr}(s,\langle \bar{a}_{-i},a'\rangle) = s'\}.$$

Note that if $\mathsf{tr}(s,\bar{a}) = s'$ then $\mathsf{Susp}((s,s'),\bar{a}) = \mathsf{Agt}$, *i.e.*, every agent is suspect if there is no deviation from the suggested move. For a play ρ and a strategy profile $\bar{\sigma}$, the set of suspect agents for ρ and $\bar{\sigma}$ is given by the set of suspect agents along each transition of ρ:

$$\mathsf{Susp}(\rho,\bar{\sigma}) = \{i \in \mathsf{Agt} \mid \forall j \in \mathbb{N}.\ i \in \mathsf{Susp}((\rho[j],\rho[j+1]),\bar{\sigma}(\rho[0,j])\}.$$

The idea is that agent i is a suspect for a pair (s,s') and move $\bar{a}$ if she can unilaterally deviate from her action a_i in $\bar{a}$ to trigger the transition (s,s'). It follows from the above definitions that agent i is in $\mathsf{Susp}(\rho,\bar{\sigma})$ if, and only if, there is a strategy σ' for agent i such that $\mathsf{Out}(\langle \bar{\sigma}_{-i},\sigma'\rangle) = \rho$.

For a fixed play π in a concurrent game $\mathcal{G}$, we build the *suspect game* $\mathcal{H}(\mathcal{G},\pi)$, a two-player turn-based zero-sum game between Eve and Adam as follows. The set of states of $\mathcal{H}(\mathcal{G},\pi)$ is the disjoint union of the states $V_\exists \subseteq \mathsf{St} \times 2^{\mathsf{Agt}}$ owned by Eve, and the set $V_\forall \subset \mathsf{St} \times 2^{\mathsf{Agt}} \times \mathsf{Act}^{\mathsf{Agt}}$ owned by Adam. The game proceeds in the following way: from a state (s,P) in $V_\exists$, Eve chooses a legal move $\bar{a}$ from s in $\mathcal{G}$, resulting in the new state $(s,P,\bar{a})$ in $V_\forall$. Adam then chooses a move $\bar{a}'$ that will actually apply in $\mathcal{G}$ leading to a state s' in St; the resulting state in $\mathcal{H}(\mathcal{G},\pi)$ is $(s',P \cap \mathsf{Susp}((s,s'),\bar{a}))$. In the special case when the state chosen by Adam is such that $s' = \mathsf{tr}(s,\bar{a})$, we say that Adam obeys Eve. In this case, the new state is given by (s',P).

We define the two projections $proj_1$ and $proj_2$ from $V_\exists$ to St and 2^{Agt}, respectively, by $proj_1(s,P) = s$ and $proj_2(s,P) = P$. These projections are extended to plays in $\mathcal{H}(\mathcal{G},\pi)$ in a natural way but only using Eve's states – for example, $proj_1((s_0,P_0)(s_0,P_0,\bar{a})(s_1,P_1))\cdots = s_0 s_1 \cdots$. For any play ρ in $\mathcal{H}(\mathcal{G},\pi)$, $proj_2(\rho)$, which is a sequence of sets of agents in $\mathcal{G}$, is non-increasing, and hence its limit $\lambda(\rho)$ is well defined. Note that if $\lambda(\rho) \neq \emptyset$ then $proj_1(\rho)$ is a play in $\mathcal{G}$. A play ρ is winning for Eve, if for all $i \in \lambda(\rho)$ the play π is as good as or better than $proj_1(\rho)$ for agent i in $\mathcal{G}$, *i.e.*., $\mathsf{Payoff}_i(\pi) \geq \mathsf{Payoff}_i(proj_1(\rho))$. The *winning region* $W(\mathcal{G},\pi)$ is the set of states of $\mathcal{H}(\mathcal{G},\pi)$ from which Eve has a winning strategy.

The correctness of the suspect game construction is captured by the following theorem from [3].

