



Heriot-Watt University  
Research Gateway

## Equilibrium design for concurrent games

### Citation for published version:

Gutierrez, J, Najib, M, Perelli, G & Wooldridge, M 2019, Equilibrium design for concurrent games. in W Fokkink & R van Glabbeek (eds), *30th International Conference on Concurrency Theory (CONCUR 2019)*., 22, Leibniz International Proceedings in Informatics, vol. 140, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 30th International Conference on Concurrency Theory 2019, Amsterdam, Netherlands, 27/08/19. <https://doi.org/10.4230/LIPIcs.CONCUR.2019.22>

### Digital Object Identifier (DOI):

[10.4230/LIPIcs.CONCUR.2019.22](https://doi.org/10.4230/LIPIcs.CONCUR.2019.22)

### Link:

[Link to publication record in Heriot-Watt Research Portal](#)

### Document Version:

Publisher's PDF, also known as Version of record

### Published In:

30th International Conference on Concurrency Theory (CONCUR 2019)

### Publisher Rights Statement:

© Julian Guitierrez, Muhammad Najib, Giuseppe Perelli, and Michael Wooldridge.

### General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

### Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact [open.access@hw.ac.uk](mailto:open.access@hw.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

# Equilibrium Design for Concurrent Games

**Julian Gutierrez**

Department of Computer Science, University of Oxford, UK  
julian.gutierrez@cs.ox.ac.uk

**Muhammad Najib**

Department of Computer Science, University of Oxford, UK  
mnajib@cs.ox.ac.uk

**Giuseppe Perelli**

Department of Computer Science, University of Göteborg, Sweden  
giuseppe.perelli@gu.se

**Michael Wooldridge**

Department of Computer Science, University of Oxford, UK  
michael.wooldridge@cs.ox.ac.uk

---

## Abstract

In game theory, *mechanism design* is concerned with the design of incentives so that a desired outcome of the game can be achieved. In this paper, we study the design of incentives so that a desirable equilibrium is obtained, for instance, an equilibrium satisfying a given temporal logic property – a problem that we call *equilibrium design*. We base our study on a framework where system specifications are represented as temporal logic formulae, games as quantitative concurrent game structures, and players’ goals as mean-payoff objectives. In particular, we consider system specifications given by LTL and GR(1) formulae, and show that implementing a mechanism to ensure that a given temporal logic property is satisfied on some/every Nash equilibrium of the game, whenever such a mechanism exists, can be done in PSPACE for LTL properties and in  $\text{NP}/\Sigma_2^P$  for GR(1) specifications. We also study the complexity of various related decision and optimisation problems, such as optimality and uniqueness of solutions, and show that the complexities of all such problems lie within the polynomial hierarchy. As an application, equilibrium design can be used as an alternative solution to the rational synthesis and verification problems for concurrent games with mean-payoff objectives whenever no solution exists, or as a technique to repair, whenever possible, concurrent games with undesirable rational outcomes (Nash equilibria) in an optimal way.

**2012 ACM Subject Classification** Theory of computation → Modal and temporal logics; Computing methodologies → Multi-agent systems; Theory of computation → Algorithmic game theory

**Keywords and phrases** Games, Temporal logic, Synthesis, Model checking, Nash equilibrium

**Digital Object Identifier** 10.4230/LIPIcs.CONCUR.2019.22

**Acknowledgements** Najib acknowledges the financial support of the Indonesia Endowment Fund for Education (LPDP), and Perelli the support of the project “dSynMA”, funded by the ERC under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 772459).

## 1 Introduction

Over the past decade, there has been increasing interest in the use of game-theoretic equilibrium concepts such as Nash equilibrium in the analysis of concurrent and multi-agent systems (see, e.g., [3, 4, 8, 14, 15, 17, 23]). This work views a concurrent system as a game, with system components (agents) corresponding to players in the game, which are assumed to be acting rationally in pursuit of their individual preferences. Preferences may be specified by associating with each player a temporal logic goal formula, which the player desires to see satisfied, or by assuming that players receive rewards in each state the system visits, and seek to maximise the average reward they receive (the *mean payoff*). A further



© Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael Wooldridge;  
licensed under Creative Commons License CC-BY

30th International Conference on Concurrency Theory (CONCUR 2019).

Editors: Wan Fokkink and Rob van Glabbeek; Article No. 22; pp. 22:1–22:16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

possibility is to combine goals and rewards: players primarily seek the satisfaction of their goal, and only secondarily seek to maximise their mean payoff. The key decision problems in such settings relate to what temporal logic properties hold on computations of the system that may be generated by players choosing strategies that form a game-theoretic (Nash) equilibrium. These problems are typically computationally complex, since they subsume temporal logic synthesis [32]. If players have LTL goals, for example, then checking whether an LTL formula holds on some Nash equilibrium path in a concurrent game is 2EXPTIME-complete [14, 16, 17], rather than only PSPACE-complete as it is the case for model checking, certainly a computational barrier for the practical analysis and automated verification of reactive, concurrent, and multi-agent systems modelled as multi-player games.

Within this game-theoretic reasoning framework, a key issue is that individually rational choices can cause outcomes that are highly undesirable, and concurrent games also fall prey to this problem. This has motivated the development of techniques for modifying games, in order to avoid bad equilibria, or to facilitate good equilibria. *Mechanism design* is the problem of designing a game such that, if players behave rationally, then a desired outcome will be obtained [26]. Taxation and subsidy schemes are probably the most important class of techniques used in mechanism design. They work by levying taxes on certain actions (or providing subsidies), thereby incentivising players away from some outcomes towards others. The present paper studies the design of subsidy schemes (incentives) for concurrent games, so that a desired outcome (a Nash equilibrium in the game) can be obtained – a problem that we call *Equilibrium design*. We model agents as synchronously executing concurrent processes, with each agent receiving an integer payoff for every state the overall system visits; the overall payoff an agent receives over an infinite computation path is then defined to be the mean payoff over this path. While agents (naturally) seek to maximise their individual mean payoff, the designer of the subsidy scheme wishes to see some temporal logic formula satisfied, either on some or on every Nash equilibrium of the game.

With this model, we assume that the designer – an external principal – has a finite budget that is available for making subsidies, and this budget can be allocated across agent/state pairs. By allocating this budget appropriately, the principal can incentivise players away from some states and towards others. Since the principal has some temporal logic goal formula, it desires to allocate subsidies so that players are rationally incentivised to choose strategies so that the principal’s temporal logic goal formula is satisfied in the path that would result from executing the strategies. For this general problem, following [24], we identify two variants of the principal’s mechanism design problem, which we refer to as WEAK IMPLEMENTATION and STRONG IMPLEMENTATION. In the WEAK variant, we ask whether the principal can allocate the budget so that the goal is achieved on *some* computation path that would be generated by Nash equilibrium strategies in the resulting system; in the STRONG variation, we ask whether the principal can allocate the budget so that the resulting system has at least one Nash equilibrium, and moreover the temporal logic goal is satisfied on *all* paths that could be generated by Nash equilibrium strategies. For these two problems, we consider goals specified by LTL formulae or GR(1) formulae [5], give algorithms for each case, and classify the complexity of the problem. While LTL is a natural language for the specification of properties of concurrent and multi-agent systems, GR(1) is an LTL fragment that can be used to easily express several prefix-independent properties of computation paths of reactive systems, such as  $\omega$ -regular properties often used in automated formal verification. We then go on to study variations of these two problems, for example considering *optimality* and *uniqueness* of solutions, and show that the complexities of all such problems lie within the polynomial hierarchy, thus making them potentially amenable to efficient practical implementations. Table 1 summarises the main computational complexity results in the paper.

■ **Table 1** Summary of main complexity results.

