





DESIGNING EQUILIBRIA IN CONCURRENT GAMES WITH SOCIAL WELFARE AND TEMPORAL LOGIC CONSTRAINTS *

JULIAN GUTIERREZ ^a, MUHAMMAD NAJIB ^b, GIUSEPPE PERELLI ^c,
AND MICHAEL WOOLDRIDGE ^d

^a University of Sussex
e-mail address: J.Gutierrez@sussex.ac.uk

^b School of Mathematical and Computer Sciences, Heriot-Watt University
e-mail address: m.najib@hw.ac.uk

^c Department of Computer Science, Sapienza University of Rome
e-mail address: perelli@di.uniroma1.it

^d Department of Computer Science, University of Oxford
e-mail address: mjw@cs.ox.ac.uk

ABSTRACT. In game theory, *mechanism design* is concerned with the design of incentives so that a desirable outcome will be achieved under the assumption that players act rationally. In this paper, we explore the concept of *equilibrium design*, where incentives are designed to obtain a desirable equilibrium that satisfies a specific temporal logic property. Our study is based 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. We consider system specifications given by LTL and GR(1) formulae, and show that designing incentives to ensure that a given temporal logic property is satisfied on some/every Nash equilibrium of the game can be achieved in PSPACE for LTL properties and in NP/ Σ_2^P for GR(1) specifications. We also examine the complexity of related decision and optimisation problems, such as optimality and uniqueness of solutions, as well as considering social welfare, and show that the complexities of these problems lie within the polynomial hierarchy. Equilibrium design can be used as an alternative solution to rational synthesis and verification problems for concurrent games with mean-payoff objectives when no solution exists or as a technique to repair concurrent games with undesirable Nash equilibria in an optimal way.

Key words and phrases: game theory; concurrency; temporal logic; verification; model checking; Nash equilibrium; social welfare.

* Preliminary version appeared in the *Proceedings of 30th International Conference on Concurrency Theory (CONCUR 2019)* [GNPW19a].

$y \backslash x$	0	1	2
0	○		■
1			

Figure 1: Graphical representation from Example 1.1.

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.*, [AKP18, AMMR16, BBMU15, FKL10, GHPW17, GHW17, KPV16]). 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 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 (e.g., Nash) equilibrium. These problems are typically computationally complex, since they subsume temporal logic synthesis [PR89]. 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 [FKL10, GHW15, GHW17], rather than only PSPACE-complete as it is the case for model checking. This represents a major computational barrier for the practical analysis and automated verification of reactive, concurrent, and multi-agent systems modelled as multi-player games.

A classic problem in game theory is that individually rational choices can result in outcomes that are highly undesirable, and concurrent games also fall prey to this problem. To illustrate this, consider the scenario in the following example.

Example 1.1. Consider a system with two robot agents operating in an environment modelled as a 3×2 grid world. Initially, the robots are located at two corners, as shown in Figure 1.1. Each robot can move one square horizontally, vertically, or diagonally (similar to the way a king moves in chess). Each move costs the robot 1 unit of energy and incurs a payment of -1 . The task of robot \circ (resp. \blacksquare) is to visit square $(2, 1)$ (resp. $(0, 1)$), for example, to deliver parcels. To model this objective, we give the robots a payment of 3 when they reach their target squares. We assume that at each time-step, each robot must make a move and cannot remain in the same position for two consecutive time-steps. Furthermore, each robot wants to maximise the sum of payments it receives.

Suppose that having two robots occupying the same grid square (i.e., having the same coordinates) is considered undesirable from a global perspective because it increases the likelihood of collisions. To maximise their total payoffs, the robots will choose routes that minimise the number of steps needed to reach their target squares. Observe that the minimum number of steps required for the robots to reach their respective target squares is 2, and there are two routes that achieve this: via $(1, 0)$ and $(1, 1)$. However, if both robots

choose the same shortest route (e.g., via $(1, 1)$), this results in an undesirable outcome. Moreover, since the shortest routes correspond to the best strategies for each robot, these outcomes are considered stable from a game-theoretic perspective.

These concerns have motivated the development of techniques for modifying games, in order to avoid undesirable equilibria, or to facilitate desirable equilibria. In game theory, the field of *mechanism design* is concerned with designing a game such that, if players behave rationally, then a desired outcome will be obtained [OR94]. Direct incentives, for example in the form of taxation or subsidy, are probably the most important tools used in mechanism design.

This paper explores the design of incentive schemes for concurrent games so as to achieve a desired outcome, a concept we refer to as *equilibrium design*. Specifically, we use *reward schemes* to incentivise players so that the Nash equilibria of the game satisfy the desired property. In our model, agents are represented as concurrently executing processes that operate synchronously. Each agent receives an integer payoff for every state visited by the overall system. The total payoff an agent receives over an infinite computation path is defined as the *mean-payoff* over that path. While agents naturally seek to maximise their individual mean-payoffs, the designer of the reward scheme aims to ensure that a specific temporal logic formula is satisfied on some or all Nash equilibria of the resulting game.

With this model, we assume that the designer – an external principal – has a finite budget that is available for designing reward schemes, 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 rewards 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 [M. 13], 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 [BJP⁺12], give algorithms for each case, and classify the complexity of the problem. While LTL is a natural language for the specification

	LTL Spec.	GR(1) Spec.
WEAK IMPL.	PSPACE-c (Thm. 4.4)	NP-c (Thm. 4.6)
STRONG IMPL.	PSPACE-c (Cor. 5.2)	Σ_2^P -c (Thm. 5.3)
OPT-WI	FPSPACE-c (Thm. 6.4)	FP^{NP} -c (Thm. 6.6)
OPT-SI	FPSPACE-c (Thm. 6.12)	$\text{FP}^{\Sigma_2^P}$ -c (Thm. 6.15)
EXACT-WI	PSPACE-c (Cor. 6.5)	D^P -c (Cor. 6.7)
EXACT-SI	PSPACE-c (Cor. 6.13)	D_2^P -c (Cor. 6.16)
UOPT-WI	PSPACE-c (Cor. 6.8)	Δ_2^P -c (Cor. 6.9)
UOPT-SI	PSPACE-c (Cor. 6.17)	Δ_3^P -c (Cor. 6.18)
UT-WI	PSPACE-c (Thm. 7.3)	NP-c (Thm. 7.5)
UT-SI	PSPACE-c (Thm. 7.3)	Σ_2^P -c (Thm. 7.5)
ET-WI	PSPACE-c (Thm. 7.4)	NP-c (Thm. 7.6)
ET-SI	PSPACE-c (Thm. 7.4)	Σ_2^P -c (Thm. 7.6)

Table 1: Summary of main complexity results.

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 ω -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. We also examine a setting in which a (benevolent) principal considers the welfare of the players in the design of a reward scheme. To capture this setting, we introduce two concepts: *utilitarian* and *egalitarian* social welfare measures. We show that, while the problems associated with LTL specifications are in PSPACE (or FPSPACE), the ones with GR(1) specifications lie within the polynomial hierarchy, thus making them potentially amenable to more efficient practical implementations. Table 1 summarises the main computational complexity results in the paper.

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

- Section 2 presents the relevant background on games, logic, and Nash equilibrium.
- In Section 3 we formalise the concept of reward schemes.
- In Section 4 and 5 we describe the main problems of interest and present the proofs to obtain tight computational complexity bounds.
- In Section 6 we study variations of the main problems, including optimality and uniqueness of solutions, and show their respective computational complexity classes.
- In Section 7 we consider two of the most important social welfare measures, and examine the related computational problems.
- In Section 8 we conclude, discuss related work, and propose some directions for further research.

2. PRELIMINARIES

Complexity Classes. Here we briefly describe the different complexity classes used in this paper. We assume that the reader is familiar with the classes NP, PSPACE and the notation for complexities relative to an oracle within the polynomial hierarchy [Pap94]. In particular, we assume the following

- $\Sigma_0^P = \Pi_0^P = \Delta_0^P = P = \Delta_1^P$;
- $\Sigma_{i+1}^P = \text{NP}^{\Sigma_i^P}$;
- $\Pi_{i+1}^P = \text{co-NP}^{\Sigma_i^P}$;
- $\Delta_{i+1}^P = \text{P}^{\Sigma_i^P}$.

For instance, $\Sigma_2^P = \text{NP}^{\text{NP}}$ is the class of problems that can be solved in polynomial time by a non-deterministic Turing machine that can invoke an oracle to solve another NP problem.

For a given decision problem in the complexity class \mathcal{C} , \mathcal{FC} denotes the complexity class of the corresponding *function* problem. Consider, for example, the class NP and a problem $\mathcal{P} \in \text{NP}$. The problem of finding a solution to an instance of \mathcal{P} is in FNP.

Finally, D_i^P denotes the class of languages that are the intersection of a language in Σ_i^P and a language in Π_i^P (note that this is not the same as $\Sigma_i^P \cap \Pi_i^P$). Intuitively, this corresponds to the class of problems that require two consecutive and independent calls to a Σ_i^P procedure and a Π_i^P procedure. This is typically used to refer to problems in which the solution is, in a sense, *unique* or *exact*. For instance, consider the problem:

$$\text{EXACT-CLIQUE} = \{(G, k) : \text{the largest clique in graph } G \text{ is of size } k\}.$$

This requires solving CLIQUE for k to decide whether a clique of size k *exists*, and solving CO-CLIQUE for $k + 1$ to decide whether a clique of size $k + 1$ does *not exist*. For more in-depth presentation of the class D_i^P , see [Pap94, PY84, ACHI17].

Linear Temporal Logic. LTL [Pnu77] extends classical propositional logic with two operators, **X** (“next”) and **U** (“until”), that can be used to express properties of infinite paths. The syntax of LTL is defined with respect to a set AP of atomic propositions by the following grammar:

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

where $p \in \text{AP}$. As is conventional in the LTL literature, we introduce some further classical and temporal operators via the following equivalences:

$$\varphi_1 \wedge \varphi_2 \equiv \neg(\neg\varphi_1 \vee \neg\varphi_2) \quad \varphi_1 \rightarrow \varphi_2 \equiv \neg\varphi_1 \vee \varphi_2 \quad \mathbf{F}\varphi \equiv \top \mathbf{U} \varphi \quad \mathbf{G}\varphi \equiv \neg\mathbf{F}\neg\varphi$$

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

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

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

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 [BJP⁺12]:

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

where each subformula ψ_i and φ_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 , Ac , and 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* $\pi = (s_0, \vec{a}^0), (s_1, \vec{a}^1), \dots$ is an infinite sequence in $(\text{St} \times \vec{\text{Ac}})^\omega$ such that $\text{tr}(s_k, \vec{a}^k) = s_{k+1}$ for all k . In particular, \vec{a}_i^k is the action of player i in step k . Sometimes, we call the single iteration (s_k, \vec{a}^k) a *configuration*.