Theorem 2.2. *Let $\mathcal{G}$ be a concurrent game, s a state of $\mathcal{G}$ and π a play in $\mathcal{G}$. Then the following two conditions are equivalent.*

1. *There is an NE $\bar{\sigma}$ from s in $\mathcal{G}$ whose outcome is π.*
2. *There is a play ρ from (s,Agt) in $\mathcal{H}(\mathcal{G},\pi)$ satisfying*
   (a) *Adam obeys Eve along ρ,*
   (b) *$proj_1(\rho) = \pi$, and*
   (c) *for all $i \in \mathbb{N}$, there is a strategy $\sigma^i_\exists$ for Eve, for which any play in $\rho[0,i] \cdot \mathsf{Out}(\rho[i],\sigma^i_\exists)$ is winning for Eve.*

The following theorem from [3] asserts that not only is the suspect game construction correct, but it does not result in an exponential blow-up in size as well. Note that the infinite play π is specified by the pair of finite sets $(occ(\pi), inf(\pi))$.

Theorem 2.3. *Let $\mathcal{G}$ be a concurrent game and π a play in $\mathcal{G}$. The number of reachable states from $\mathsf{St} \times \mathsf{Agt}$ in $\mathcal{H}(\mathcal{G},\pi)$ is polynomial in the size of $\mathcal{G}$.*

Note that for two plays π and π' in a concurrent game $\mathcal{G}$ the suspect games $\mathcal{H}(\mathcal{G},\pi)$ and $\mathcal{H}(\mathcal{G},\pi')$ have identical CGSs and differ only in the winning conditions. Thus, the CGS $\mathcal{J}(\mathcal{G})$ of $\mathcal{H}(\mathcal{G},\pi)$ depends solely on $\mathcal{G}$ and is polynomial in size. Also, if we denote the set of losing agents in π by $\mathrm{Los}(\pi)$, the winning condition for Eve in $\mathcal{H}(\mathcal{G},\pi)$ can be stated as follows: for every $i \in \lambda(\rho) \cap \mathrm{Los}(\pi)$, agent i loses $proj_1(\rho)$ in $\mathcal{G}$. Thus, the winning condition depends only on $\mathrm{Los}(\pi)$ and not the exact sequence π. Henceforth we denote the suspect game by $\mathcal{H}(\mathcal{G},L)$, where $L \subseteq \mathsf{Agt}$ and Eve wins the play ρ if for every $i \in \lambda(\rho) \cap L$, agent i loses the play $proj_1(\rho)$ in $\mathcal{G}$.

3 Relevant Equilibria for Omega-regular Concurrent Games

3.1 Reachability Games

Social Welfare Problem We begin by showing that SWDP, the social welfare decision problem, is NP-complete for reachability objectives. The reduction from SAT is based on a construction by Bouyer *et al.* [3] for the constrained NE existence problem for turn-based games.

Lemma 3.1. SWDP *for reachability objectives is* NP*-complete.*

Proof. We reduce SAT to SWDP with reachability objectives. Let $\phi = C_1 \wedge \ldots \wedge C_m$ be an instance of SAT with $C_i = \ell_{i,1} \vee \ell_{i,2} \vee \ell_{i,3}$ over a set of variables $\{x_1,\ldots,x_n\}$. Then construct the $(m+1)$-player turn-based reachability game as shown in Fig. 1. In this game Player 0 owns all the states, shown as circles in the figure. The target set for Player 0 includes all the states, *i.e.*, Player 0 always wins. For Player i with $i \in \{1,\ldots,m\}$ the target set includes x_k if x_k appears in C_i, and $\neg x_k$ if $\neg x_k$ appears in C_i, *i.e.*, the target