	LTL Spec.	GR(1) Spec.
WEAK IMPLEMENTATION	PSPACE-complete (Thm. 6)	NP-complete (Thm. 7)
STRONG IMPLEMENTATION	PSPACE-complete (Cor. 9)	$\Sigma_2^P$ -complete (Thm. 10)
OPT-WI	FPSPACE-complete (Thm. 14)	$\text{FP}^{\text{NP}}$ -complete (Thm. 16)
OPT-SI	FPSPACE-complete (Thm. 22)	$\text{FP}^{\Sigma_2^P}$ -complete (Thm. 25)
EXACT-WI	PSPACE-complete (Cor. 15)	$\text{D}^P$ -complete (Cor. 17)
EXACT-SI	PSPACE-complete (Cor. 23)	$\text{D}_2^P$ -complete (Cor. 26)
UOPT-WI	PSPACE-complete (Cor. 18)	$\Delta_2^P$ -complete (Cor. 19)
UOPT-SI	PSPACE-complete (Cor. 27)	$\Delta_3^P$ -complete (Cor. 28)

## 2 Preliminaries

**Linear Temporal Logic.** LTL [31] extends classical propositional logic with two operators,  $\mathbf{X}$  (“next”) and  $\mathbf{U}$  (“until”), that can be used to express properties of paths. The syntax of LTL is defined with respect to a set AP of atomic propositions as follows:

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

where  $p \in \text{AP}$ . As commonly found in the LTL literature, we use of the following abbreviations:  $\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2)$ ,  $\varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2$ ,  $\mathbf{F}\varphi \equiv \top \mathbf{U} \varphi$ , and  $\mathbf{G}\varphi \equiv \neg\mathbf{F}\neg\varphi$ .

We interpret formulae of LTL with respect to pairs  $(\alpha, t)$ , where  $\alpha \in (2^{\text{AP}})^\omega$  is an infinite sequence of atomic proposition evaluations that indicates which propositional variables are true in every time point and  $t \in \mathbb{N}$  is a temporal index into  $\alpha$ . Formally, the semantics of LTL formulae is given by the following rules:

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

If  $(\alpha, 0) \models \varphi$ , we write  $\alpha \models \varphi$  and say that  $\alpha$  *satisfies*  $\varphi$ .

**General Reactivity of rank 1.** The language of *General Reactivity of rank 1*, denoted GR(1), is the fragment of LTL given by formulae written in the following form [5]:

$$(\mathbf{GF}\psi_1 \wedge \dots \wedge \mathbf{GF}\psi_m) \rightarrow (\mathbf{GF}\varphi_1 \wedge \dots \wedge \mathbf{GF}\varphi_n),$$

where each subformula  $\psi_i$  and  $\varphi_i$  is a Boolean combination of atomic propositions.

**Mean-Payoff.** For a sequence  $r \in \mathbb{R}^\omega$ , let  $\text{mp}(r)$  be the *mean-payoff* value of  $r$ , that is,

$$\text{mp}(r) = \liminf_{n \rightarrow \infty} \text{avg}_n(r)$$

where, for  $n \in \mathbb{N} \setminus \{0\}$ , we define  $\text{avg}_n(r) = \frac{1}{n} \sum_{j=0}^{n-1} r_j$ , with  $r_j$  the  $(j+1)$ th element of  $r$ .

**Arenas.** An *arena* is a tuple  $A = \langle N, \text{Ac}, \text{St}, s_0, \text{tr}, \lambda \rangle$  where  $N$ ,  $\text{Ac}$ , and  $\text{St}$  are finite non-empty sets of *players* (write  $N = |N|$ ), *actions*, and *states*, respectively; if needed, we write  $\text{Ac}_i(s)$ , to denote the set of actions available to player  $i$  at  $s$ ;  $s_0 \in \text{St}$  is the *initial state*;  $\text{tr} : \text{St} \times \vec{\text{Ac}} \rightarrow \text{St}$  is a *transition function* mapping each pair consisting of a state  $s \in \text{St}$  and an *action profile*  $\vec{\mathbf{a}} \in \vec{\text{Ac}} = \text{Ac}^N$ , one for each player, to a successor state; and  $\lambda : \text{St} \rightarrow 2^{\text{AP}}$  is a labelling function, mapping every state to a subset of *atomic propositions*.

We sometimes call an action profile  $\vec{\mathbf{a}} = (\mathbf{a}_1, \dots, \mathbf{a}_n) \in \vec{\text{Ac}}$  a *decision*, and denote  $\mathbf{a}_i$  the action taken by player  $i$ . We also consider *partial* decisions. For a set of players  $C \subseteq N$  and action profile  $\vec{\mathbf{a}}$ , we let  $\vec{\mathbf{a}}_C$  and  $\vec{\mathbf{a}}_{-C}$  be two tuples of actions, respectively, one for all players in  $C$  and one for all players in  $N \setminus C$ . We also write  $\vec{\mathbf{a}}_i$  for  $\vec{\mathbf{a}}_{\{i\}}$  and  $\vec{\mathbf{a}}_{-i}$  for  $\vec{\mathbf{a}}_{N \setminus \{i\}}$ . For two decisions  $\vec{\mathbf{a}}$  and  $\vec{\mathbf{a}}'$ , we write  $(\vec{\mathbf{a}}_C, \vec{\mathbf{a}}'_{-C})$  to denote the decision where the actions for players in  $C$  are taken from  $\vec{\mathbf{a}}$  and the actions for players in  $N \setminus C$  are taken from  $\vec{\mathbf{a}}'$ .

A *path*  $\pi = (s_0, \vec{\mathbf{a}}^0), (s_1, \vec{\mathbf{a}}^1) \dots$  is an infinite sequence in  $(\text{St} \times \vec{\text{Ac}})^\omega$  such that  $\text{tr}(s_k, \vec{\mathbf{a}}^k) = s_{k+1}$  for all  $k$ . Paths are generated in the arena by each player  $i$  selecting a *strategy*  $\sigma_i$  that will define how to make choices over time. We model strategies as finite state machines with output. Formally, for arena  $A$ , a strategy  $\sigma_i = (Q_i, q_i^0, \delta_i, \tau_i)$  for player  $i$  is a finite state machine with output (a transducer), where  $Q_i$  is a finite and non-empty set of *internal states*,  $q_i^0$  is the *initial state*,  $\delta_i : Q_i \times \vec{\text{Ac}} \rightarrow Q_i$  is a deterministic *internal transition function*, and  $\tau_i : Q_i \rightarrow \text{Ac}_i$  an *action function*. Let  $\text{Str}_i$  be the set of strategies for player  $i$ . Note that this definition implies that strategies have *perfect information*<sup>1</sup> and finite memory (although we impose no bounds on memory size).

A *strategy profile*  $\vec{\sigma} = (\sigma_1, \dots, \sigma_n)$  is a vector of strategies, one for each player. As with actions,  $\vec{\sigma}_i$  denotes the strategy assigned to player  $i$  in profile  $\vec{\sigma}$ . Moreover, by  $(\vec{\sigma}_B, \vec{\sigma}'_C)$  we denote the combination of profiles where players in disjoint  $B$  and  $C$  are assigned their corresponding strategies in  $\vec{\sigma}$  and  $\vec{\sigma}'$ , respectively. Once a state  $s$  and profile  $\vec{\sigma}$  are fixed, the game has an *outcome*, a path in  $A$ , denoted by  $\pi(\vec{\sigma}, s)$ . Because strategies are deterministic,  $\pi(\vec{\sigma}, s)$  is the unique path induced by  $\vec{\sigma}$ , that is, the sequence  $s_0, s_1, s_2, \dots$  such that

- $s_{k+1} = \text{tr}(s_k, (\tau_1(q_1^k), \dots, \tau_n(q_n^k)))$ , and
- $q_i^{k+1} = \delta_i(s_k, (\tau_1(q_1^k), \dots, \tau_n(q_n^k)))$ , for all  $k \geq 0$ .

Furthermore, we simply write  $\pi(\vec{\sigma})$  for  $\pi(\vec{\sigma}, s_0)$ .