Paths are generated in the arena by each player i selecting a *strategy* σ_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 Str_i be the set of strategies for player i . Note that this definition implies that strategies have *perfect information* 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
- $q_i^{k+1} = \delta_i(s_i^k, (\tau_1(q_1^k), \dots, \tau_n(q_n^k)))$, 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 feature that is essential for game theory: *preferences*, which give games their strategic structure. A *multi-player game* is obtained from an arena A by associating each player with a *goal*, which represents that player's preferences. We consider multi-player games with **mp** goals. A multi-player **mp** game is a tuple $\mathcal{G} = \langle A, (\mathbf{w}_i)_{i \in \mathbb{N}} \rangle$, 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 π in A induces a sequence $\lambda(\pi) = \lambda(s_0)\lambda(s_1) \cdots$ of sets of atomic propositions; if, in addition, A is the arena of an **mp** game, then, for each player i , the sequence $\mathbf{w}_i(\pi) = \mathbf{w}_i(s_0)\mathbf{w}_i(s_1) \cdots$ of weights is also induced. Unless stated otherwise, for a game \mathcal{G} and a path π in it, the payoff of player i is $\text{pay}_i(\pi) = \text{mp}(\mathbf{w}_i(\pi))$.

Nash equilibrium. Using payoff functions, we can define the game-theoretic concept of Nash equilibrium [OR94]. 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 σ'_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} . Observe that in a given game, there may be more than one Nash equilibria, and as such, different equilibrium outcomes may behave differently—some of which may not be desirable. This was intuitively illustrated in Example 1.1. To frame this problem more appropriately within our framework, consider the following example modified from Example 1.1.

Example 2.1. Consider the same setting as in Example 1.1, but with the tasks of the robots set as follows. The task for robot \circ (resp. \blacksquare) is to visit $(0, 0)$ and $(2, 1)$ (resp. $(2, 0)$ and $(0, 1)$) in an alternating fashion infinitely often. Each movement costs 1 unit of energy and gives a payoff of -1. When a robot reaches a target corner after visiting the other one, it gets 3. Each robot wants to maximise the mean-payoff of its infinite run.




Figure 2: Transition system representing the movements and payoff function for robot \circ .

We can model the movements and payoff function of robot \circ as a transition system in Figure 2. The vertices are marked with $(x, y)_f$, where $f \in \{0, 1\}$ is a flag to mark the last corner robot \circ visited (0 for $(0, 0)$ and 1 for $(2, 1)$.) We set the payoff function for robot \circ as follows. $w_{\circ}(e_i) = -1$ for $i \in \{1, 2, 7, 8, 9, 10\}$ and $w_{\circ}(e_j) = 2$ for $j \in \{3, 4, 5, 6\}$ ¹. We can model the movements and payoff function of robot \blacksquare in a similar manner. The graphical representation of the game can be obtained by taking the cross product of the two transition systems.

From the game we have obtained, we can see that there are many Nash equilibrium runs. We define “bad” Nash equilibrium run as one in which the robots occupy the same location simultaneously and infinitely often. From the set of all Nash equilibrium runs, some are bad and some are not. For example, consider a run where robot \circ ’s sequence of moves is $((0, 0)_0(1, 0)_0(2, 1)_1(1, 0)_1)^\omega$ and robot \blacksquare ’s is $((2, 0)_0(1, 1)_0(0, 1)_1(1, 1)_1)^\omega$. This is a Nash equilibrium run since both robots get the mean-payoff of $\frac{1}{2}$ and cannot obtain better rewards by changing their actions. Furthermore, it is not a bad Nash equilibrium run, since the robots never simultaneously occupy the same position. Now, consider a different scenario where robot \blacksquare plays the same strategy as in the previous run, but robot \circ ’s sequence of moves is $((0, 0)_0(1, 1)_0(2, 1)_1(1, 1)_1)^\omega$. The mean-payoffs for both robots are still $\frac{1}{2}$, making this a Nash equilibrium run. However, in this run, the robots will occupy $(1, 1)$ simultaneously and infinitely often, making it a bad Nash equilibrium run.

From a system design perspective, we want to eliminate such bad Nash equilibrium runs. One way to achieve this is by modifying the payoff function for each robot so that the resulting set of equilibria does not include bad runs. We can do this by providing rewards to the robots in order to “nudge” them into taking certain paths. Consider again the payoff function of robot \circ shown in Figure 2. Suppose we provide rewards for e_3 and

¹For clarity in presentation, we have excluded locations $(0, 1)$ and $(2, 0)$ because they are not part of any shortest routes. Additionally, we have placed the payoffs on the edges rather than the vertices. However, it is easy to transform the transition system and push the payoffs to the vertices.

e_5 , 1 unit of payoff each. Thus we have $w_{\circ}(e_3) = 3$, $w_{\circ}(e_5) = 3$. Now consider the run of robot \circ as follows: $h((0,0)_0(1,0)_0(2,1)_1(1,0)_1)^\omega$, where h is a finite prefix. This run is a Nash equilibrium that results in a mean-payoff of 1 for robot \circ . In fact, in *every* Nash equilibrium, robot \circ 's run corresponds to this type of run, since otherwise, the robot will get a mean-payoff of < 1 . Similarly, we can design a reward scheme for robot \blacksquare that will result in \blacksquare always choosing the run $h((2,0)_0(1,1)_0(0,1)_1(1,1)_1)^\omega$ in every Nash equilibrium. By combining these payoff functions, we obtain a new payoff function that prevents the system from getting stuck in bad equilibria.

In the next section, we formalise the problem of designing payoff functions using rewards to achieve desirable Nash equilibrium runs.

3. FROM MECHANISM DESIGN TO EQUILIBRIUM DESIGN

We now describe a method for modifying the payoff functions of players in a given game to achieve desirable Nash equilibrium runs. As discussed in the introduction, this problem is closely related to the well-known concept of *mechanism design* in game theory. Consider a system with multiple agents, represented by the set N . Each agent $i \in N$ aims 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 φ that it wants the system to satisfy. To accomplish this, the principal seeks to incentivise the agents to act collectively and rationally to bring about φ . In our model, incentives are given by *reward schemes* and goals by temporal logic formulae.

Reward Schemes: A reward scheme defines *additional* imposed payoff 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 reward scheme as they see fit. Since agents will seek to maximise their overall rewards, the principal can incentivise agents to visit certain states and avoid others. If the reward scheme is designed correctly, the agents are incentivised to choose a strategy profile $\vec{\sigma}$ such that $\pi(\vec{\sigma}) \models \varphi$. Formally, we model a reward scheme as a function $\kappa : (N \rightarrow \text{St}) \rightarrow \mathbb{N}$, where the intended interpretation is that $\kappa(i)(s)$ is the reward 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 reward for player i .

Notice that having an unlimited fund for a reward scheme would make some problems trivial, as the principal can always incentivise players to satisfy φ (provided that there is a path in A satisfying φ). 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 reward scheme κ as follows.

Definition 3.1. Given a game \mathcal{G} and reward scheme κ , we define

$$\text{cost}(\kappa) = \sum_{i \in N} \sum_{s \in \text{St}} \kappa_i(s).$$

We say that a reward scheme κ is *admissible* if it does not exceed the budget β , that is, if $\text{cost}(\kappa) \leq \beta$. Let $\mathcal{K}(\mathcal{G}, \beta)$ denote the set of admissible reward 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}, κ) to denote the resulting game after the application of reward scheme κ on game \mathcal{G} . Formally, we define the application of some reward scheme on a game as follows.

Definition 3.2. Given a game $\mathcal{G} = \langle A, (\mathbf{w}_i)_{i \in \mathbb{N}} \rangle$ and an admissible reward scheme κ , we define $(\mathcal{G}, \kappa) = \langle A, (\mathbf{w}'_i)_{i \in \mathbb{N}} \rangle$, where $\mathbf{w}'_i(s) = \mathbf{w}_i(s) + \kappa_i(s)$, 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 reward scheme that will incentivise players to collectively choose a rational outcome (a Nash equilibrium) that satisfies its temporal logic goal φ . We call this problem *equilibrium design*. Following [M. 13], 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.

Remark 3.3. For the rest of the paper, we assume that weights and reward schemes are using a binary representation for their values. This is the standard way of considering them in the context of mean-payoff games [ZP96, UW11].

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 reward 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 reward schemes over \mathcal{G} given budget β that satisfy a formula φ in at least one path π 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 4.1 (WEAK IMPLEMENTATION). Given a game \mathcal{G} , formula φ , and budget β :

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 [UW11].

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. Note that the value $\text{pun}_i(s)$ corresponds to the one of a two-player zero-sum mean-payoff game [ZP96], where the coalition $-i = \mathbb{N} \setminus \{i\}$ is playing adversarially against i . Thus, computing $\text{pun}_i(s)$ amounts to computing the winning value of i in such two-player zero-sum mean-payoff game, which can then be done in $\text{NP} \cap \text{co-NP}$. Also, note that the coalition $\mathbb{N} \setminus \{i\}$ 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 $(s, \vec{\mathbf{a}})$ is z -secure for player i if $\text{pun}_i(\text{tr}(s, (\vec{\mathbf{a}}_{-i}, \mathbf{a}'_i))) \leq z$ 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 4.2. For every mp game \mathcal{G} and ultimately periodic path $\pi = (s_0, \vec{\mathbf{a}}_0), (s_1, \vec{\mathbf{a}}_1), \dots$, 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 $z \in \mathbb{R}^{\mathbb{N}}$, where $z_i \in \text{pun}_i(\mathcal{G})$ such that, for every $i \in \mathbb{N}$

- (a) for all $k \in \mathbb{N}$, the pair $(s_k, \bar{\mathbf{a}}^k)$ is z_i -secure for i , and
- (b) $z_i \leq \text{pay}_i(\pi)$.

Proof. For (1) implies (2): Let z_i be the largest value player i can get by deviating from π . Let $k \in \mathbb{N}$ be such that $z_i = \text{pun}_i(\text{tr}(s_k, (\bar{\mathbf{a}}_{-i}, \mathbf{a}'_i)))$. Suppose further that $\text{pay}_i(\pi) < z_i$. Thus, player i would deviate at s_k , which is a contradiction to π being a path induced by a Nash equilibrium.