Figure 1: Reduction from SAT to SWDP: Reachability

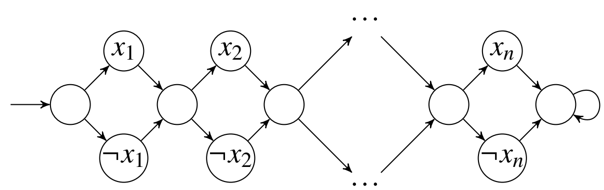

set is $\{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\}$. The threshold value v for the social welfare function in the SWDP instance is set to $m+1$, the number of players. Clearly, the game has an NE with a payoff of 1 for each player if, and only if, ϕ is satisfiable. If τ is a satisfying valuation for ϕ, the strategy for Player 0 is simple: between x_k and $\neg x_k$ choose x_k if $\tau(x_k) = 1$ and $\neg x_k$ otherwise. Conversely, if all the players win, a satisfying valuation τ can be similarly obtained from Player 0's strategy, by setting $\tau(x_k)$ to 1 if Player 0 chooses to move to x_k and to 0 otherwise.

We show that SWDP for reachability objectives is in NP by mirroring the Algorithm in Section 5.1.2 in [3] with a small modification for adapting the solution for the constrained NE existence problem to SWDP. The algorithm shown below makes essential use of the suspect game $\mathcal{H}(\mathcal{G}, L)$ and its reduction to the safety game $\mathcal{J}(\mathcal{G})$ with safety objective Ω_L. See [3] for the technical details.

1. Given a value $v \in [0,n]$, first guess a lasso-shaped play $\rho = \alpha_1 \cdot \alpha_2^{\omega}$ where $|\alpha_i|^2 \leq 2|\mathsf{St}|^2$ in $\mathcal{J}(\mathcal{G})$ such that Adam obeys Eve along ρ, and the play $\pi = proj_1(\rho)$ in $\mathcal{G}$ satisfies the constraint that at least v players are winning in it. This condition on the number of winning players is the only change from Section 5.1.2 in [3],
2. Then compute the set $W(\mathcal{G}, \mathrm{Los}(\pi))$ of the winning states for Eve in the suspect game $\mathcal{H}(\mathcal{G}, \mathrm{Los}(\pi))$, where $\mathrm{Los}(\pi)$ is the set of losing players along π.
3. Finally, check that ρ always stays in $W(\mathcal{G}, \mathrm{Los}(\pi))$.

We refer to [3] for the proof that this nondeterministic algorithm runs in polynomial time.

□

Pareto Optimal Decision Problem It is clear that the construction in the proof of Lemma 3.1 yields an NE that is Pareto-optimal if, and only if, the formula ϕ is satisfiable. Hence the Pareto Optimal Decision Problem for reachability objectives is NP-hard using the same reduction as in the lemma.

Lemma 3.2. PODP *for reachability objectives is* NP*-hard.*

For the upper bound, we show that PODP for reachability objectives is in the class $\mathsf{P}^{\mathsf{NP}} = \Delta_2^p$ in the polynomial hierarchy as follows. First, use binary search for the threshold value v in the range $[0,n]$ and the NP oracle for deciding SWDP in the proof of Lemma 3.5 above to determine the maximum value m of v for which the procedure returns yes. Then, to check if there is a run ρ where more than m players are winners, search for any strongly connected component (SCC) C reachable from the initial state s in the underlying CGS that satisfies the following condition: there are more than m sets in $F_1, \ldots, F_n$ with a nonempty intersection with C. Since this can be done in polynomial time using Tarjan's algorithm for finding all SCCs in a directed graph [13], the entire procedure runs in P^{NP} time. However, we leave the question of whether PODP for reachability objectives is P^{NP}-hard open.

3.2 Safety Games

Social Welfare Problem We show that SWDP is NP-complete for safety objectives by reduction from SAT. We use a modification of a construction by Bouyer *et al.* [3] for the value problem for ordered Büchi objectives with the counting preorder.