Arenas define the dynamic structure of games, but lack a central aspect of a game: preferences, which give games their strategic structure. A *multi-player game* is obtained from an arena  $A$  by associating each player with a goal. We consider multi-player games with **mp** goals. A multi-player **mp** game is a tuple  $\mathcal{G} = \langle A, (\mathbf{w}_i)_{i \in N} \rangle$ , where  $A$  is an arena and  $\mathbf{w}_i : \text{St} \rightarrow \mathbb{Z}$  is a function mapping, for every player  $i$ , every state of the arena into an integer number. In any game with arena  $A$ , a path  $\pi$  in  $A$  induces a sequence  $\lambda(\pi) = \lambda(s_0)\lambda(s_1) \dots$  of sets of atomic propositions; if, in addition,  $A$  is the arena of an **mp** game, then, for each player  $i$ , the sequence  $\mathbf{w}_i(\pi) = \mathbf{w}_i(s_0)\mathbf{w}_i(s_1) \dots$  of weights is also induced. Unless stated otherwise, for a game  $\mathcal{G}$  and a path  $\pi$  in it, the payoff of player  $i$  is  $\text{pay}_i(\pi) = \text{mp}(\mathbf{w}_i(\pi))$ .

<sup>1</sup> Mean-payoff games with imperfect information are generally undecidable [13].

**Nash equilibrium.** Using payoff functions, we can define the game-theoretic concept of Nash equilibrium [26]. For a multi-player game  $\mathcal{G}$ , a strategy profile  $\vec{\sigma}$  is a *Nash equilibrium* of  $\mathcal{G}$  if, for every player  $i$  and strategy  $\sigma'_i$  for player  $i$ , we have

$$\text{pay}_i(\pi(\vec{\sigma})) \geq \text{pay}_i(\pi((\vec{\sigma}_{-i}, \sigma'_i))) .$$

Let  $\text{NE}(\mathcal{G})$  be the set of Nash equilibria of  $\mathcal{G}$ .

### 3 From Mechanism Design to Equilibrium Design

We now describe the two main problems that are our focus of study. As discussed in the introduction, such problems are closely related to the well-known problem of *mechanism design* in game theory. Consider a system populated by agents  $\mathbb{N}$ , where each agent  $i \in \mathbb{N}$  wants to maximise its payoff  $\text{pay}_i(\cdot)$ . As in a mechanism design problem, we assume there is an external *principal* who has a goal  $\varphi$  that it wants the system to satisfy, and to this end, wants to incentivise the agents to act collectively and rationally so as to bring about  $\varphi$ . In our model, incentives are given by *subsidy schemes* and goals by temporal logic formulae.

**Subsidy Schemes.** A subsidy scheme defines additional imposed rewards over those given by the weight function  $w$ . While the weight function  $w$  is fixed for any given game, the principal is assumed to be at liberty to define a subsidy scheme as they see fit. Since agents will seek to maximise their overall rewards, the principal can incentivise agents away from performing visiting some states and towards visiting others; if the principal designs the subsidy scheme correctly, the agents are incentivised to choose a strategy profile  $\vec{\sigma}$  such that  $\pi(\vec{\sigma}) \models \varphi$ . Formally, we model a subsidy scheme as a function  $\kappa : \mathbb{N} \rightarrow \text{St} \rightarrow \mathbb{N}$ , where the intended interpretation is that  $\kappa(i)(s)$  is the subsidy in the form of a natural number  $k \in \mathbb{N}$  that would be imposed on player  $i$  if such a player visits state  $s \in \text{St}$ . For instance, if we have  $w_i(s) = 1$  and  $\kappa(i)(s) = 2$ , then player  $i$  gets  $1 + 2 = 3$  for visiting such a state. For simplicity, hereafter we write  $\kappa_i(s)$  instead of  $\kappa(i)(s)$  for the subsidy for player  $i$ .

Notice that having an unlimited fund for a subsidy scheme would make some problems trivial, as the principal can always incentivise players to satisfy  $\varphi$  (provided that there is a path in  $A$  satisfying  $\varphi$ ). A natural and more interesting setting is that the principal is given a constraint in the form of *budget*  $\beta \in \mathbb{N}$ . The principal then can only spend within the budget limit. To make this clearer, we first define the *cost* of a subsidy scheme  $\kappa$  as follows.

► **Definition 1.** Given a game  $\mathcal{G}$  and subsidy scheme  $\kappa$ , we let  $\text{cost}(\kappa) = \sum_{i \in \mathbb{N}} \sum_{s \in \text{St}} \kappa_i(s)$ .

We say that a subsidy scheme  $\kappa$  is *admissible* if it does not exceed the budget  $\beta$ , that is, if  $\text{cost}(\kappa) \leq \beta$ . Let  $\mathcal{K}(\mathcal{G}, \beta)$  denote the set of admissible subsidy schemes over  $\mathcal{G}$  given budget  $\beta \in \mathbb{N}$ . Thus we know that for each  $\kappa \in \mathcal{K}(\mathcal{G}, \beta)$  we have  $\text{cost}(\kappa) \leq \beta$ . We write  $(\mathcal{G}, \kappa)$  to denote the resulting game after the application of subsidy scheme  $\kappa$  on game  $\mathcal{G}$ . Formally, we define the application of some subsidy scheme on a game as follows.

► **Definition 2.** Given a game  $\mathcal{G} = \langle A, (w_i)_{i \in \mathbb{N}} \rangle$  and an admissible subsidy scheme  $\kappa$ , we define  $(\mathcal{G}, \kappa) = \langle A, (w'_i)_{i \in \mathbb{N}} \rangle$ , where  $w'_i(s) = w_i(s) + \kappa_i(s)$ , for each  $i \in \mathbb{N}$  and  $s \in \text{St}$ .

We now come to the main question(s) that we consider in the remainder of the paper. We ask whether the principal can find a subsidy scheme that will incentivise players to collectively choose a rational outcome (a Nash equilibrium) that satisfies its temporal logic goal  $\varphi$ . We call this problem *equilibrium design*. Following [24], we define two variants of this problem, a *weak* and a *strong* implementation of the equilibrium design problem. The formal definition of the problems and the analysis of their respective computational complexity are presented in the next sections.

#### 4 Equilibrium Design: Weak Implementation

In this section, we study the weak implementation of the equilibrium design problem, a logic-based computational variant of the principal's mechanism design problem in game theory. We assume that the principal has full knowledge of the game  $\mathcal{G}$  under consideration, that is, the principal uses all the information available of  $\mathcal{G}$  to find the appropriate subsidy scheme, if such a scheme exists. We now formally define the weak variant of the implementation problem, and study its respective computational complexity, first with respect to goals (specifications) given by LTL formulae and then with respect to GR(1) formulae.

Let  $\text{WI}(\mathcal{G}, \varphi, \beta)$  denote the set of subsidy schemes over  $\mathcal{G}$  given budget  $\beta$  that satisfy a formula  $\varphi$  in at least one path  $\pi$  generated by  $\vec{\sigma} \in \text{NE}(\mathcal{G})$ . Formally

$$\text{WI}(\mathcal{G}, \varphi, \beta) = \{\kappa \in \mathcal{K}(\mathcal{G}, \beta) : \exists \vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa) \text{ s.t. } \pi(\vec{\sigma}) \models \varphi\}.$$

► **Definition 3** (WEAK IMPLEMENTATION). *Given a game  $\mathcal{G}$ , formula  $\varphi$ , and budget  $\beta$ :*

*Is it the case that  $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ ?*

In order to solve WEAK IMPLEMENTATION, we first characterise the Nash equilibria of a multi-player concurrent game in terms of punishment strategies. To do this in our setting, we recall the notion of secure values for mean-payoff games [33].

For a player  $i$  and a state  $s \in \text{St}$ , by  $\text{pun}_i(s)$  we denote the punishment value of  $i$  over  $s$ , that is, the maximum payoff that  $i$  can achieve from  $s$ , when all other players behave adversarially. Such a value can be computed by considering the corresponding two-player zero-sum mean-payoff game [35]. Thus, it is in  $\text{NP} \cap \text{coNP}$ , and note that both player  $i$  and coalition  $N \setminus \{i\}$  can achieve the optimal value of the game using *memoryless* strategies. Then, for a player  $i$  and a value  $z \in \mathbb{R}$ , a pair  $(s, \vec{a})$  is  $z$ -secure for player  $i$  if  $\text{pun}_i(\text{tr}(s, (\vec{a}_{-i}, \mathbf{a}'_i))) \leq z$  for every  $\mathbf{a}'_i \in \text{Ac}$ . Write  $\text{pun}_i(\mathcal{G})$  for the set of punishment values for player  $i$  in  $\mathcal{G}$ .