For (2) implies (1): Define strategy profile $\vec{\sigma}$ that follows π as long as no-one has deviated from π . In such a case where player i deviates on the k -th iteration, the strategy profile $\vec{\sigma}_{-i}$ starts playing the z_i -secure strategy for player i that guarantees the payoff of player i to be less than z_i . Therefore, we have $\text{pay}_i(\pi(\vec{\sigma}_{-i}, \sigma'_i)) \leq z_i \leq \text{pay}_i(\pi)$, for every possible strategy σ'_i of player i (the second inequality is due to condition 2(b)). Thus, there is no beneficial deviation for player i and π is a path induced by a Nash equilibrium. Indeed, by contradiction, assume for player i that a strategy σ'_i is a beneficial deviation from the strategy profile $\vec{\sigma}$. Then, we would have $z_i \leq \text{pay}_i(\vec{\sigma}) < \text{pay}_i(\vec{\sigma}_{-i}, \sigma'_i) < z_i$, the last inequality following from the fact that $\vec{\sigma}_{-i}$ is z_i -secure for coalition $-i$. Clearly, the sequence of inequality makes it a contradiction. \square

The characterisation of Nash Equilibria provided in Theorem 4.2 will allow us to turn the WEAK IMPLEMENTATION problem into a *path finding* problem over (\mathcal{G}, κ) . On the other hand, with respect to the budget β that the principal has at its disposal, the definition of reward scheme function κ implies that the size of $\mathcal{K}(\mathcal{G}, \beta)$ is bounded, and particularly, it is bounded by β and the number of agents and states in the game \mathcal{G} , in the following way.

Proposition 4.3. *Given a game \mathcal{G} with $|N|$ players and $|\text{St}|$ states and budget β , 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.

Proof. For a fixed budget b , the number of reward schemes of budget exactly b corresponds to the number of *weak compositions* of b in m parts, which is given by $\binom{b+m-1}{b}$ [S. 09]. Therefore, the number of reward schemes of budget at most β is the sum

$$|\mathcal{K}(\mathcal{G}, \beta)| = \sum_{b=0}^{\beta} \binom{b+m-1}{b}.$$

We now prove that

$$\sum_{b=0}^{\beta} \binom{b+m-1}{b} = \frac{\beta+1}{m} \binom{\beta+m}{\beta+1}.$$

By induction on β , as base case, for $\beta = 0$, we have that

$$\binom{\beta+m-1}{\beta} = 1 = \frac{\beta+1}{m} \binom{\beta+m}{\beta+1}.$$

For the inductive case, let us assume that the assertion hold for some β and let us prove for $\beta + 1$. We have the following:

$$\sum_{b=0}^{\beta+1} \binom{b+m-1}{b} = \sum_{b=0}^{\beta} \binom{b+m-1}{b} + \binom{\beta+m-1}{\beta+1} = \frac{\beta+1}{m} \binom{\beta+m}{\beta+1} + \binom{\beta+m}{\beta+1}.$$

Algorithm 1 Algorithm for WEAK IMPLEMENTATION

input: $(\mathcal{G}, \varphi, \beta)$

- 1: guess:
 - a reward scheme $\kappa \in \mathcal{K}(\mathcal{G}, \beta)$
 - a state $s \in \text{St}$ for every player $i \in \mathbb{N}$
 - punishment memoryless strategies $(\vec{\sigma}_{-1}, \dots, \vec{\sigma}_{-n})$ for all players $i \in \mathbb{N}$
- 2: Compute (\mathcal{G}, κ)
- 3: Compute the vector z where $z_i = \text{pun}_i(s)$ w.r.t. the punishment strategy $\vec{\sigma}_{-i}$
- 4: Build $(\mathcal{G}, \kappa)[z]$
- 5: **if** There exists π in $(\mathcal{G}, \kappa)[z]$ such that $\pi \models \varphi$ and $z_i \leq \text{pay}_i(\pi)$ for every player $i \in \mathbb{N}$
 - then**
 - 6: **return** YES
 - 7: **else**
 - 8: **return** NO

Therefore we have

$$\begin{aligned}
\frac{\beta+1}{m} \binom{\beta+m}{\beta+1} + \binom{\beta+m}{\beta+1} &= \binom{\beta+m}{\beta+1} \left(\frac{\beta+1}{m} + 1 \right) = \\
&= \binom{\beta+m}{\beta+1} \frac{\beta+1+m}{m} = \frac{\beta+1+m}{m} \cdot \frac{(\beta+m)!}{(\beta+1)!(\beta+m-\beta-1)!} = \\
\frac{(\beta+m+1)!}{(\beta+1)!m!} &= \frac{(\beta+m+1)!}{(\beta+1)!m!} \cdot \frac{\beta+2}{\beta+2} \cdot \frac{m}{m} = \frac{\beta+2}{m} \cdot \frac{(\beta+m+1)!}{(\beta+2)!(m-1)!} = \\
&= \frac{\beta+2}{m} \cdot \frac{(\beta+m+1)!}{(\beta+2)!(\beta+m+1-\beta-2)!} = \frac{\beta+2}{m} \binom{\beta+m+1}{\beta+2}
\end{aligned}$$

This proves the assertion. \square

From Proposition 4.3 we derive that the number of possible reward schemes is *polynomial* in the budget β 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 Algorithm 1. Note that, in Line 4, the algorithm builds a game $(\mathcal{G}, \kappa)[z]$, which is obtained from (\mathcal{G}, κ) by removing the states s such that $\text{pun}_i(s) \leq z_i$ for some player i , and transitions $(s, \vec{\alpha}_{-i})$ that are not z_i secure for player i .

Theorem 4.4. WEAK IMPLEMENTATION with LTL specifications is PSPACE-complete.

Proof. Firstly, notice that the correctness of Algorithm 1 directly follows from the characterisation provided in Theorem 4.2 and the definition of WEAK IMPLEMENTATION. Specifically, the path π in line 5 corresponds to a NE run of (\mathcal{G}, κ) . In other words, $\pi = \vec{\sigma} \in \text{NE}((\mathcal{G}, \kappa))$ because it satisfies condition (2) in Theorem 4.2. Moreover, π also satisfies the property φ . According to the definition of WEAK IMPLEMENTATION, this provides a witness for the non-emptiness of $\text{WI}(\mathcal{G}, \varphi, \beta)$. Therefore, we can conclude that Algorithm 1 correctly solves WEAK IMPLEMENTATION.

Regarding the complexity, observe the following. Since the set $\mathcal{K}(\mathcal{G}, \beta)$ is finitely bounded (Proposition 4.3), and punishment strategies only need to be memoryless, thus also finitely bounded, clearly line 1 can be guessed non-deterministically. Moreover, each of the guessed elements is of polynomial size, thus this step can be done (deterministically) in polynomial space. Line 2 clearly can be done in polynomial time. Line 3 can also be done in polynomial

time since, given $(\vec{\sigma}_{-1}, \dots, \vec{\sigma}_{-n})$, we can compute the vector $z = (z_1, \dots, z_n)$ by solving $|N|$ number of *one-player mean-payoff games*, one for each player i —this can be done in polynomial time for each i [ZP96, Thm. 6]—then set $z_i = \text{pun}_i(s)$ (see Theorem 4.2). For line 5, we will use Theorem 4.2 and consider two cases, one for LTL specifications and one for GR(1) specifications. We first consider the case with LTL specifications, and in the next subsection with GR(1) specifications. For LTL specifications, consider the formula

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

written in LTL^{Lim} [BCHK14], an extension of LTL where statements about mean-payoff values over a given weighted arena can be made.² The semantics of the temporal operators of LTL^{Lim} is just like the one for LTL over infinite computation paths $\pi = s_0, s_1, s_2, \dots$. 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 π 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 π if and only if $\text{pay}_i(\pi) = \text{mp}(\mathbf{w}_i(\pi))$ is greater or equal than constant value z_i . Formula φ_{WI} corresponds exactly to 2(b) in Theorem 4.2. Furthermore, since every path in $(\mathcal{G}, \kappa)[z]$ satisfies condition 2(a) of Theorem 4.2, every computation path of $(\mathcal{G}, \kappa)[z]$ that satisfies φ_{WI} is a witness to the WEAK IMPLEMENTATION problem.

Therefore, membership follows from the algorithm and the fact that model checking for LTL^{Lim} is PSPACE-complete [BCHK14].

Hardness follows from the fact that LTL model checking is a special case of WEAK IMPLEMENTATION. To see this, consider an instance of WEAK IMPLEMENTATION $(\mathcal{G}, \varphi, \beta)$ where $N = \{1\}$, $\beta = 0$ and for every $s \in \text{St}$, $w_1(s) = 0$. Clearly, $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ if and only if the LTL formula φ is satisfied in the underlying arena A of \mathcal{G} , which constitutes model checking φ against A . \square

Remark 4.5. Note that the formula φ_{WI} in the proof of the theorem is used to check whether there is a path satisfying the formula. We refer to this as “existential” LTL^{Lim} model checking. This notion is not directly addressed in [BCHK14], where the discussion is centered around “universal” model checking. However, one can easily be derived from the other by negating the formula and flipping the answer of model checking, all whilst remaining within PSPACE. We further note that the language LTL^{Lim} is closed under negation, and (as such) strict inequalities are also expressible. Indeed, strict inequalities are also explicitly used in [BCHK14].

Case with GR(1) specifications. One of the main bottlenecks of our algorithm to solve WEAK IMPLEMENTATION lies in line 5, where we solve an LTL^{Lim} model checking problem. To reduce the complexity of our decision procedure, we consider WEAK IMPLEMENTATION with the specification φ expressed in the GR(1) sublanguage of LTL. With this specification language, we can avoid model checking LTL^{Lim} in line 5. Indeed, with GR(1) specifications, we can solve line 5 in polynomial time. This is made possible by a linear program (LP) that we define, drawing inspiration from Kosaraju and Sullivan’s technique for detecting zero cycles [KS88]. The LP yields a solution if and only if line 5 returns true.

Theorem 4.6. WEAK IMPLEMENTATION *with GR(1) specifications is NP-complete.*

²The formal semantics of LTL^{Lim} can be found in [BCHK14]. We prefer to give only an informal description here.

Proof. For the upper bound, observe that in Algorithm 1, line 1 can be done *non-deterministically* in polynomial time. Furthermore, lines 2–4 can be done *deterministically* in polynomial time. Then, in order to solve line 5 we define a linear program of size polynomial in (\mathcal{G}, κ) having a solution if and only if there exists an ultimately periodic path π such that $z_i \leq \text{pay}_i(\pi)$ and satisfies the GR(1) specification.