Lemma 3.3. SWDP *for safety objectives is* NP*-complete.*

Proof. We reduce SAT to SWDP with safety objectives for a turn-based game. Consider an instance $\phi = C_1 \wedge \ldots \wedge C_m$ of SAT with $C_j = \ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3}$ over a set of variables $\{x_1, \ldots, x_n\}$. We associate a $(2n+1)$-player turn-based game $\mathcal{G}_\phi$ with ϕ. The set of states is given by the union of

$$V_0 = \{s\} \text{ where } s \text{ is the initial state in the SWDP instance,}$$
$$V_k = \{x_k, \neg x_k\} \text{ for each } 1 \leq k \leq n, \text{ and}$$
$$V_{n+j} = \{\ell_{j,1}, \ell_{j,2}, \ell_{j,3}\} \text{ for each } 1 \leq j \leq m.$$

We add a transition from each state in V_i to each state in V_{i+1} for $0 \leq i \leq n+m$, assuming $V_{n+m+1} = V_0$. The game has $2n+1$ players $P_0, \ldots, P_{2n}$. P_0 owns all the states and her safety objective is given by the unsafe set $F_0 = \emptyset$, *i.e.*, she always wins. This implies that all strategy profiles are NE, as the other players have no choice at any state. For $1 \leq k \leq n$, P_{2k-1} and P_{2k} have safety objectives given by the unsafe states

$$F_{2k-1} = \{x_k\} \cup \{\ell_{j,p} \mid \ell_{j,p} \text{ is the literal } x_k\}$$
$$F_{2k} = \{\neg x_k\} \cup \{\ell_{j,p} \mid \ell_{j,p} \text{ is the literal } \neg x_k\}$$

At least n of these sets $F_1, \ldots, F_{2n}$ will be visited along any infinite play and thus at least n of these $2n$ players $P_1, \ldots, P_{2n}$ will always lose. We show that ϕ is satisfiable if, and only if, there exists an NE for which at most (and hence exactly) n of these $2n$ sets $F_1, \ldots, F_{2n}$ are visited, *i.e.*, at least $n+1$ players win. In other words, ϕ is satisfiable if, and only if, there is an NE σ in $\mathcal{G}_\phi$ with $\mathsf{sw}(\sigma) \geq n+1$.

Assume ϕ is satisfiable and let τ be a satisfying valuation. The strategy σ_0 for P_0 simply follows τ, *i.e.*, for states in V_{k-1} for $1 \leq k \leq n$, the strategy chooses x_k if $\tau(x_k) = \mathsf{true}$ and it chooses $\neg x_k$ otherwise. From a state in V_{n+k-1} for $1 \leq k \leq m$, it chooses one of the $\ell_{j,p}$ that evaluates to true under τ, say the one with the least index. As a result, the number of sets in $F_1, \ldots, F_n$ that are visited at least once is n. Since P_0 owns all the states and wins using σ_0 and the strategies of all the other players are empty, we have an NE σ where $n+1$ players win.

Conversely, pick a play in $\mathcal{G}$ resulting from a strategy σ_0 of P_0 such that at most (hence exactly) n of the sets $F_1, \ldots, F_{2n}$ are visited at least once. In particular, this play never visits one of x_k and $\neg x_k$ for any $1 \leq k \leq n$. Clearly, such a strategy is part of an NE σ (since all strategy profiles are) in which the number of winning players is $n+1$. One can define a truth valuation τ over $\{x_1, \ldots, x_n\}$ from the play – simply set $\tau(x_k)$ to true if x_k is visited at least once and to false otherwise. Also, any state of V_{n+j} with $1 \leq j \leq m$ that is visited at least once must correspond to a literal that is assigned the value true by τ, otherwise there would be more than n states among $F_1, \ldots, F_n$ visited at least once. Hence, each clause of ϕ evaluates to true under τ, and therefore ϕ is satisfiable.