► **Theorem 4.** *For every mp game  $\mathcal{G}$  and ultimately periodic path  $\pi = (s_0, \vec{a}_0), (s_1, \vec{a}_1), \dots$ , the following are equivalent:*

1. *There is  $\vec{\sigma} \in \text{NE}(\mathcal{G})$  such that  $\pi = \pi(\vec{\sigma}, s_0)$ ;*
2. *There exists  $z \in \mathbb{R}^N$ , where  $z_i \in \text{pun}_i(\mathcal{G})$  such that, for every  $i \in N$* 
  - a. *for all  $k \in \mathbb{N}$ , the pair  $(s_k, \vec{a}^k)$  is  $z_i$ -secure for  $i$ , and*
  - b.  *$z_i \leq \text{pay}_i(\pi)$ .*

The characterisation of Nash Equilibria provided in Theorem 4 will allow us to turn the WEAK IMPLEMENTATION problem into a *path finding* problem over  $(\mathcal{G}, \kappa)$ . On the other hand, with respect to the budget  $\beta$  that the principal has at its disposal, the definition of subsidy scheme function  $\kappa$  implies that the size of  $\mathcal{K}(\mathcal{G}, \beta)$  is bounded, and particularly, it is bounded by  $\beta$  and the number of agents and states in the game  $\mathcal{G}$ , in the following way.

► **Proposition 5.** *Given a game  $\mathcal{G}$  with  $|N|$  players and  $|\text{St}|$  states and budget  $\beta$ , it holds that*

$$|\mathcal{K}(\mathcal{G}, \beta)| = \frac{\beta + 1}{m} \binom{\beta + m}{\beta + 1},$$

*with  $m = |N \times \text{St}|$  being the number of pairs of possible agents and states.*

From Proposition 5 we derive that the number of possible subsidy schemes is *polynomial* in the budget  $\beta$  and singly *exponential* in both the number of agents and states in the game. At this point, solving WEAK IMPLEMENTATION can be done with the following procedure:

1. Guess:
  - a subsidy scheme  $\kappa \in \mathcal{K}(\mathcal{G}, \beta)$ ,
  - a state  $s \in \text{St}$  for every player  $i \in N$ , and
  - punishment memoryless strategies  $(\vec{\sigma}_{-1}, \dots, \vec{\sigma}_{-n})$  for all players  $i \in N$ ;
2. Compute  $(\mathcal{G}, \kappa)$ ;
3. Compute  $z \in \mathbb{R}^N$ ;
4. Compute the game  $(\mathcal{G}, \kappa)[z]$  by removing the states  $s$  such that  $\text{pun}_i(s) \leq z_i$  for some player  $i$  and the transitions  $(s, \vec{a}_{-i})$  that are not  $z_i$  secure for player  $i$ ;
5. Check whether there exists an ultimately periodic path  $\pi$  in  $(\mathcal{G}, \kappa)[z]$  such that  $\pi \models \varphi$  and  $z_i \leq \text{pay}_i(\pi)$  for every player  $i \in N$ .

Since the set  $\mathcal{K}(\mathcal{G}, \beta)$  is finitely bounded (Proposition 5), and punishment strategies only need to be memoryless, thus also finitely bounded, clearly step 1 can be guessed nondeterministically. Moreover, each of the guessed elements is of polynomial size, thus this step can be done (deterministically) in polynomial space. Step 2 clearly can be done in polynomial time. Step 3 can also be done in polynomial time since, given  $(\vec{\sigma}_{-1}, \dots, \vec{\sigma}_{-n})$ , we can compute  $z$  solving  $|N|$  one-player mean-payoff games, one for each player  $i$  [35, Thm. 6]. For step 5, we will use Theorem 4 and consider two cases, one for LTL specifications and one for GR(1) specifications. Firstly, for LTL specifications, consider the formula

$$\varphi_{\text{WI}} := \varphi \wedge \bigwedge_{i \in N} (\text{mp}(i) \geq z_i)$$

written in  $\text{LTL}^{\text{Lim}}$  [7], an extension of LTL where statements about mean-payoff values over a given weighted arena can be made.<sup>2</sup> The semantics of the temporal operators of  $\text{LTL}^{\text{Lim}}$  is just like the one for LTL over infinite computation paths  $\pi = s_0, s_1, s_3, \dots$ . On the other hand, the meaning of  $\text{mp}(i) \geq z_i$  is simply that such an atomic formula is true if, and only if, the mean-payoff value of  $\pi$  with respect to player  $i$  is greater or equal to  $z_i$ , a constant real value; that is,  $\text{mp}(i) \geq z_i$  is true in  $\pi$  if and only if  $\text{pay}_i(\pi) = \text{mp}(w_i(\pi))$  is greater or equal than constant value  $z_i$ . Formula  $\varphi_{\text{WI}}$  corresponds exactly to 2(b) in Theorem 4. Furthermore, since every path in  $(\mathcal{G}, \kappa)[z]$  satisfies condition 2(a) of Theorem 4, every computation path of  $(\mathcal{G}, \kappa)[z]$  that satisfies  $\varphi_{\text{WI}}$  is a witness to the WEAK IMPLEMENTATION problem.

► **Theorem 6.** WEAK IMPLEMENTATION *with LTL specifications is PSPACE-complete.*

**Proof.** Membership follows from the procedure above and the fact that model checking for  $\text{LTL}^{\text{Lim}}$  is PSPACE-complete [7]. Hardness follows from the fact that LTL model checking is a special case of WEAK IMPLEMENTATION. For instance, consider the case in which all weights for all players are set to the same value, say 0, and the principal has budget  $\beta = 0$ . ◀

**Case with GR(1) specifications.** One of the main bottlenecks of our procedure to solve WEAK IMPLEMENTATION lies in step 5, where we solve an  $\text{LTL}^{\text{Lim}}$  model checking problem. To reduce the complexity of our decision procedure, we consider WEAK IMPLEMENTATION with the specification  $\varphi$  expressed in the GR(1) sublanguage of LTL. With this specification language, the path finding problem can be solved without model-checking the  $\text{LTL}^{\text{Lim}}$  formula given before. In order to do this, we can define a linear program (LP) such that the LP has a solution if and only if  $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ . From our previous procedure, observe that step 1 can be done nondeterministically in polynomial time, and steps 2–4 can be done

<sup>2</sup> The formal semantics of  $\text{LTL}^{\text{Lim}}$  can be found in [7]. We prefer to give only an informal description here.



(deterministically) in polynomial time. Furthermore, using LP, we also can check step 5 deterministically in polynomial time. For the lower-bound, we use [33] and note that if  $\varphi = \top$  and  $\beta = 0$ , then the problem reduces to checking whether the underlying **mp** game has a Nash equilibrium. Based on the above observations, we have the following result.