To do this, first recall that φ 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}, κ) that satisfy the boolean combinations ψ_l and θ_r , respectively. Observe that property φ is satisfied over a path π if, and only if, either π visits every $V(\theta_r)$ infinitely many times or visits some of the $V(\psi_l)$ only a finite number of times.

For the game $(\mathcal{G}, \kappa)[z]$, let $W = (V, E, (w'_i)_{i \in \mathbb{N}})$ be the underlying multi-weighted graph, where $w'_i(v) = w_i(s) - z_i$ for every $i \in \mathbb{N}$, $v \in V$, and $s \in \text{St}$ such that v corresponds to s . Furthermore, we introduce a variable x_e for every edge $e \in E$, where the value x_e corresponds to the number of times that the edge e is used on a cycle. Let $\text{src}(e)$ and $\text{trg}(e)$ be the *source* and *target* of the edge e , respectively; $\text{out}(v) = \{e \in E : \text{src}(e) = v\}$; and $\text{in}(v) = \{e \in E : \text{trg}(e) = v\}$.

Consider ψ_l for some $1 \leq l \leq m$, and define the linear program $\text{LP}(\psi_l)$ with the following inequalities and equations:

- Eq1: $x_e \geq 0$ for each edge e — an edge cannot be used a negative number of times;
- Eq2: $\sum_{e \in E} x_e \geq 1$ — ensures that at least one edge is chosen;
- Eq3: for each $i \in \mathbb{N}$, $\sum_{e \in E} w'_i(\text{trg}(e))x_e \geq 0$ — this enforces that the total sum of any solution is non-negative³;
- Eq4: $\sum_{\text{trg}(e) \cap V(\psi_l) \neq \emptyset} x_e = 0$ — this ensures that 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.

Now, by construction, it follows that $\text{LP}(\psi_l)$ admits a solution if and only if there exists a path π 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*. Now, consider the linear program $\text{LP}(\theta_1, \dots, \theta_n)$ defined with the following inequalities and equations:

- Eq1: $x_e \geq 0$ for each edge e — an edge cannot be used a negative number of times;
- Eq2: $\sum_{e \in E} x_e \geq 1$ — ensures that at least one edge is chosen;
- Eq3: for each $i \in \mathbb{N}$, $\sum_{e \in E} w'_i(\text{trg}(e))x_e \geq 0$ — this enforces that the total sum of any solution is non-negative;
- Eq4: for all $1 \leq r \leq n$, $\sum_{\text{trg}(e) \cap V(\theta_r) \neq \emptyset} x_e \geq 1$ — this ensures that for every $V(\theta_r)$ at least one state is in the cycle;
- 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.

In this case, $\text{LP}(\theta_1, \dots, \theta_n)$ admits a solution if and only if there exists a path π such that $z_i \leq \text{pay}_i(\pi)$ for every player i and visits every $V(\theta_r)$ *infinitely many times*. As highlighted

³Notice that by using $w'_i(\text{src}(e)) = w_i(\text{src}(e)) - z_i$ we ensure that a non-negative cycle corresponds to paths where agent i ensures a payoff greater or equal than z_i , which is the z -secure value for them, thus preventing deviations.

in [KS88], any positive multiple of a solution to the above LP problem also constitutes a valid solution. Consequently, by appropriately scaling any given solution, we can always obtain an integral solution.

Since the constructions above are polynomial in the size of both (\mathcal{G}, κ) and φ , we can conclude it is possible to check in NP the statement that there is a path π satisfying φ 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, observe that if $\varphi := \top$ and $\beta = 0$, then the problem is equivalent to checking whether the mp game has a Nash equilibrium, which is NP-hard [UW11]. \square

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

While it may be good news for the principal to find that $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$, it may not be sufficient. Even if there is a desirable Nash equilibrium, it is possible that other equilibria may be undesirable. In such cases, the principal may want to avoid the risk of the system getting stuck in bad equilibria. This motivates us to consider the *strong implementation* variant of equilibrium design. In a strong implementation, we require that *every* Nash equilibrium outcome satisfies the specification φ , for a *non-empty* set of outcomes. Let $\text{SI}(\mathcal{G}, \varphi, \beta)$ denote the set of reward schemes κ given budget β over \mathcal{G} such that:

- (1) (\mathcal{G}, κ) has at least one Nash equilibrium outcome,
- (2) every Nash equilibrium outcome of (\mathcal{G}, κ) satisfies φ .

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 5.1 (STRONG IMPLEMENTATION). Given a game \mathcal{G} , formula φ , and budget β :

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

STRONG IMPLEMENTATION can be solved with a similar procedure as in Algorithm 1 for WEAK IMPLEMENTATION where lines 1–4 are exactly the same, but with line 5 modified as follows:

Check whether:

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

Observe that a positive answer to both (a) and (b) above implies that $\text{NE}(\mathcal{G}, \kappa) \neq \emptyset$ and for every $\vec{\sigma} \in \text{NE}(\mathcal{G}, \kappa)$ we have $\pi(\vec{\sigma}) \models \varphi$. Thus, $\kappa \in \text{SI}(\mathcal{G}, \varphi, \beta)$ and as such $\text{SI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$. The complete algorithm is shown in Algorithm 2. For LTL specifications, to solve line 5 in Algorithm 2, consider the following LTL^{Lim} formulae that correspond, respectively, to

Algorithm 2 Algorithm for STRONG IMPLEMENTATION

input: $(\mathcal{G}, \varphi, \beta)$

1: guess:

- a reward scheme $\kappa \in \mathcal{K}(\mathcal{G}, \beta)$
- a state $s \in \text{St}$ for every player $i \in \mathbb{N}$
- punishment memoryless strategies $(\vec{\sigma}_{-1}, \dots, \vec{\sigma}_{-n})$ for all players $i \in \mathbb{N}$

2: Compute (\mathcal{G}, κ)

3: Compute the vector z where $z_i = \text{pun}_i(s)$ w.r.t. the punishment strategy $\vec{\sigma}_{-i}$

4: Build $(\mathcal{G}, \kappa)[z]$

5: **if** (a) there exists π in $(\mathcal{G}, \kappa)[z]$ such that $z_i \leq \text{pay}_i(\pi)$ for every player $i \in \mathbb{N}$
and
 (b) there is no π in $(\mathcal{G}, \kappa)[z]$ such that $z_i \leq \text{pay}_i(\pi)$ for every player $i \in \mathbb{N}$ and $\pi \models \neg\varphi$
then

6: **return** YES

7: **else**

8: **return** NO

conditions (a) and (b):

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

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

Notice that the expression $\text{NE}(\mathcal{G}, \kappa) \neq \emptyset$ can be expressed as “there exists a path π in \mathcal{G} that satisfies formula φ_{\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 π in \mathcal{G} , if π satisfies formula φ_{\exists} , then π also satisfies formula φ ”. Thus, using these two formulae to solve line 5, we obtain the following result.

Corollary 5.2. STRONG IMPLEMENTATION *with LTL specifications is PSPACE-complete.*

Proof. Using an argument similar to that in Theorem 4.4, we can directly infer the correctness of Algorithm 2 from Theorem 4.2 and the definition of STRONG IMPLEMENTATION. In particular, in line 5, the existence of path π in (a), and the absence of path π in (b), correspond to the non-emptiness of $\text{SI}(\mathcal{G}, \varphi, \beta)$.

Membership follows from the fact that for line 5, (a) can be solved by existential LTL^{Lim} model checking, and (b) by universal LTL^{Lim} model checking—both clearly in PSPACE by Savitch’s theorem. Hardness is similar to the construction in Theorem 4.4. \square

Case with GR(1) specifications. Notice that line 5 (a) in Algorithm 2 is essentially $\text{NE}(\mathcal{G}, \kappa) \neq \emptyset$, that is, checking whether the set of Nash equilibrium in a mean-payoff game is not empty—this can be solved in NP [UW11]. For the (b) 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 the formula above, we modify the LP in Theorem 4.6. Specifically, we modify Eq4 in $\text{LP}(\theta_1, \dots, \theta_n)$ to encode the θ -part of $\neg\varphi$. Thus, we have the following equation in $\text{LP}'(\theta_1, \dots, \theta_n)$:

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

All other equations remain the same.

In this case, $\text{LP}'(\theta_1, \dots, \theta_n)$ has a solution if and only if there is a path π 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 π 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 5.3. STRONG IMPLEMENTATION *with GR(1) specifications is Σ_2^P -complete.*

Proof. For membership, observe that by rearranging the problem statement, we have the following question:

Check whether the following expression is true

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

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

and

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

Statement (2) can be checked in NP (Theorem 4.2). Whereas, verifying statement (3) is in coNP; to see this, notice that we can rephrase (3) as follows: $\exists z \in \{\text{pun}_i(s) : s \in \text{St}\}^N$ such that both $\text{LP}(\psi_l)$ and $\text{LP}'(\theta_1, \dots, \theta_n)$ have a solution in $(\mathcal{G}, \kappa)[z]$. Thus Σ_2^P membership follows.

We prove hardness by a reduction from QSAT_2 (satisfiability of quantified Boolean formula with 2 alternations) [Pap94]. Let $\psi(\mathbf{x}, \mathbf{y})$ be an $n + m$ variable Boolean 3DNF formula, where $\mathbf{x} = \{x_1, \dots, x_n\}$ and $\mathbf{y} = \{y_1, \dots, y_m\}$, with t_1, \dots, t_k terms. Write \mathbf{t}_j for the set of literals in j -th term and \mathbf{t}_i^j for the i -th literal in \mathbf{t}_j . Moreover write x_i^j and y_i^j for variable $x_i \in \mathbf{x}$ and $y_i \in \mathbf{y}$ that appears in j -th term, respectively. For instance, if the fifth term is of the form of $(x_2 \wedge \neg x_3 \wedge y_4)$, then we have $\mathbf{t}_5 = \{x_2^5, x_3^5, y_4^5\}$ and $\mathbf{t}_5^1 = x_2^5$. Let $\mathbf{T} = \{\mathbf{t}_i \cap \mathbf{y} : 1 \leq i \leq k\}$, that is, the set of subset of \mathbf{t}_i that contains only y -literals.