To show that SWDP for safety objectives is in NP, we follow the Algorithm in Section 5.2.3 in [3] with a small modification, just as in the reachability case in Section 3.1. The algorithm reduces the suspect game $\mathcal{H}(\mathcal{G}, L)$ to the safety game $\mathcal{J}(\mathcal{G})$ with safety objective Ω_L. See [3] for the technical details.

1. Given a value $v \in [0, n]$, first guess a lasso-shaped play $\rho = \alpha_1 \cdot \alpha_2^\omega$ where $|\alpha_i|^2 \leq |\mathsf{St}|^2$ in $\mathcal{J}(\mathcal{G})$ such that Adam obeys Eve along ρ, and the play $\pi = proj_1(\rho)$ in $\mathcal{G}$ satisfies the constraint that at least v players are winning in it.

2. Then check that any deviation by Adam along ρ, say at position i leads to a state from which Eve has a strategy σ^i that ensures that any play in $\rho[0,i] \cdot \mathsf{Out}(\sigma^i)$ is winning.

We refer to [3] for the details of step 2 above and the proof that this nondeterministic algorithm runs in polynomial time.

□

Pareto Optimal Decision Problem It is clear that the construction in the proof of Lemma 3.3 yields an NE that is Pareto-optimal if, and only if, the formula ϕ is satisfiable, just as in the case of reachability.

Lemma 3.4. PODP *for safety objectives is* NP*-hard.*

For the upper bound, we show that PODP for safety objectives is in the class $\mathsf{P}^{\mathsf{NP}} = \Delta_2^p$ in the polynomial hierarchy using the same procedure as in the reachability case, except now we use the NP oracle for deciding SWDP for safety objectives. As in the case of reachability, we leave the question of whether PODP for safety objectives is P^{NP}-hard open.

3.3 Büchi Games

Social Welfare Problem We show that the social welfare problem for Büchi objectives can be solved in polynomial time by giving a polynomial time algorithm that invokes the procedure for Constrained NE Existence Problem from [3]. A polynomial-time algorithm for the latter was presented in [3].

Lemma 3.5. SWDP *for Büchi objectives is in* P*.*

Proof. The following is a polynomial-time algorithm for deciding whether there is a Nash equilibrium with v or more winners in a given concurrent game with Büchi objectives, starting from a state s.

1. Find all the reachable SCCs in the underlying digraph of $\mathcal{G}$ using Tarjan's algorithm [13]. Call the number of agents that meet their Büchi objectives in the SCC C the *rank of* C. Agent i meets her objective in the SCC C if $F_i \cap C \neq \emptyset$.
2. Sort the SCCs in non-increasing order according to their rank.
3. For each rank r starting from the highest down to v, check whether there is an NE with r winners using the algorithm for constrained NE Existence from [3] by setting both the lower and upper thresholds $\mathbf{v}$ and $\mathbf{u}$ to the winner profile of each SCC C of rank r one by one. If such an NE exists then return 'yes', else return 'no'.

□

The correctness of the algorithm depends on Proposition 2.1 on the lasso characterization of NEs in concurrent games.

Pareto Optimal Decision Problem

Lemma 3.6. PODP *for Büchi objectives is in* P*.*

Proof. A small modification in the algorithm in the proof of Lemma 3.5 gives a polynomial-time algorithm for deciding whether a Pareto-optimal NE exists. After sorting the SCCs in non-increasing order according to their rank, starting from the highest rank we check if any of the SCCs with the given rank is an NE. If a rank r is found for which all SCCs are non-NEs, then return 'no'. Otherwise, if an SCC corresponding to an NE is found, then return 'yes'. □

3.4 CoBüchi Games

Social Welfare Problem We show that SWDP is NP-complete for coBüchi objectives by reduction from SAT. The reduction is the same as in Section 3.2 for safety objectives, with the unsafe states now playing the role of coBüchi objectives.