► **Theorem 7.** WEAK IMPLEMENTATION *with GR(1) specifications is NP-complete.*

**Proof sketch.** For the upper bound, we define an LP of size polynomial in  $(\mathcal{G}, \kappa)$  having a solution if and only if there is an ultimately periodic path  $\pi$  such that  $z_i \leq \text{pay}_i(\pi)$  and satisfies the GR(1) specification. Recall that  $\varphi$  has the following form

$$\varphi = \bigwedge_{l=1}^m \mathbf{GF}\psi_l \rightarrow \bigwedge_{r=1}^n \mathbf{GF}\theta_r,$$

and let  $V(\psi_l)$  and  $V(\theta_r)$  be the subset of states in  $(\mathcal{G}, \kappa)$  that satisfy the Boolean combinations  $\psi_l$  and  $\theta_r$ , respectively. Property  $\varphi$  is satisfied on  $\pi$  if, and only if, either  $\pi$  visits every state in  $V(\theta_r)$  infinitely often or some of the states in  $V(\psi_l)$  only a finite number of times. For the game  $(\mathcal{G}, \kappa)[z]$ , let  $W = (V, E, (w_a)_{a \in \mathbb{N}})$  be the underlying multi-weighted graph, and for every edge  $e \in E$  introduce a variable  $x_e$ . Informally, the value of  $x_e$  is the number of times that  $e$  is used on a cycle. Formally, let  $\text{src}(e) = \{v \in V : \exists w e = (v, w) \in E\}$ ;  $\text{trg}(e) = \{v \in V : \exists w e = (w, v) \in E\}$ ;  $\text{out}(v) = \{e \in E : \text{src}(e) = v\}$ ; and  $\text{in}(v) = \{e \in E : \text{trg}(e) = v\}$ . Now, consider  $\psi_l$  for some  $1 \leq l \leq m$ , and define the following linear program  $\text{LP}(\psi_l)$ :

**Eq1:**  $x_e \geq 0$  for each edge  $e$  – a basic consistency criterion;

**Eq2:**  $\sum_{e \in E} x_e \geq 1$  – at least one edge is chosen;

**Eq3:** for each  $a \in \mathbb{N}$ ,  $\sum_{e \in E} w_a(\text{src}(e))x_e \geq 0$  – total sum of any solution is non-negative;

**Eq4:**  $\sum_{\text{src}(e) \cap V(\psi_l) \neq \emptyset} x_e = 0$  – no state in  $V(\psi_l)$  is in the cycle associated with the solution;

**Eq5:** for each  $v \in V$ ,  $\sum_{e \in \text{out}(v)} x_e = \sum_{e \in \text{in}(v)} x_e$  – this condition says that the number of times one enters a vertex is equal to the number of times one leaves that vertex.

$\text{LP}(\psi_l)$  has a solution if and only if there is a path  $\pi$  in  $\mathcal{G}$  such that  $z_i \leq \text{pay}_i(\pi)$  for every player  $i$  and visits  $V(\psi_l)$  only *finitely many times*. Consider now the linear program  $\text{LP}(\theta_1, \dots, \theta_n)$  defined as follows. Eq1–Eq3 as well as Eq5 are as in  $\text{LP}(\psi_l)$ , and:

**Eq4:** for all  $1 \leq r \leq n$ ,  $\sum_{\text{src}(e) \cap V(\theta_r) \neq \emptyset} x_e \geq 1$  – this condition says that, for every  $V(\theta_r)$ , at least one state in  $V(\theta_r)$  is in the cycle associated with the solution of the linear program.

In this case,  $\text{LP}(\theta_1, \dots, \theta_n)$  has a solution if and only if there exists a path  $\pi$  such that  $z_i \leq \text{pay}_i(\pi)$  for every player  $i$  and visits every  $V(\theta_r)$  *infinitely many times*. Since the constructions above are polynomial in the size of both  $(\mathcal{G}, \kappa)$  and  $\varphi$ , we can conclude it is possible to check in NP the statement that there is a path  $\pi$  satisfying  $\varphi$  such that  $z_i \leq \text{pay}_i(\pi)$  for every player  $i$  in the game if and only if one of the two linear programs defined above has a solution. For the lower-bound, we use [33] as discussed before. ◀

We now turn our attention to the strong implementation of the equilibrium design problem. As in this section, we first consider LTL specifications and then GR(1) specifications.

## 5 Equilibrium Design: Strong Implementation

Although the principal may find  $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$  to be good news, it might not be good enough. It could be that even though there is a desirable Nash equilibrium, the others might be undesirable. This motivates us to consider the *strong implementation* variant of equilibrium design. Intuitively, in a strong implementation, we require that *every* Nash equilibrium outcome satisfies the specification  $\varphi$ , for a *non-empty* set of outcomes. Then, let  $\text{SI}(\mathcal{G}, \varphi, \beta)$  denote the set of subsidy schemes  $\kappa$  given budget  $\beta$  over  $\mathcal{G}$  such that:

1.  $(\mathcal{G}, \kappa)$  has at least one Nash equilibrium outcome,
2. every Nash equilibrium outcome of  $(\mathcal{G}, \kappa)$  satisfies  $\varphi$ .

Formally we define it as follows:

$$\text{SI}(\mathcal{G}, \varphi, \beta) = \{\kappa \in \mathcal{K}(\mathcal{G}, \beta) : \text{NE}(\mathcal{G}, \kappa) \neq \emptyset \wedge \forall \vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa) \text{ s.t. } \pi(\vec{\sigma}) \models \varphi\}.$$

This gives us the following decision problem:

► **Definition 8** (STRONG IMPLEMENTATION). *Given a game  $\mathcal{G}$ , formula  $\varphi$ , and budget  $\beta$ :*

*Is it the case that  $\text{SI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ ?*

STRONG IMPLEMENTATION can be solved with a 5-step procedure where the first four steps are as in WEAK IMPLEMENTATION, and the last step (step 5) is as follows:

5. Check whether:

- (a) there is no ultimately periodic path  $\pi$  in  $(\mathcal{G}, \kappa)[z]$  such that  $z_i \leq \text{pay}_i(\pi)$  for each  $i \in \mathbb{N}$ ;
- (b) there is an ultimately periodic path  $\pi$  in  $(\mathcal{G}, \kappa)[z]$  such that  $\pi \models \neg\varphi$  and  $z_i \leq \text{pay}_i(\pi)$ , for each  $i \in \mathbb{N}$ .

For step 5, observe that a positive answer to 5(a) or 5(b) is a counterexample to  $\kappa \in \text{SI}(\mathcal{G}, \varphi, \beta)$ . Then, to carry out this procedure for the STRONG IMPLEMENTATION problem with LTL specifications, consider the following  $\text{LTL}^{\text{Lim}}$  formulae:

$$\varphi_{\exists} = \bigwedge_{i \in \mathbb{N}} (\text{mp}(i) \geq z_i);$$

$$\varphi_{\forall} = \varphi_{\exists} \rightarrow \varphi.$$

Notice that the expression  $\text{NE}(\mathcal{G}, \kappa) \neq \emptyset$  can be expressed as “there exists a path  $\pi$  in  $\mathcal{G}$  that satisfies formula  $\varphi_{\exists}$ ”. On the other hand, the expression  $\forall \vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa)$  such that  $\pi(\vec{\sigma}) \models \varphi$  can be expressed as “for every path  $\pi$  in  $\mathcal{G}$ , if  $\pi$  satisfies formula  $\varphi_{\exists}$ , then  $\pi$  also satisfies formula  $\varphi$ ”. Thus, using these two formulae, we obtain the following result.

► **Corollary 9.** STRONG IMPLEMENTATION *with LTL specifications is PSPACE-complete.*

**Proof.** Membership follows from the fact that step 5(a) can be solved by existential  $\text{LTL}^{\text{Lim}}$  model checking, whereas step 5(b) by universal  $\text{LTL}^{\text{Lim}}$  model checking – both clearly in PSPACE by Savitch’s theorem. Hardness is similar to the construction in Theorem 6. ◀

**Case with GR(1) specifications.** Notice that the first part, *i.e.*,  $\text{NE}(\mathcal{G}, \kappa) \neq \emptyset$  can be solved in NP [33]. For the second part, observe that

$$\forall \vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa) \text{ such that } \pi(\vec{\sigma}) \models \varphi$$

is equivalent to

$$\neg \exists \vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa) \text{ such that } \pi(\vec{\sigma}) \models \neg\varphi.$$

Thus we have

$$\neg\varphi = \bigwedge_{l=1}^m \mathbf{GF}\psi_l \wedge \neg \left( \bigwedge_{r=1}^n \mathbf{GF}\theta_r \right).$$

To check this, we modify the LP in Theorem 7. Specifically, we modify Eq4 in  $\text{LP}(\theta_1, \dots, \theta_n)$  to encode the  $\theta$ -part of  $\neg\varphi$ . Thus, we have the following equation in  $\text{LP}'(\theta_1, \dots, \theta_n)$ :

**Eq4:** there exists  $r$ ,  $1 \leq r \leq n$ ,  $\Sigma_{\text{src}(e) \cap V(\theta_r) \neq \emptyset} x_e = 0$  – this condition ensures that at least one set  $V(\theta_r)$  does not have any state in the cycle associated with the solution.