For a formula $\psi(\mathbf{x}, \mathbf{y})$ we construct an instance of STRONG IMPLEMENTATION such that $\text{SI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ if and only if there is an $\vec{\mathbf{x}} \in \{0, 1\}^n$ such that $\psi(\mathbf{x}, \mathbf{y})$ is true for every $\vec{\mathbf{y}} \in \{0, 1\}^m$. Let \mathcal{G} be such a game where

- $N = \{1, 2\}$,
- $\text{St} = \{\bigcup_{j \in [1, k]} (\mathbf{t}_j \times \{0, 1\}^3)\} \cup \{\mathbf{T} \times \{0\}^3\} \cup \{\langle \text{source}, \{0\}^3 \rangle, \langle \text{sink}, \{0\}^3 \rangle\}$,
- $s_0 = \text{source}$,
- for each state $s \in \text{St}$
 - $\text{Ac}_1(s) = \{\{\mathbf{T} \cup \{\text{sink}\}\} \times \{0\}^3\}$, $\text{Ac}_2(s) = \{\varepsilon\}$, if $s = \langle \text{source}, \{0\}^3 \rangle$,
 - $\text{Ac}_1(s) = \{\mathbf{t}_i^1 : s[0] \subseteq \mathbf{t}_i \wedge i \in [1, k]\}$, $\text{Ac}_2(s) = \{0, 1\}^3$, if $s \in \{\mathbf{T} \times \{0\}^3\}$,
 - $\text{Ac}_1(s) = \{\varepsilon\}$, $\text{Ac}_2(s) = \{\varepsilon\}$, if $s \in \bigcup_{j \in [1, k]} (\mathbf{t}_j \times \{0, 1\}^3)$,
- for an action profile $\vec{\mathbf{a}} = (\mathbf{a}_1, \mathbf{a}_2)$
 - $\text{tr}(s, \vec{\mathbf{a}}) = \mathbf{a}_1$, if $s = \langle \text{source}, \{0\}^3 \rangle$,
 - $\text{tr}(s, \vec{\mathbf{a}}) = \langle \mathbf{a}_1, \mathbf{a}_2 \rangle$, if $s \in \{\mathbf{T} \times \{0\}^3\}$,

- $\text{tr}(s, \bar{\mathbf{a}}) = \langle \mathbf{t}_j^{(i \bmod 3)+1}, s[1] \rangle$, if $s = \langle \mathbf{t}_j^i, s[1] \rangle \in \bigcup_{j \in [1, k]} (\mathbf{t}_j \times \{0, 1\}^3)$;
- $\text{tr}(s, \bar{\mathbf{a}}) = s$, otherwise;
- for each state $s \in \text{St}$, $\lambda(s) = s[0]$,
- for each state $s \in \text{St}$
 - $w_1(s) = \frac{2}{3}$, if $s[0] = \mathbf{sink}^4$,
 - $w_1(s) = 0$, otherwise;
- the payoff of player $i \in \mathbf{N}$ for an ultimately periodic path π in \mathcal{G} is
 - $\text{pay}_1(\pi) = \text{mp}(w_1(\pi))$,
 - $\text{pay}_2(\pi) = -\text{mp}(w_1(\pi))$,

Furthermore, let $\beta = |\mathbf{x}|$ and the $\text{GR}(1)$ property to be $\varphi := \mathbf{GF} \neg \mathbf{sink}$. Define a (partial) reward scheme $\kappa : \mathbf{x} \rightarrow \{0, 1\}$. The weights are updated with respect to κ as follows: for each $s \in \text{St}$ such that $s[0] \in \mathbf{t}_j \setminus \mathbf{y}$, that is, an x -literal that appears in term t_j

$$w_1(s) = \begin{cases} 1, & \text{if } \kappa(s) = 1 \wedge s[0] \text{ is not negated in } t_j \\ 1, & \text{if } \kappa(s) = 0 \wedge s[0] \text{ is negated in } t_j \\ 0, & \text{if } \kappa(s) = 1 \wedge s[0] \text{ is negated in } t_j \\ 0, & \text{otherwise;} \end{cases}$$

for each $s \in \text{St}$ such that $s[0] \in \mathbf{t}_j \setminus \mathbf{x}$, that is, a y -literal that appears in term t_j , $s[0] = \mathbf{t}_j^i$

$$w_1(s) = \begin{cases} 1, & \text{if } s[1][i] = 1 \wedge s[0] \text{ is not negated in } t_j \\ 1, & \text{if } s[1][i] = 0 \wedge s[0] \text{ is negated in } t_j \\ 0, & \text{if } s[1][i] = 1 \wedge s[0] \text{ is negated in } t_j \\ 0, & \text{otherwise;} \end{cases}$$

the weights of other states remain unchanged.

The construction is now complete, and polynomial to the size of formula $\psi(\mathbf{x}, \mathbf{y})$. We claim that $\text{SI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ if and only if there is an $\vec{\mathbf{x}} \in \{0, 1\}^n$ such that $\psi(\mathbf{x}, \mathbf{y})$ is true for every $\vec{\mathbf{y}} \in \{0, 1\}^m$. From left to right, consider a reward scheme $\kappa \in \text{SI}(\mathcal{G}, \varphi, \beta)$ which implies that there exists no Nash equilibrium run in (\mathcal{G}, κ) that ends up in \mathbf{sink} . This means that for every action $\bar{\mathbf{a}}_2 \in \text{Ac}_2(s)$, there exists $\bar{\mathbf{a}}_1 \in \text{Ac}_1(s)$, $s \in \{\mathbf{T} \times \{0\}^3\}$, such that $\text{pay}_1(\pi) = 1$, where π is the resulting path of the joint action. Observe that this corresponds to the existence of (at least) a term t_i , which evaluates to true under assignment $\vec{\mathbf{x}}$, regardless the value of $\vec{\mathbf{y}}$. From right to left, consider an assignment $\vec{\mathbf{x}} \in \{0, 1\}^n$ such that for all $\vec{\mathbf{y}} \in \{0, 1\}^m$, the formula $\psi(\mathbf{x}, \mathbf{y})$ is true. This means that for every $\vec{\mathbf{y}}$, there exists (at least one) term t_i in $\psi(\mathbf{x}, \mathbf{y})$ that evaluates to true. By construction, specifically the weight updating rules, for every $\bar{\mathbf{a}}_2$ corresponding to assignment $\vec{\mathbf{y}}$, there exists \mathbf{t}_j such that $\forall i \in [1, 3], w_1(\mathbf{t}_j^i) = 1$. This means that player 1 can always get payoff equals to 1, therefore, any run that ends in \mathbf{sink} is not sustained by Nash equilibrium. \square

6. OPTIMALITY AND UNIQUENESS OF SOLUTIONS

Having asked the questions studied in the previous sections, the principal may want to dig deeper. 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.)

⁴This can be implemented by a macrostate with three substates—2 substates with weight of 1, and 1 with weight of 0—forming a simple cycle.

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 reward 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 6.1 (OPT-WI). Given a game \mathcal{G} and a specification formula φ :

What is the optimum budget β 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 φ , 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 6.2 (EXACT-WI). Given a game \mathcal{G} , a specification formula φ , 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 κ' 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 6.3. *Let z_i be the largest payoff that player i can get after deviating from a path π . The optimum budget is an integer between 0 and $\sum_{i \in \mathcal{N}} z_i \cdot (|\text{St}| - 1)$.*

Proof. The lower-bound is straightforward. The upper-bound follows from the fact that the maximum value the principal has to pay to player i is when the path π is a simple cycle and formed from all states in St , apart from 1 deviation state. \square

Using Proposition 6.3, 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 6.3 to go through all possible solutions in exponential time, but using only nondeterministic polynomial space. Formally, we have the following results.

Theorem 6.4. *OPT-WI with LTL specifications is FSPACE-complete.*

Proof. Since the search space is bounded (Proposition 6.3), by using WEAK IMPLEMENTATION an an oracle we can iterate through every instance and return the smallest β such that $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$. Moreover, each instance is of polynomial size in the size of the input. Thus membership in PSPACE follows. Hardness is straightforward. \square

Corollary 6.5. *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 6.6. *OPT-WI with GR(1) specifications is FP^{NP} -complete.*

Proof. Membership follows from the fact that the search space, which is bounded as in Proposition 6.3, can be fully explored using binary search and WEAK IMPLEMENTATION as an oracle. More precisely, we can find the smallest budget β such that $\text{WI}(\mathcal{G}, \varphi, \beta) \neq \emptyset$ by checking every possible value for β , which lies between 0 and 2^n , where n is the length of the encoding of the instance. Since we need logarithmically many calls to the NP oracle (to WEAK IMPLEMENTATION), in the end we have searching procedure that runs in polynomial time.

For hardness we reduce from TSP COST (optimal travelling salesman problem) that is known to be FP^{NP} -complete [Pap94]. Given a TSP COST instance $\langle G, c \rangle$, $G = \langle V, E \rangle$ is a graph, $c : E \rightarrow \mathbb{Z}$ is a cost function. We assume that $\text{WI}(\mathcal{G}, \varphi, \beta)$ is efficient. To encode TSP COST instance, we construct a game \mathcal{G} and GR(1) formula φ , such that the optimum budget β corresponds to the value of optimum tour. Let \mathcal{G} be such a game where

- $N = \{1\}$,
- $\text{St} = \{\langle v, e \rangle : v \in V \wedge e \in \text{in}(v)\} \cup \{\langle \text{sink}, \varepsilon \rangle\}$,
- s_0 can be chosen arbitrarily from $\text{St} \setminus \{\langle \text{sink}, \varepsilon \rangle\}$,
- for each state $\langle v, e \rangle \in \text{St}$ and edge $e' \in E \cup \{\varepsilon\}$
 - $\text{tr}(\langle v, e \rangle, e') = \langle \text{trg}(e'), e' \rangle$, if $v \neq \text{sink}$ and $e' \neq \varepsilon$,
 - $\text{tr}(\langle v, e \rangle, e') = \langle \text{sink}, \varepsilon \rangle$, otherwise;
- for each state $\langle v, e \rangle \in \text{St}$
 - $w_1(\langle v, e \rangle) = \max\{c(e') : e' \in E\} - c(e)$, if $v \neq \text{sink}$,
 - $w_1(\langle v, e \rangle) = \max\{c(e') : e' \in E\}$, otherwise;
- the payoff of player 1 for a path π in \mathcal{G} is $\text{pay}_1(\pi) = \text{mp}(w_1(\pi))$,
- for each state $\langle v, e \rangle \in \text{St}$, the set of actions available to player 1 is $\text{out}(v) \cup \{\varepsilon\}$,
- for each state $\langle v, e \rangle \in \text{St}$, $\lambda(\langle v, e \rangle) = v$.

Furthermore, let $\varphi := \bigwedge_{v \in V} \mathbf{GF} v$. The construction is now complete, and is polynomial to the size of $\langle G, c \rangle$.