Lemma 3.7. SWDP *for coBüchi objectives is* NP*-complete.*

Proof. We reduce SAT to SWDP with coBüchi objectives for a turn-based game using the same construction as in the proof of Lemma 3.3, with the unsafe states $F_1,\dots,F_n$ being now designated as the coBüchi objectives.

At least n of these sets $F_1,\dots,F_{2n}$ will be visited infinitely often along any infinite play and thus at least n of these $2n$ players $P_1,\dots,P_{2n}$ will always lose. We show that ϕ is satisfiable if, and only if, there exists an NE for which at most (and hence exactly) n of these $2n$ sets $F_1,\dots,F_{2n}$ are visited infinitely often, *i.e.*, at least $n+1$ players win. In other words, ϕ is satisfiable if, and only if, there is an NE σ in $\mathcal{G}_\phi$ with $\mathsf{sw}(\sigma) \geq n+1$.

Assume ϕ is satisfiable and let τ be a satisfying valuation. The strategy σ_0 for P_0 simply follows τ, *i.e.*, for states in V_{k-1} for $1 \leq k \leq n$, the strategy chooses x_k if $\tau(x_k) = \mathsf{true}$ and it chooses $\neg x_k$ otherwise. From a state in V_{n+k-1} for $1 \leq k \leq m$, it chooses one of the $\ell_{j,p}$ that evaluates to true under τ, say the one with the least index. This way the number of sets in $F_1,\dots,F_n$ that are visited infinitely often is n and the other sets are not visited at all. Since P_0 owns all the states and wins using σ_0 and the strategies of all the other players are empty, we have an NE σ where $n+1$ players win.

Conversely, pick a play in $\mathcal{G}$ resulting from a strategy σ_0 of P_0 such that at most (hence exactly) n of the sets $F_1,\dots,F_{2n}$ are visited infinitely often. In particular, this play always visits one of x_k and $\neg x_k$ finitely often for any $1 \leq k \leq n$. Clearly, such a strategy is part of an NE σ (since all strategy profiles are) in which the number of winning players is $n+1$. One can define a truth valuation τ over $\{x_1,\dots,x_n\}$ from the play – simply set $\tau(x_k)$ to true if x_k is visited infinitely often and to false otherwise. Also, any state of V_{n+j} with $1 \leq j \leq m$ that is visited infinitely often must correspond to a literal that is assigned the value true by τ, otherwise there would be more than n states among $F_1,\dots,F_n$ visited infinitely often. Hence each clause of ϕ evaluates to true under τ, and hence ϕ is satisfiable.

To show that SWDP for coBüchi objectives is in NP, we follow the Algorithm in Section 5.4.3 in [3] with a small modification as in the reachability case above. The algorithm also uses the suspect game $\mathcal{H}(\mathcal{G},L)$ and its reduction to the safety game $\mathcal{J}(\mathcal{G})$ with safety objective Ω_L. See [3] for the technical details.

1. Given a value $v \in [0,n]$, first guess a lasso-shaped play $\rho = \alpha_1 \cdot \alpha_2^\omega$ where $|\alpha_i|^2 \leq |\mathsf{St}|^2$ in $\mathcal{J}(\mathcal{G})$ such that Adam obeys Eve along ρ, and the play $\pi = proj_1(\rho)$ in $\mathcal{G}$ satisfies the constraint that at least v players are winning in it.
2. Then check that any deviation by Adam along ρ, say at position i leads to a state from which Eve has a strategy σ^i that ensures that any play in $\rho[0,i] \cdot \mathsf{Out}(\sigma^i)$ is winning.

We refer to [3] for the details of step 2 above and the proof that this nondeterministic algorithm runs in polynomial time.

□

Pareto Optimal Decision Problem As in the case of the other objectives considered above, the construction in the proof of Lemma 3.7 yields an NE which is Pareto-optimal if and only if the formula ϕ is satisfiable.