In this case,  $\text{LP}'(\theta_1, \dots, \theta_n)$  has a solution if and only if there is a path  $\pi$  such that  $z_i \leq \text{pay}_i(\pi)$  for every player  $i$  and, for at least one  $V(\theta_r)$ , its states are visited only *finitely many times*. Thus, we have a procedure that checks if there is a path  $\pi$  that satisfies  $\neg\varphi$  such that  $z_i \leq \text{pay}_i(\pi)$  for every player  $i$ , if and only if both linear programs have a solution. Using this new construction, we can now prove the following result.

► **Theorem 10.** STRONG IMPLEMENTATION *with GR(1) specifications is  $\Sigma_2^P$ -complete.*

**Proof sketch.** For membership, observe that by rearranging the problem statement, we have the following question: Check whether the following expression is true

$$\exists \kappa \in \mathcal{K}(\mathcal{G}, \beta), \quad (1)$$

$$\exists \vec{\sigma} \in \sigma_1 \times \dots \times \sigma_n, \text{ such that } \vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa), \quad (2)$$

and

$$\forall \vec{\sigma}' \in \sigma_1 \times \dots \times \sigma_n, \text{ if } \vec{\sigma}' \in \text{NE}(\mathcal{G}, \kappa) \text{ then } \pi(\vec{\sigma}') \models \varphi. \quad (3)$$

Statement (2) can be checked in NP (Theorem 4), whereas verifying statement (3) is in coNP; to see this, notice that we can rephrase (3) as follows:  $\neg \exists z \in \{\text{pun}_i(s) : s \in \text{St}\}^N$  such that both  $\text{LP}(\psi_l)$  and  $\text{LP}'(\theta_1, \dots, \theta_n)$  have a solution in  $(\mathcal{G}, \kappa)[z]$ . Thus, membership in  $\Sigma_2^P$  follows. We prove hardness via a reduction from  $\text{QSAT}_2$  (satisfiability of quantified Boolean formulae with 2 alternations), which is known to be  $\Sigma_2^P$ -complete [28]. ◀

## 6 Optimality and Uniqueness of Solutions

Having asked the questions studied in the previous sections, the principal – the *designer* in the equilibrium design problem – may want to explore further information. Because the power of the principal is limited by its budget, and because from the point of view of the system, it may be associated with a reward (*e.g.*, money, savings, etc.) or with the inverse of the amount of a finite resource (*e.g.*, time, energy, etc.) an obvious question is asking about *optimal* solutions. This leads us to *optimisation* variations of the problems we have studied. Informally, in this case, we ask what is the least budget that the principal needs to ensure that the implementation problems have positive solutions. The principal may also want to know whether a given subsidy scheme is *unique*, so that there is no point in looking for any other solutions to the problem. In this section, we investigate these kind of problems, and classify our study into two parts, one corresponding to the WEAK IMPLEMENTATION problem and another one corresponding to the STRONG IMPLEMENTATION problem.

### 6.1 Optimality and Uniqueness in the Weak Domain

We can now define formally some of the problems that we will study in the rest of this section. To start, the optimisation variant for WEAK IMPLEMENTATION is defined as follows.

► **Definition 11** (OPT-WI). *Given a game  $\mathcal{G}$  and a specification formula  $\varphi$ :*

*What is the optimum budget  $\beta$  such that  $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ ?*

Another natural problem, which is related to OPT-WI, is the “exact” variant – a membership question. In this case, in addition to  $\mathcal{G}$  and  $\varphi$ , we are also given an integer  $b$ , and ask whether it is indeed the smallest amount of budget that the principal has to spend for some optimal weak implementation. This decision problem is formally defined as follows.

► **Definition 12** (EXACT-WI). *Given a game  $\mathcal{G}$ , a specification formula  $\varphi$ , and an integer  $b$ :*

*Is  $b$  equal to the optimum budget for  $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ ?*

To study these problems, it is useful to introduce some concepts first. More specifically, let us introduce the concept of *implementation efficiency*. We say that a WEAK IMPLEMENTATION (resp. STRONG IMPLEMENTATION) is *efficient* if  $\beta = \text{cost}(\kappa)$  and there is no  $\kappa'$  such that  $\text{cost}(\kappa') < \text{cost}(\kappa)$  and  $\kappa' \in \text{WI}(\mathcal{G}, \varphi, \beta)$  (resp.  $\kappa' \in \text{SI}(\mathcal{G}, \varphi, \beta)$ ). In addition to the concept of efficiency for an implementation problem, it is also useful to have the following result.

► **Proposition 13.** *Let  $z_i$  be the largest payoff that player  $i$  can get after deviating from a path  $\pi$ . The optimum budget is an integer between 0 and  $\sum_{i \in N} z_i \cdot (|\text{St}| - 1)$ .*

Using Proposition 13, we can show that both OPT-WI and EXACT-WI can be solved in PSPACE for LTL specifications. Intuitively, the reason is that we can use the upper bound given by Proposition 13 to go through all possible solutions in exponential time, but using only nondeterministic polynomial space. Formally, we have the following results.

► **Theorem 14.** *OPT-WI with LTL specifications is FSPACE-complete.*

► **Corollary 15.** *EXACT-WI with LTL specifications is PSPACE-complete.*

The fact that both OPT-WI and EXACT-WI with LTL specifications can be answered in, respectively, FSPACE and PSPACE does not come as a big surprise: checking an instance can be done using polynomial space and there are only exponentially many instances to be checked. However, for OPT-WI and EXACT-WI with GR(1) specifications, these two problems are more interesting.

► **Theorem 16.** *OPT-WI with GR(1) specifications is  $\text{FP}^{\text{NP}}$ -complete.*

**Proof sketch.** Membership follows from the fact that the search space, which is bounded as in Proposition 13, can be explored using binary search and WEAK IMPLEMENTATION as an oracle. More precisely, we can find the smallest budget  $\beta$  such that  $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$  by checking every possible value for  $\beta$ , which lies between 0 and  $2^n$ , where  $n$  is the length of the encoding of the instance. Since, due to the binary search routine, we need logarithmically many calls to the NP oracle (*i.e.*, to WEAK IMPLEMENTATION), in the end we have a searching procedure that would run in polynomial time. For the lower bound, we reduce from TSP COST (the optimal travelling salesman problem), which is  $\text{FP}^{\text{NP}}$ -complete [28]. ◀

► **Corollary 17.** *EXACT-WI with GR(1) specifications is  $\text{D}^{\text{P}}$ -complete.*

**Proof.** For membership, observe that an input is a “yes” instance of EXACT-WI if and only if it is a “yes” instance of WEAK IMPLEMENTATION *and* a “yes” instance of WEAK IMPLEMENTATION COMPLEMENT (the problem where one asks whether  $\text{WI}(\mathcal{G}, \varphi, \beta) = \emptyset$ ). Since the former problem is in NP and the latter problem is in coNP, membership in  $\text{D}^{\text{P}}$  follows. For the lower bound, we use the same reduction technique as in Theorem 16, and reduce from EXACT TSP, a problem known to be  $\text{D}^{\text{P}}$ -hard [28, 29]. ◀

Following [27], we may naturally ask whether the optimal solution given by OPT-WI is unique. We call this problem UOPT-WI. For some fixed budget  $\beta$ , it may be the case that for two subsidy schemes  $\kappa, \kappa' \in \text{WI}(\mathcal{G}, \varphi, \beta)$  – we assume the implementation is efficient – we have  $\kappa \neq \kappa'$  and  $\text{cost}(\kappa) = \text{cost}(\kappa')$ . With LTL specifications, it is not difficult to see that we can solve UOPT-WI in polynomial space. Therefore, we have the following result.

► **Corollary 18.** *UOPT-WI with LTL specifications is PSPACE-complete.*

For GR(1) specifications, we reason about UOPT-WI using the following procedure:

1. Find the exact budget using binary search and WEAK IMPLEMENTATION as an oracle;
2. Use an NP oracle once to guess two distinct subsidy schemes with precisely this budget; if no such subsidy schemes exist, return “yes”; otherwise, return “no”.