Now, consider the smallest $\text{cost}(\kappa)$, $\kappa \in \text{WI}(\mathcal{G}, \varphi, \beta)$. We argue that $\text{cost}(\kappa)$ is indeed the lowest value such that a tour in G is attainable. Suppose for contradiction, that there exists κ' such that $\text{cost}(\kappa') < \text{cost}(\kappa)$. Let π' be a path in (\mathcal{G}, κ') and $z_1 = w_1(\langle \text{sink}, \varepsilon \rangle)$ the largest value player 1 can get by deviating from π' . We have $\text{pay}_1(\pi') < z_1$, and since for every $\langle v, e \rangle \in \text{St}$ there exists an edge to $\langle \text{sink}, \varepsilon \rangle$, thus player 1 would deviate to $\langle \text{sink}, \varepsilon \rangle$ and stay there forever. This deviation means that φ is not satisfied, which is a contradiction to $\kappa' \in \text{WI}(\mathcal{G}, \varphi, \beta)$. The construction of φ also ensures that the path is a valid tour, i.e., the tour visits every city at least once. Notice that φ does not guarantee a Hamiltonian cycle. However, removing the condition of visiting each city *only once* does not remove the hardness, since *Euclidean* TSP is NP-hard [GGJ76, Pap77]. Therefore, in the planar case there is an optimal tour that visits each city only once, or otherwise, by the triangle inequality, skipping a repeated visit would not increase the cost. Finally, since $\text{WI}(\mathcal{G}, \varphi, \beta)$ is efficient, we have β to be exactly the value of the optimum tour in the corresponding TSP COST instance. \square

Corollary 6.7. EXACT-WI with GR(1) specifications is D^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 $WI(\mathcal{G}, \varphi, \beta) = \emptyset$). Since the former problem is in NP and the latter problem is in coNP, membership in D^P follows. For the lower bound, we use the same reduction technique as in Theorem 6.6, and reduce from EXACT TSP, a problem known to be D^P -hard [Pap94, PY84]. \square

Following [Pap84], we may naturally ask whether the optimal solution given by OPT-WI is unique. We call this problem UOPT-WI. For some fixed budget β , it may be the case that for two reward schemes $\kappa, \kappa' \in WI(\mathcal{G}, \varphi, \beta)$ – 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 6.8. 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 reward schemes with precisely this budget; if no such reward schemes exist, return “yes”; otherwise, return “no”.

The above decision procedure clearly is in Δ_2^P (for the upper bound). Furthermore, since Theorem 6.6 implies Δ_2^P -hardness [Kre88] (for the lower bound), we have the following corollary.

Corollary 6.9. UOPT-WI with GR(1) specifications is Δ_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 6.10 (OPT-SI). Given a game \mathcal{G} and a specification formula φ :

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

Definition 6.11 (EXACT-SI). Given a game \mathcal{G} , a specification formula φ , and an integer b :

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

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 6.12. OPT-SI with LTL specifications is FSPACE-complete.

Proof. The proof is analogous to that of Theorem 6.4. \square

Corollary 6.13. 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 $FP^{\Sigma_2^P}$. For hardness, we define an $FP^{\Sigma_2^P}$ -complete problem, namely WEIGHTED MINQSAT₂. Recall that in QSAT₂ 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 $\text{WEIGHTED MINQSAT}_2$ as follows. Given $\psi(\mathbf{x}, \mathbf{y})$ and a weight function $\mathbf{c} : \mathbf{x} \rightarrow \mathbb{Z}^{\geq}$, $\text{WEIGHTED MINQSAT}_2$ 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 $\text{WEIGHTED MINQSAT}_2$ generalises MINQSAT_2 , which is known to be $\text{FP}^{\Sigma_2^{\text{P}}[\log n]}$ -hard [CH04], *i.e.*, MINQSAT_2 is an instance of $\text{WEIGHTED MINQSAT}_2$, where all weights are 1.

Theorem 6.14. $\text{WEIGHTED MINQSAT}_2$ is $\text{FP}^{\Sigma_2^{\text{P}}}$ -complete.

Proof. Membership follows from the upper-bound of MINQSAT_2 [CH04]: since we have an exponentially large input with respect to that of MINQSAT_2 , by using binary search we will need polynomially many calls to the Σ_2^{P} oracle. Hardness is immediate [CH04]. \square

Now that we have an $\text{FP}^{\Sigma_2^{\text{P}}}$ -hard problem in hand, we can proceed to determine the complexity class of OPT-SI with $\text{GR}(1)$ specifications. For the upper bound we one can use arguments analogous to those in Theorem 6.6. For the lower bound, one can reduce from $\text{WEIGHTED MINQSAT}_2$. Formally, we have:

Theorem 6.15. OPT-SI with $\text{GR}(1)$ specifications is $\text{FP}^{\Sigma_2^{\text{P}}}$ -complete.

Proof. Membership uses arguments analogous to those in Theorem 6.6. For hardness, we reduce $\text{WEIGHTED MINQSAT}_2$ to OPT-SI using the same techniques used in Theorem 5.3 with few modifications. Given a $\text{WEIGHTED MINQSAT}_2$ instance $\langle \psi(\mathbf{x}, \mathbf{y}), \mathbf{c} \rangle$, we construct a game \mathcal{G} and $\text{GR}(1)$ formula φ , such that the optimum budget β corresponds to the value of optimal solution to $\langle \psi(\mathbf{x}, \mathbf{y}), \mathbf{c} \rangle$. To this end, we may assume that $\text{SI}(\mathcal{G}, \varphi, \beta)$ is efficient and construct \mathcal{G} with exactly the same rules as in Theorem 5.3 except for the following:

- clearly the value of β is unknown,
- the initial weight for each state $s \in \text{St}$
 - $w_1(s) = \frac{2}{3}$, if $s[0] = \mathbf{sink}$,

$$w_1(s) = \begin{cases} -\mathbf{c}(s[0]) + 1, & \text{if } s[0] \in \mathbf{t}_j \setminus \mathbf{y} \wedge s[0] \text{ is not negated in } t_j \\ 1, & \text{if } s[0] \in \mathbf{t}_j \setminus \mathbf{y} \wedge s[0] \text{ is negated in } t_j; \end{cases}$$

- $w_1(s) = 0$, otherwise;
- given a reward scheme κ , we update the weight for each $s \in \text{St}$ such that $s[0] \in \mathbf{t}_j \setminus \mathbf{y}$, that is, an x -literal that appears in term t_j

$$w_1(s) = \begin{cases} w_1(s) + \kappa(s), & \text{if } s[0] \text{ is not negated in } t_j \\ w_1(s), & \text{otherwise;} \end{cases}$$

the construction is complete and polynomial to the size of $\langle \psi(\mathbf{x}, \mathbf{y}), \mathbf{c} \rangle$.

Let o be the optimal solution to $\text{WEIGHTED MINQSAT}_2$ given the input $\langle \psi(\mathbf{x}, \mathbf{y}), \mathbf{c} \rangle$. We claim that β is exactly o . To see this, consider the smallest $\text{cost}(\kappa)$, $\kappa \in \text{SI}(\mathcal{G}, \varphi, \beta)$. We argue that this is indeed the least total weight of an assignment $\vec{\mathbf{x}}$ such that $\psi(\mathbf{x}, \mathbf{y})$ is true for every $\vec{\mathbf{y}}$. Assume towards a contradiction that $\text{cost}(\kappa) < o$. By the construction of $w_1(\cdot)$, there exists no π such that $\text{pay}_1(\pi) > \frac{2}{3}$. Therefore, any run π' that ends up in \mathbf{sink} is sustained by Nash equilibrium, which is a contradiction to $\kappa \in \text{SI}(\mathcal{G}, \varphi, \beta)$. Now, since $\text{SI}(\mathcal{G}, \varphi, \beta)$ is efficient, by definition, there exists no $\kappa' \in \text{SI}(\mathcal{G}, \varphi, \beta)$ such that $\text{cost}(\kappa') < \text{cost}(\kappa)$. Thus we have β equals to o as required. \square

Corollary 6.16. EXACT-SI with GR(1) specifications is 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 $SI(\mathcal{G}, \varphi, \beta) = \emptyset$ instead. The lower bound follows from the hardness of STRONG IMPLEMENTATION and STRONG IMPLEMENTATION COMPLEMENT problems, which immediately implies D_2^P -hardness [ACHI17, Lemma 3.2]. \square

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

Corollary 6.17. UOPT-SI with LTL specifications is PSPACE-complete.

Corollary 6.18. UOPT-SI with GR(1) specifications is Δ_3^P -complete.

7. EQUILIBRIUM DESIGN WITH SOCIAL WELFARE

Until this point, we have only considered problems that primarily concern the satisfaction of temporal logic specifications. Indeed, this is one of the key differences between equilibrium design and mechanism design. However, a benevolent principal may not only be concerned with the satisfaction of specifications (and optimality of solutions) but also the well-being of agents and the fairness of outcomes. Well-being and fairness can influence the equilibria of a game in many ways. For example, if agents are subject to *inequity aversion* [FS99] (i.e., they value fairness and are willing to forego some personal gain), then this can affect the strategies they choose and the resulting equilibria of the game. To account for this, we extend implementation problems to include *social welfare* measures.

One well-known measure of social welfare is *utilitarian social welfare*, which provides a measure of overall and average benefit for society. However, this measure may be appropriate for certain circumstances but problematic for others. For instance, it may lead to an unfair reward scheme where the principal allocates all the budget to one player, and none to the others. An alternative measure that takes fairness into consideration is *egalitarian social welfare*. It measures the welfare of a society by the well-being of the worst-off individual (i.e., the maximin criterion). A notable argument in defence of this system is Rawls’ *veil of ignorance* [Raw99]. From the players’ perspective, this notion of welfare may be more relevant for our setting: prior to the application of a reward scheme, the players do not know which scheme will be chosen and implemented since the principal computes and chooses it “behind the veil of ignorance”⁵.

Egalitarian welfare may not only be preferable in a philosophical sense, but also useful in a practical way. To see this, let us revisit our previous Example 2.1 and make some changes to the scenario as follows. For each step robot \blacksquare takes, it spends 2 units of energy and receives a payment of -2 . This means that the payoff of robot \blacksquare in Nash equilibria is $-\frac{1}{2}$, while robot \circ ’s remains $\frac{1}{2}$. If we consider the robots’ payoffs as their energy levels, then the Nash equilibria in such a situation are not attainable: robot \blacksquare does not have enough energy to realise Nash equilibrium runs. In this setting, the egalitarian social welfare concept is more useful since we want to ensure that every robot has enough resources to

⁵In [EM04] a similar argument for egalitarian system in an *artificial* society is presented using an example of fair access agreement of a common resource (satellite) [LVB99].

carry out its task⁶. However, it is important to note that this analysis comes with some caveats: (a) it generally works only for prefix-independent specifications, and (b) it ignores any finite prefixes, i.e., on some finite prefixes, the payoff may (temporarily) be negative, but is positive as the number of steps approaches infinity. These limitations are typically not a concern for prefix-independent specifications such as $\text{GR}(1)$, which is used in this paper.