Lemma 3.8. PODP *for coBüchi objectives is* NP*-hard.*

For the upper bound, we show that PODP for coBüchi objectives is in the class $\mathsf{P}^{\mathsf{NP}} = \Delta_2^p$ in the polynomial hierarchy using the same procedure as in the safety case, except now we use the NP oracle for deciding SWDP for coBüchi objectives. As in the case of safety, we leave the question whether PODP for coBüchi objectives is P^{NP}-hard open.

3.5 Parity Games

Social Welfare Problem Bouyer *et al.* [3] showed that the Constrained NE Existence Problem is $\mathsf{P}_{\|}^{\mathsf{NP}}$-complete for parity objectives by reduction from ⊕SAT. Intuitively, $\mathsf{P}_{\|}^{\mathsf{NP}}$ is the class of all languages accepted by some deterministic polynomial time Turing machine M using an oracle for solving NP problems, such that on any input the machine M builds a set of queries to the oracle before making the queries just once. For a formal definition see [14]. In the ⊕SAT problem, given a finite set of instances of SAT, the goal is to decide whether the number of satisfiable instances is even. The problem is known to be $\mathsf{P}_{\|}^{\mathsf{NP}}$-complete [8].

We show the same upper bound for SWDP for parity objectives, *i.e.*, it is in $\mathsf{P}_{\|}^{\mathsf{NP}}$. The proof is essentially the same algorithm as in [3] – see Section 5.6.2 in the paper for Rabin objectives. We first translate the parity objectives to corresponding Rabin ones with half as many pairs as the number of priorities and then apply the algorithm in [3]. The only modification in the algorithm is identical to the ones for the reachability, safety and coBüchi objectives, namely step 1, where we check that the play $\pi = proj_1(\rho)$ in $\mathcal{G}$ satisfies the constraint that at least v players are winning in it, where v is the threshold input to SWDP.

Lemma 3.9. SWDP *for parity objectives is in* $\mathsf{P}_{\|}^{\mathsf{NP}}$.

However, for the lower bound we leave the question whether SWDP for parity objectives is $\mathsf{P}_{\|}^{\mathsf{NP}}$-hard as open. However, we can show the weaker result that the problem is NP-hard just by coding the coBüchi condition by a parity condition with two colours.

Pareto Optimal Decision Problem We show that PODP for parity objectives is in the class $\mathsf{P}^{\mathsf{P}^{\mathsf{NP}}} = \mathsf{P}^{\mathsf{NP}}$ in the polynomial hierarchy using the same procedure as in the coBüchi case, except now we use the $\mathsf{P}_{\|}^{\mathsf{NP}}$ oracle for deciding SWDP for parity objectives. Again, we leave the question of whether PODP for parity objectives is P^{NP}-hard open.

3.6 Muller Games

Social Welfare Problem We show that SWDP is PSPACE-complete for Muller objectives by a reduction from TQBF.

Lemma 3.10. SWDP *for Muller objectives is* PSPACE*-complete.*

Proof. We reduce TQBF to SWDP with Muller objectives for a three-player turn-based game by reusing the proof of PSPACE-hardness of deciding the winner of a zero-sum two-player Muller game by Hunter and Dawar [10]. First, note that the complementary objective of a zero-sum Muller game is also a Muller game, simply by changing the colouring function appropriately. Then given a TQBF formula φ we take the two-player Muller game $\mathcal{G}_\varphi$ in [10] and simply add one more player (Player 2) who has the same Muller objective as Player 0 but controls no vertex. In other words, for all plays ρ, Player 2 wins ρ iff Player 0 wins ρ iff Player 1 loses ρ. Thus, setting the threshold value v in SWDP to 2 and using the construction in [10] with the above modification reduces TQBF to SWDP with Muller objectives for three players.