The above decision procedure clearly is in  $\Delta_2^P$  (for the upper bound). Furthermore, since Theorem 16 implies  $\Delta_2^P$ -hardness [22] (for the lower bound), we have the following corollary.

► **Corollary 19.** *UOPT-WI with GR(1) specifications is  $\Delta_2^P$ -complete.*

## 6.2 Optimality and Uniqueness in the Strong Domain

In this subsection, we study the same problems as in the previous subsection but with respect to the STRONG IMPLEMENTATION variant of the equilibrium design problem. We first formally define the problems of interest and then present the two first results.

► **Definition 20** (OPT-SI). *Given a game  $\mathcal{G}$  and a specification formula  $\varphi$ :*

*What is the optimum budget  $\beta$  such that  $\text{SI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ ?*

► **Definition 21** (EXACT-SI). *Given a game  $\mathcal{G}$ , a specification formula  $\varphi$ , and an integer  $b$ :*

*Is  $b$  equal to the optimum budget for  $\text{SI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ ?*

For the same reasons discussed in the weak versions of these two problems, we can prove the following two results with respect to games with LTL specifications.

► **Theorem 22.** *OPT-SI with LTL specifications is FSPACE-complete.*

► **Corollary 23.** *EXACT-SI with LTL specifications is PSPACE-complete.*

For GR(1) specifications, observe that using the same arguments for the upper-bound of OPT-WI with GR(1) specifications, we obtain the upper-bound for OPT-SI with GR(1) specifications. Then, it follows that OPT-SI is in  $\text{FP}^{\Sigma_2^P}$ . For hardness, we define an  $\text{FP}^{\Sigma_2^P}$ -complete problem, namely WEIGHTED MINQSAT<sub>2</sub>. Recall that in QSAT<sub>2</sub> we are given a Boolean 3DNF formula  $\psi(\mathbf{x}, \mathbf{y})$  and sets  $\mathbf{x} = \{x_1, \dots, x_n\}$ ,  $\mathbf{y} = \{y_1, \dots, y_m\}$ , with a set of terms  $T = \{t_1, \dots, t_k\}$ . Define WEIGHTED MINQSAT<sub>2</sub> as follows. Given  $\psi(\mathbf{x}, \mathbf{y})$  and a weight function  $c : \mathbf{x} \rightarrow \mathbb{Z}^{\geq}$ , WEIGHTED MINQSAT<sub>2</sub> is the problem of finding an assignment  $\vec{\mathbf{x}} \in \{0, 1\}^n$  with the least total weight such that  $\psi(\mathbf{x}, \mathbf{y})$  is true for every  $\vec{\mathbf{y}} \in \{0, 1\}^m$ . Observe that WEIGHTED MINQSAT<sub>2</sub> generalises MINQSAT<sub>2</sub>, which is known to be  $\text{FP}^{\Sigma_2^P[\log n]}$ -hard [12], *i.e.*, MINQSAT<sub>2</sub> is an instance of WEIGHTED MINQSAT<sub>2</sub>, where all weights are 1.

► **Theorem 24.** *WEIGHTED MINQSAT<sub>2</sub> is  $\text{FP}^{\Sigma_2^P}$ -complete.*

**Proof.** Membership follows from the upper-bound of MINQSAT<sub>2</sub> [12]: since we have an exponentially large input with respect to that of MINQSAT<sub>2</sub>, by using binary search we will need polynomially many calls to the  $\Sigma_2^P$  oracle. Hardness is immediate [12]. ◀

Now that we have an  $\text{FP}^{\Sigma_2^P}$ -hard problem in our hands, we can proceed to determine the complexity class of OPT-SI with GR(1) specifications. For the upper bound we one can use arguments analogous to those in Theorem 16. For the lower bound, one can reduce from WEIGHTED MINQSAT<sub>2</sub>. Formally, we have:

► **Theorem 25.** *OPT-SI with GR(1) specifications is  $\text{FP}^{\Sigma_2^P}$ -complete.*

► **Corollary 26.** *EXACT-SI with GR(1) specifications is  $\text{D}_2^P$ -complete.*

**Proof.** Membership follows from the fact that an input is a “yes” instance of EXACT-SI (with GR(1) specifications) if and only if it is a “yes” instance of STRONG IMPLEMENTATION and a “yes” instance of STRONG IMPLEMENTATION COMPLEMENT, the decision problem where we ask  $\text{SI}(\mathcal{G}, \varphi, \beta) = \emptyset$  instead. The lower bound follows from the hardness of STRONG IMPLEMENTATION and STRONG IMPLEMENTATION COMPLEMENT problems, which immediately implies  $\text{D}_2^P$ -hardness [1, Lemma 3.2]. ◀

Furthermore, analogous to UOPT-WI, we also have the following corollaries.

► **Corollary 27.** *UOPT-SI with LTL specifications is PSPACE-complete.*

► **Corollary 28.** *UOPT-SI with GR(1) specifications is  $\Delta_3^P$ -complete.*

## 7 Conclusions & Related and Future Work

### Equilibrium design vs. mechanism design – connections with Economic theory

Although equilibrium design is closely related to mechanism design, as typically studied in game theory [21], the two are not exactly the same. Two key features in mechanism design are the following. Firstly, in a mechanism design problem, the designer is not given a game structure, but instead is asked to provide one; in that sense, a mechanism design problem is closer to a rational synthesis problem [14, 16]. Secondly, in a mechanism design problem, the designer is only interested in the game’s outcome, which is given by the payoffs of the players in the game; however, in equilibrium design, while the designer is interested in the payoffs of the players as these may need to be perturbed by its budget, the designer is also interested – and in fact primarily interested – in the satisfaction of a temporal logic goal specification, which the players in the game do not take into consideration when choosing their individual rational choices; in that sense, equilibrium design is closer to rational verification [17] than to mechanism design. Thus, equilibrium design is a new computational problem that sits somewhere in the middle between mechanism design and rational verification/synthesis. Technically, in equilibrium design we go beyond rational synthesis and verification through the additional design of subsidy schemes for incentivising behaviours in a concurrent and multi-agent system, but we do not require such subsidy schemes to be incentive compatible mechanisms, as in mechanism design theory, since the principal may want to reward only a group of players in the game so that its temporal logic goal is satisfied, while rewarding other players in the game in an unfair way – thus, leading to a game with a suboptimal social welfare measure. In this sense, equilibrium design falls short with respect to the more demanding social welfare requirements often found in mechanism design theory.

### Equilibrium design vs. rational verification – connections with Computer science

Typically, in rational synthesis and verification [14, 16, 17, 23] we want to check whether a property is satisfied on some/every Nash equilibrium computation run of a reactive, concurrent, and multi-agent system. These verification problems are primarily concerned with qualitative properties of a system, while assuming rationality of system components. However, little attention is paid to quantitative properties of the system. This drawback has been recently identified and some work has been done to cope with questions where both

qualitative and quantitative concerns are considered [3, 6, 9, 10, 11, 18, 20, 34]. Equilibrium design is new and different approach where this is also the case. More specifically, as in a mechanism design problem, through the introduction of an external principal – the designer in the equilibrium design problem – we can account for overall qualitative properties of a system (the principal’s goal given by an LTL or a GR(1) specification) as well as for quantitative concerns (optimality of solutions constrained by the budget to allocate additional rewards/resources). Our framework also mixes qualitative and quantitative features in a different way: while system components are only interested in maximising a quantitative payoff, the designer is primarily concerned about the satisfaction of a qualitative (logic) property of the system, and only secondarily about doing it in a quantitatively optimal way.