Formally, we represent the measure of social welfare as a function $S : \mathbb{R}^{|\mathcal{N}|} \rightarrow \mathbb{R}$ mapping a tuple of real numbers into a real number, representing a *payoff aggregation*. More specifically, for a strategy profile $\vec{\sigma}$, the social welfare measure of $\vec{\sigma}$ is given by $S(\text{pay}_1(\vec{\sigma}), \dots, \text{pay}_n(\vec{\sigma}))$. With a slight abuse of notation, we write $S(\vec{\sigma})$ for the social welfare measure of $\vec{\sigma}$. The two aforementioned concepts of social welfare thus can be defined as follows.

Definition 7.1. For a game \mathcal{G} and a strategy profile $\vec{\sigma}$, define

- (1) the *utilitarian social welfare* of $\vec{\sigma}$ w.r.t. \mathcal{G} as $\text{usw}(\vec{\sigma}) = \sum_{i \in \mathcal{N}} \text{pay}_i(\vec{\sigma})$;
- (2) the *egalitarian social welfare* of $\vec{\sigma}$ w.r.t. \mathcal{G} as $\text{esw}(\vec{\sigma}) = \min_{i \in \mathcal{N}} \{\text{pay}_i(\vec{\sigma})\}$;

We also denote the minimum utilitarian and egalitarian social welfare for a given set $\text{NE}(\mathcal{G})$ by:

$$\text{MinNE}_S(\mathcal{G}) = \min\{S(\vec{\sigma}) : \vec{\sigma} \in \text{NE}(\mathcal{G})\}$$

where $S \in \{\text{usw}, \text{esw}\}$.

We now define the decision problem of *threshold social welfare*. This problem involves deciding whether there is a reward scheme with a social welfare measure of no less than a given value. Formally, the problem is defined as follows.

Definition 7.2 (Threshold social welfare). For a given WEAK IMPLEMENTATION (resp. STRONG IMPLEMENTATION) instance $(\mathcal{G}, \varphi, \beta)$, a social welfare function S , and a threshold value t , decide whether there is $\kappa \in \text{WI}(\mathcal{G}, \varphi, \beta)$ (resp. $\kappa \in \text{SI}(\mathcal{G}, \varphi, \beta)$) such that $t \leq \text{MinNE}_S((\mathcal{G}, \kappa))$. We write UT-WI and UT-SI, for threshold problems with utilitarian social welfare function in WEAK IMPLEMENTATION and STRONG IMPLEMENTATION domains, respectively. Similarly, we write ET-WI and ET-SI for threshold problems with egalitarian social welfare function in WEAK IMPLEMENTATION and STRONG IMPLEMENTATION, respectively.

To solve the threshold problems, we utilise the procedures for WEAK IMPLEMENTATION and STRONG IMPLEMENTATION presented in previous sections. We first show how to check the threshold problem where the specification is given in LTL formulae. For the utilitarian social welfare measure, we have the following.

Theorem 7.3. UT-WI and UT-SI are PSPACE-complete for LTL specifications.

Proof. We begin with the upper bound for UT-WI. We apply a slight modification of the WEAK IMPLEMENTATION problem with LTL specifications. Consider the arena $A' = (\mathcal{N} \cup \{n+1\}, \text{Ac}, \text{St}, s_0, \text{tr}', \lambda)$, with tr' defined as

$$\text{tr}'(\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}_{n+1}) = \text{tr}(\mathbf{a}_1, \dots, \mathbf{a}_n)$$

for every $(\mathbf{a}_1, \dots, \mathbf{a}_n, \mathbf{a}_{n+1}) \in \text{Ac}^{|\mathcal{N}|+1}$, and the game $\mathcal{G}' = \langle A', (\mathbf{w}_i)_{i \in \mathcal{N}}, (\mathbf{w}_{n+1}) \rangle$ with $\mathbf{w}_{n+1}(s) = \sum_{i \in \mathcal{N}} (\mathbf{w}_i(s))$ for each $s \in \text{St}$. Player $n+1$ is a dummy player who does not affect the play

⁶Indeed, the relationship between limited resources and their usage by system components has been extensively studied in the form of games in many research papers, and such games are closely related to games with mean-payoff goals (see e.g., [CdAHS03, CDHR10, CD12, DDG⁺10, VCD⁺15]).

of the game. Observe that for every strategy profile $\vec{\sigma}$ in \mathcal{G}' , it holds that

$$\text{pay}'_{n+1}(\vec{\sigma}) = \sum_{i \in \mathbb{N}} \text{pay}'_i(\vec{\sigma}) = \sum_{i \in \mathbb{N}} \text{pay}_i(\vec{\sigma}_{-(n+1)}) = \text{usw}(\vec{\sigma}_{-(n+1)}).$$

We can adapt the same procedure for solving the WEAK IMPLEMENTATION problem in order to solve UT-WI. To this end, we replace the LTL^{Lim} formula φ_{WI} with

$$\varphi_{\text{WI,usw}} := \varphi_{\text{WI}} \wedge (\text{mp}(n+1) \geq t).$$

Observe that if the formula above is satisfied in $(\mathcal{G}', \kappa)[z]$, then exists $\kappa \in \text{WI}(\mathcal{G}, \varphi, \beta)$ such that $t \leq \text{MinNE}_{\text{usw}}((\mathcal{G}, k))$. Thus, the PSPACE upper bound follows.

The PSPACE upper bound for UT-SI case can be obtained in a similar way by adapting the procedure for solving the STRONG IMPLEMENTATION problem. The construction is similar to the one for UT-WI explained above, and it suffices to replace the LTL^{Lim} formula φ_{\exists} with

$$\varphi_{\exists,\text{usw}} := \varphi_{\exists} \wedge (\text{mp}(n+1) \geq t),$$

and φ_{\forall} with

$$\varphi_{\forall,\text{usw}} := \varphi_{\exists,\text{usw}} \rightarrow \varphi.$$

The lower bounds for both cases immediately follow from the lower bound of WEAK IMPLEMENTATION and STRONG IMPLEMENTATION problems since we can reduce those problems to UT-WI and UT-SI, respectively, by fixing t to be the smallest weight value appeared in the arena. \square

For egalitarian social welfare measure, we obtain the following result.

Theorem 7.4. *ET-WI and ET-SI are PSPACE-complete for LTL specifications.*

Proof. The upper bounds for ET-WI and ET-SI follow from the adaptation of Algorithms 1 and 2 used to solve WEAK IMPLEMENTATION and STRONG IMPLEMENTATION, respectively. For ET-WI, we begin with the fact that for a given WEAK IMPLEMENTATION instance $(\mathcal{G}, \varphi, \beta)$ and a threshold t , we can solve ET-WI as follows: check if there is a κ such that

- (1) $\kappa \in \text{WI}(\mathcal{G}, \varphi, \beta)$, and
- (2) $\text{MinNE}_{\text{esw}}((\mathcal{G}, \kappa)) \geq t$.

Notice that (1) is exactly WEAK IMPLEMENTATION, and thus can be solved by Algorithm 1. For (2), observe that the constraint $\text{MinNE}_{\text{esw}}((\mathcal{G}, \kappa)) \geq t$ can be “embedded” already in line 5 of the algorithm by requiring that $\text{pay}_i(\pi) \geq t, \forall i \in \mathbb{N}$. This can be done by replacing the formula φ_{WI} corresponding to line 5 with the following formula

$$\varphi_{\text{WI,esw}} := \varphi_{\text{WI}} \wedge \bigwedge_{i \in \mathbb{N}} (\text{mp}(i) \geq t).$$

We can also solve ET-SI in a similar way to the above. We modify line 5 in Algorithm 2 by adding an extra requirement that the path π satisfies $\text{pay}_i(\pi) \geq t, \forall i \in \mathbb{N}$. To this end, we replace φ_{\exists} with

$$\varphi_{\exists,\text{esw}} := \varphi_{\exists} \wedge \bigwedge_{i \in \mathbb{N}} (\text{mp}(i) \geq t),$$

and φ_{\forall} with

$$\varphi_{\forall,\text{esw}} := \varphi_{\exists,\text{esw}} \rightarrow \varphi.$$

Lower bounds follow directly from the hardness of WEAK IMPLEMENTATION and STRONG IMPLEMENTATION by the same argument as in Theorem 7.3. \square

We now address the threshold problems where the specifications are given in the GR(1) fragment. For implementations with LTL specifications, the bottleneck comes from the LTL model checking problem. As a result, adding an extra constraint for social welfare threshold would not affect the overall complexity. Surprisingly, this is also true for GR(1) specifications: adding the constraint does not incur significant computational cost.

Theorem 7.5. *UT-WI and UT-SI are NP-complete and Σ_2^P -complete, respectively, for GR(1) specifications.*

Proof. Again, we use a similar construction as in the proof of Theorem 7.3 to build A' and \mathcal{G}' , and adapt the procedures for solving the WEAK IMPLEMENTATION and STRONG IMPLEMENTATION with GR(1) specifications. In both domains, we construct the corresponding multi-weighted graph $W = (V, E, (w'_a)_{a \in N})$ where $w'_{n+1}(v) = w_{n+1}(s) - t$. The query to this procedure corresponds exactly to the threshold social welfare problem, giving us the upper bounds. Lower bounds can be obtained by setting $t = \min\{w_{n+1}(s) : s \in \text{St}\}$. \square

Theorem 7.6. *ET-WI and ET-SI are NP-complete and Σ_2^P -complete, respectively, for GR(1) specifications.*

Proof. To solve ET-WI and ET-SI, we directly adapt from the procedures for solving the WEAK IMPLEMENTATION and STRONG IMPLEMENTATION with GR(1) specifications (Theorems 4.6 and 5.3). For ET-WI, from the game $\mathcal{G}[z]$, we build the underlying graph $\langle V, E, (w'_i)_{i \in N} \rangle$ where $w'_i(v) = w_i(s) - (\max\{z_i, t\})$. Then we define the linear programs $\text{LP}(\psi_i)$ and $\text{LP}(\theta_1, \dots, \theta_n)$ in the same way. Observe that, one of the two linear programs has a solution if and only if there is a path π satisfying φ such that for every player i , $z_i \leq \text{pay}_i(\pi)$ and $t \leq \text{pay}_i(\pi)$. For ET-SI, we also employ a similar construction. To obtain the lower bounds, we reduce from WEAK IMPLEMENTATION and STRONG IMPLEMENTATION with GR(1) specifications. The reduction simply follows from the fact that by fixing $t = \min\{w_i(s) : i \in N, s \in \text{St}\}$, we can encode WEAK IMPLEMENTATION and STRONG IMPLEMENTATION with GR(1) specifications into their corresponding social threshold problems. \square

8. 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 [HR06], 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 [FKL10, GHW15]. 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 [GHW17] 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 reward schemes for incentivising behaviours in a concurrent and multi-agent system, but we do not require such reward 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. To remedy this issue, we added social welfare constraints in the design of reward schemes, and showed that such additions do not incur extra cost from computational complexity perspective.