The proof of membership of SWDP for Muller objectives in PSPACE is as follows. Given a value $v \in [0,n]$, first guess a set W of v winning players. This can clearly be done in polynomial space. Then use the procedure for checking membership in PSPACE for the corresponding Constrained NE Problem from [3] (see the third paragraph on page 25 of the paper) with appropriate lower and upper threshold tuples of bits, $\mathbf{v}$ and $\mathbf{u}$, respectively. Here, $\mathbf{v}$ contains 1's only for the v winners and 0's elsewhere, and $\mathbf{u}$ contains all 1's. □

Pareto Optimal Decision Problem As in the previous cases, the construction in the proof of Lemma 3.10 produces an NE that is Pareto-optimal for Muller objectives if, and only if, the formula ϕ is satisfiable. Membership in the class PSPACE follows using the same algorithm for all earlier cases, except using the PSPACE oracle to decide SWDP for Muller objectives.

Lemma 3.11. PODP *for Muller objectives is* PSPACE*-complete.*

4 Conclusion

In this work, we have extended the complexity results for rational synthesis problems for concurrent games to the case of relevant equilibria. We restrict ourselves to pure strategy Nash equilibria satisfying a social welfare or a Pareto optimality condition.

This work can be extended in many possible directions. One can consider solution concepts other than Nash equilibria, such as subgame perfect equilibria and admissible strategy profiles. It will also be fruitful to extend these results to the case of quantitative games such as those with mean payoff and discounted-sum objectives. Finally, the case of mixed rather than pure strategy equilibria would be an interesting extension.

References

[1] Rajeev Alur, Thomas A Henzinger & Orna Kupferman (1997): *Alternating-time temporal logic*. In: *International Symposium on Compositionality*, Springer, pp. 23–60.

[2] Roderick Bloem, Krishnendu Chatterjee & Barbara Jobstmann (2018): *Graph Games and Reactive Synthesis*. In: *Handbook of Model Checking*, Springer, pp. 921–962.

[3] Patricia Bouyer, Romain Brenguier, Nicolas Markey & Michael Ummels (2015): *Pure Nash equilibria in concurrent deterministic games*. *Logical Methods in Computer Science* 11.

[4] Thomas Brihaye, Véronique Bruyère, Aline Goeminne & Nathan Thomasset (2021): *On relevant equilibria in reachability games*. *Journal of Computer and System Sciences* 119, pp. 211–230.

[5] Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini & Jean-Francois Raskin (2016): *The complexity of rational synthesis*. In: *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, pp. 1–15.

[6] Bernd Finkbeiner (2016): *Synthesis of reactive systems*. In: *Dependable Software Systems Engineering*, IOS Press, pp. 72–98.

[7] Dana Fisman, Orna Kupferman & Yoad Lustig (2010): *Rational synthesis*. In: *Tools and Algorithms for the Construction and Analysis of Systems: TACAS 2010*, *LNCS* 6015, Springer, pp. 190–204.

[8] Georg Gottlob (1995): *NP trees and Carnap's modal logic*. *Journal of the ACM (JACM)* 42(2), pp. 421–457.

[9] Julian Gutierrez, Muhammad Najib, Giuseppe Perelli & Michael Wooldridge (2023): *On the complexity of rational verification*. Annals of Mathematics and Artificial Intelligence 91(4), pp. 409–430.

[10] Paul Hunter & Anuj Dawar (2005): *Complexity bounds for regular games*. In: International Symposium on Mathematical Foundations of Computer Science, Springer, pp. 495–506.

[11] Orna Kupferman, Giuseppe Perelli & Moshe Y Vardi (2016): *Synthesis with rational environments*. Annals of Mathematics and Artificial Intelligence 78(1), pp. 3–20.

[12] John F Nash Jr (1950): *Equilibrium points in n-person games*. Proceedings of the National Academy of Sciences 36(1), pp. 48–49.

[13] Robert Tarjan (1972): *Depth-first search and linear graph algorithms*. SIAM Journal on Computing 1(2), pp. 146–160.

[14] Klaus W Wagner (1990): *Bounded query classes*. SIAM Journal on Computing 19(5), pp. 833–846.