### Equilibrium design vs. repair games and normative systems – connections with AI

In recent years, there has been an interest in the analysis of rational outcomes of multi-agent systems modelled as multi-player games. This has been done both with modelling and with verification purposes. In those multi-agent settings, where AI agents can be represented as players in a multi-player game, a focus of interest is on the analysis of (Nash) equilibria in such games [8, 17]. However, it is often the case that the existence of Nash equilibria in a multi-player game with temporal logic goals may not be guaranteed [16, 17]. For this reason, there has been already some work on the introduction of desirable Nash equilibria in multi-player games [2, 30]. This problem has been studied as a repair problem [2] in which either the preferences of the players (given by winning conditions) or the actions available in the game are modified; the latter one also being achieved with the use of normative systems [30]. In equilibrium design, we do not directly modify the preferences of agents in the system, since we do not alter their goals or choices in the game, but we indirectly influence their rational behaviour by incentivising players to visit, or to avoid, certain states of the overall system. We studied how to do this in an (individually) optimal way with respect to the preferences of the principal in the equilibrium design problem. However, this may not always be possible, for instance, because the principal’s temporal logic specification goal is just not achievable, or because of constraints given by its limited budget.

### Future work: social welfare requirements and practical implementation

As discussed before, a key difference with mechanism design is that social welfare requirements are not considered [25]. However, a benevolent principal might not see optimality as an individual concern, and instead consider the welfare of the players in the design of a subsidy scheme. In that case, concepts such as the *utilitarian social welfare* may be undesirable as the social welfare maximising the payoff received by players might allocate all the budget to only one player, and none to the others. A potentially better option is to improve fairness in the allocation of the budget by maximising the *egalitarian social welfare*. Finally, given that the complexity of equilibrium design is much better than that of rational synthesis/verification, we should be able to have efficient implementations, for instance, as an extension of EVE [19].

---

### References

- 1 Gadi Aleksandrowicz, Hana Chockler, Joseph Y. Halpern, and Alexander Ivrii. The Computational Complexity of Structure-based Causality. *J. Artif. Int. Res.*, 58(1):431–451, January 2017. URL: <http://dl.acm.org/citation.cfm?id=3176764.3176775>.
- 2 S. Almagor, G. Avni, and O. Kupferman. Repairing Multi-Player Games. In *CONCUR*, volume 42 of *LIPICs*, pages 325–339. Schloss Dagstuhl, 2015.

- 3 S. Almagor, O. Kupferman, and G. Perelli. Synthesis of Controllable Nash Equilibria in Quantitative Objective Games. In *IJCAI*, pages 35–41, 2018.
- 4 B. Aminof, V. Malvone, A. Murano, and S. Rubin. Graded Strategy Logic: Reasoning about Uniqueness of Nash Equilibria. In *AAMAS*, pages 698–706. ACM, 2016.
- 5 R. Bloem, B. Jobstmann, N. Piterman, A. Pnueli, and Y. Sa’ar. Synthesis of Reactive(1) designs. *Journal of Computer and System Sciences*, 78(3):911–938, 2012.
- 6 A. Bohy, V. Bruyère, E. Filiot, and J. Raskin. Synthesis from LTL Specifications with Mean-Payoff Objectives. In *TACAS*, pages 169–184, 2013.
- 7 U. Boker, K. Chatterjee, T. A. Henzinger, and O. Kupferman. Temporal Specifications with Accumulative Values. *ACM Transactions on Computational Logic*, 15(4):27:1–27:25, 2014. doi:10.1145/2629686.
- 8 P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Pure Nash Equilibria in Concurrent Deterministic Games. *Logical Methods in Computer Science*, 11(2), 2015.
- 9 K. Chatterjee and L. Doyen. Energy parity games. *Theoretical Computer Science*, 458:49–60, 2012.
- 10 K. Chatterjee, L. Doyen, T. Henzinger, and J. Raskin. Generalized Mean-payoff and Energy Games. In *FSTTCS*, pages 505–516, 2010. doi:10.4230/LIPIcs.FSTTCS.2010.505.
- 11 K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Mean-Payoff Parity Games. In *LICS*, pages 178–187. IEEE Computer Society, 2005.
- 12 H. Chockler and J. Halpern. Responsibility and Blame: A Structural-Model Approach. *Journal of Artificial Intelligence Research*, 22:93–115, 2004.
- 13 Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Toruńczyk. Energy and Mean-Payoff Games with Imperfect Information. In Anuj Dawar and Helmut Veith, editors, *Computer Science Logic*, pages 260–274, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- 14 D. Fisman, O. Kupferman, and Y. Lustig. Rational Synthesis. In *TACAS*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010.
- 15 J. Gutierrez, P. Harrenstein, G. Perelli, and M. Wooldridge. Nash Equilibrium and Bisimulation Invariance. In *CONCUR*, volume 85 of *LIPIcs*, pages 17:1–17:16. Schloss Dagstuhl, 2017.
- 16 J. Gutierrez, P. Harrenstein, and M. Wooldridge. Iterated Boolean Games. *Information and Computation*, 242:53–79, 2015.
- 17 J. Gutierrez, P. Harrenstein, and M. Wooldridge. From Model Checking to Equilibrium Checking: Reactive Modules for Rational Verification. *Artificial Intelligence*, 248:123–157, 2017.
- 18 J. Gutierrez, A. Murano, G. Perelli, S. Rubin, and M. Wooldridge. Nash Equilibria in Concurrent Games with Lexicographic Preferences. In *IJCAI*, pages 1067–1073, 2017. doi:10.24963/ijcai.2017/148.
- 19 J. Gutierrez, M. Najib, G. Perelli, and M. Wooldridge. EVE: A Tool for Temporal Equilibrium Analysis. In *ATVA*, volume 11138 of *LNCS*, pages 551–557. Springer, 2018.
- 20 J. Gutierrez, M. Najib, G. Perelli, and M. Wooldridge. On Computational Tractability for Rational Verification. In *IJCAI*, 2019. To appear.
- 21 L. Hurwicz and S. Reiter. *Designing Economic Mechanisms*. Cambridge University Press, 2006.
- 22 M. Krentel. The Complexity of Optimization Problems. *Journal of Computer and System Sciences*, 36(3):490–509, 1988. doi:10.1016/0022-0000(88)90039-6.
- 23 O. Kupferman, G. Perelli, and M. Y. Vardi. Synthesis with Rational Environments. *Annals of Mathematics and Artificial Intelligence*, 78(1):3–20, 2016.
- 24 M. Wooldridge and U. Endriss and S. Kraus and J. Lang. Incentive engineering for Boolean games. *Artificial Intelligence*, 195:418–439, 2013. doi:10.1016/j.artint.2012.11.003.
- 25 M. Maschler, E. Solan, and S. Zamir. *Game Theory*. Cambridge University Press, 2013.
- 26 M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.



## 22:16 Equilibrium Design for Concurrent Games

- 27 C. Papadimitriou. On the Complexity of Unique Solutions. *Journal of the ACM*, 31(2):392–400, 1984. doi:10.1145/62.322435.
- 28 C. Papadimitriou. *Computational complexity*. Addison-Wesley, Reading, Massachusetts, 1994.
- 29 C. Papadimitriou and M. Yannakakis. The Complexity of Facets (and some Facets of Complexity). *Journal of Computer and System Sciences*, 28(2):244–259, 1984.
- 30 G. Perelli. Enforcing Equilibria in Multi-Agent Systems. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '19, pages 188–196, 2019. URL: <http://dl.acm.org/citation.cfm?id=3306127.3331692>.
- 31 A. Pnueli. The Temporal Logic of Programs. In *FOCS*, pages 46–57. IEEE, 1977.
- 32 A. Pnueli and R. Rosner. On the Synthesis of a Reactive Module. In *POPL*, pages 179–190. ACM Press, 1989.
- 33 M. Ummels and D. Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games. In *CONCUR*, pages 482–496, 2011. doi:10.1007/978-3-642-23217-6\_32.
- 34 Y. Velner, K. Chatterjee, L. Doyen, T. Henzinger, A. Rabinovich, and J. Raskin. The Complexity of Multi-Mean-Payoff and Multi-Energy Games. *Information and Computation*, 241:177–196, 2015.
- 35 U. Zwick and M. Paterson. The Complexity of Mean Payoff Games on Graphs. *Theoretical Computer Science*, 158(1):343–359, 1996. doi:10.1016/0304-3975(95)00188-3.