Equilibrium design vs. rational verification – connections with computer science.

Typically, in rational synthesis and verification [FKL10, GHW15, GHW17, KPV16] 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 [AKP18, BBFR13, CD12, CDHR10, CHJ05, GMP⁺17, GNPW19b, VCD⁺15]. 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 [BBMU15, GHW17]. 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 [GHW15, GHW17]. For this reason, there has been already some work on the introduction of desirable Nash equilibria in multi-player games [AAK15, Per19]. This problem has been studied as a repair problem [AAK15] 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 [Per19]. 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. This paper considers games with deterministic transitions, perfect information, and non-cooperative setting. There are different interesting directions to extend the work done here. For instance, considering games in which players may be able to cooperate with each other and form coalitions [GKW19, SGW21]. It would also be interesting to investigate classes of imperfect information games in which we may still obtain decidability for our problems [GPW18, DDG⁺10, BMM17, BLMR17, AMM14]. Another avenue for future research would be to look into games with probabilistic aspects [KNPS19, GHL⁺21].

On the other hand, the reward model proposed in this paper has a *memoryless* structure, meaning that the rewards depend only on the current states and not on the history of the play. This restricts the kinds of equilibria that we can design with this model. For instance, we cannot always design equilibria that implement safety specifications, since they are prefix dependent properties. One way to address this limitation is to use a model that incorporates the history of play. For example, finite state machines can be used to create a finite-memory reward scheme modelled by a transducer. This model is similar to a player’s strategy, but the output is a vector of integer rewards. However, since the space of possible schemes is unbounded, a different approach is needed for equilibrium design with this model. Additionally, this reward model is related to the concept of *reward machines* in the reinforcement learning framework [IKVM22]. Exploring this direction and incorporating probabilistic aspects of games can establish a connection between equilibrium design and multi-agent reinforcement learning. This is an interesting area of research.

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 [GNPW18].

ACKNOWLEDGEMENTS

Perelli was supported by the PNRR MUR project PE0000013-FAIR and the PRIN 2020 projects PINPOINT. He was also supported by Sapienza University of Rome under the “Progetti Grandi di Ateneo” programme, grant RG123188B3F7414A (ASGARD - Autonomous and Self-Governing Agent-Based Rule Design). Wooldridge acknowledges the support of UKRI under a Turing AI World Leading Researcher Fellowship (grant EP/W002949/1).

REFERENCES

- [AAK15] S. Almagor, G. Avni, and O. Kupferman. Repairing Multi-Player Games. In *CONCUR*, volume 42 of *LIPICs*, pages 325–339. Schloss Dagstuhl, 2015.
- [ACHI17] 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>.
- [AKP18] S. Almagor, O. Kupferman, and G. Perelli. Synthesis of Controllable Nash Equilibria in Quantitative Objective Games. In *IJCAI*, pages 35–41, 2018.
- [AMM14] Benjamin Aminof, Fabio Mogavero, and Aniello Murano. Synthesis of Hierarchical Systems. *Science of Computer Programming*, 83:56–79, 2014.
- [AMMR16] 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.
- [BBFR13] A. Bohy, V. Bruyère, E. Filiot, and J. Raskin. Synthesis from LTL Specifications with Mean-Payoff Objectives. In *TACAS*, pages 169–184, 2013.
- [BBMU15] P. Bouyer, R. Brenguier, N. Markey, and M. Ummels. Pure Nash Equilibria in Concurrent Deterministic Games. *Logical Methods in Computer Science*, 11(2), 2015.

- [BCHK14] 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.
- [BJP⁺12] 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.
- [BLMR17] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. Verification of Multi-agent Systems with Imperfect Information and Public Actions. In *AAMAS*, pages 1268–1276. ACM, 2017.
- [BMM17] Raphaël Berthon, Bastien Maubert, and Aniello Murano. Decidability Results for ATL* with Imperfect Information and Perfect Recall. In *AAMAS*, pages 1250–1258. ACM, 2017.
- [CD12] K. Chatterjee and L. Doyen. Energy Parity Games. *Theoretical Computer Science*, 458:49–60, 2012.
- [CdAHS03] Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource Interfaces. In Rajeev Alur and Insup Lee, editors, *Embedded Software*, pages 117–133, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
- [CDHR10] 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.
- [CH04] H. Chockler and J. Halpern. Responsibility and Blame: A Structural-Model Approach. *Journal of Artificial Intelligence Research*, 22:93–115, 2004.
- [CHJ05] K. Chatterjee, T. A. Henzinger, and M. Jurdzinski. Mean-Payoff Parity Games. In *LICS*, pages 178–187. IEEE Computer Society, 2005.
- [DDG⁺10] 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.
- [EM04] Ulle Endriss and Nicolas Maudet. Welfare Engineering in Multiagent Systems. In Andrea Omicini, Paolo Petta, and Jeremy Pitt, editors, *Engineering Societies in the Agents World IV*, pages 93–106, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [FKL10] D. Fisman, O. Kupferman, and Y. Lustig. Rational Synthesis. In *TACAS*, volume 6015 of *LNCS*, pages 190–204. Springer, 2010.
- [FS99] Ernst Fehr and Klaus M Schmidt. A Theory of Fairness, Competition, and Cooperation. *The quarterly journal of economics*, 114(3):817–868, 1999.
- [GGJ76] M. Garey, R. Graham, and D. Johnson. Some NP-complete Geometric Problems. In *STOC*, pages 10–22. ACM, 1976.
- [GHL⁺21] Julian Gutierrez, Lewis Hammond, Anthony W. Lin, Muhammad Najib, and Michael Wooldridge. Rational Verification for Probabilistic Systems. In *Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning*, pages 312–322, 11 2021. doi:10.24963/kr.2021/30.
- [GHPW17] 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.
- [GHW15] J. Gutierrez, P. Harrenstein, and M. Wooldridge. Iterated Boolean Games. *Information and Computation*, 242:53–79, 2015.
- [GHW17] J. Gutierrez, P. Harrenstein, and M. Wooldridge. From Model Checking to Equilibrium Checking: Reactive Modules for Rational Verification. *Artificial Intelligence*, 248:123–157, 2017.
- [GKW19] Julian Gutierrez, Sarit Kraus, and Michael J. Wooldridge. Cooperative Concurrent Games. In Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor, editors, *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’19, Montreal, QC, Canada, May 13-17, 2019*, pages 1198–1206. International Foundation for Autonomous Agents and Multiagent Systems, 2019. URL: <http://dl.acm.org/citation.cfm?id=3331822>.
- [GMP⁺17] 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.
- [GNPW18] 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.

- [GNPW19a] Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael Wooldridge. Equilibrium Design for Concurrent Games. In Wan Fokkink and Rob van Glabbeek, editors, *30th International Conference on Concurrency Theory (CONCUR 2019)*, volume 140 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:16, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.CONCUR.2019.22.
- [GNPW19b] Julian Gutierrez, Muhammad Najib, Giuseppe Perelli, and Michael Wooldridge. On Computational Tractability for Rational Verification. In *IJCAI-19*, pages 329–335, 7 2019.
- [GPW18] J. Gutierrez, G. Perelli, and M. Wooldridge. Imperfect Information in Reactive Modules Games. *Information and Computation*, 261(Part):650–675, 2018.
- [HR06] L. Hurwicz and S. Reiter. *Designing Economic Mechanisms*. Cambridge University Press, 2006.
- [IKVM22] Rodrigo Toro Icarte, Torny Q Klassen, Richard Valenzano, and Sheila A McIlraith. Reward Machines: Exploiting Reward Function Structure in Reinforcement Learning. *Journal of Artificial Intelligence Research*, 73:173–208, 2022.
- [KNPS19] Marta Kwiatkowska, Gethin Norman, David Parker, and Gabriel Santos. Equilibria-based probabilistic model checking for concurrent stochastic games. In Maurice H. ter Beek, Annabelle McIver, and José N. Oliveira, editors, *Formal Methods – The Next 30 Years*, pages 298–315, Cham, 2019. Springer International Publishing.
- [KPV16] O. Kupferman, G. Perelli, and M. Y. Vardi. Synthesis with Rational Environments. *Annals of Mathematics and Artificial Intelligence*, 78(1):3–20, 2016.
- [Kre88] 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.
- [KS88] S Rao Kosaraju and Gregory Sullivan. Detecting Cycles in Dynamic Graphs in Polynomial Time. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 398–406, 1988.
- [LVB99] Michel Lemaitre, Gérard Verfaillie, and Nicolas Bataille. Exploiting a Common Property Resource under a Fairness Constraint: A Case Study. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI '99*, page 206–211, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
- [M. 13] 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.
- [OR94] M.J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
- [Pap77] C. Papadimitriou. The Euclidean Travelling Salesman Problem is NP-complete. *Theoretical Computer Science*, 4(3):237 – 244, 1977. doi:10.1016/0304-3975(77)90012-3.
- [Pap84] C. Papadimitriou. On the Complexity of Unique Solutions. *Journal of the ACM*, 31(2):392–400, 1984. doi:10.1145/62.322435.
- [Pap94] C. Papadimitriou. *Computational complexity*. Addison-Wesley, Reading, Massachusetts, 1994.
- [Per19] 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>.
- [Pnu77] A. Pnueli. The Temporal Logic of Programs. In *FOCS*, pages 46–57. IEEE, 1977.
- [PR89] A. Pnueli and R. Rosner. On the Synthesis of a Reactive Module. In *POPL*, pages 179–190. ACM Press, 1989.
- [PY84] 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.
- [Raw99] John Rawls. *A Theory of Justice*. Harvard University Press, 1999.
- [S. 09] S. Heubach and T. Mansour. *Combinatorics of Compositions and Words: Solutions Manual*. Chapman & Hall/CRC, 2009.
- [SGW21] Thomas Steeples, Julian Gutierrez, and Michael J. Wooldridge. Mean-Payoff Games with ω -Regular Specifications. In Frank Dignum, Alessio Lomuscio, Ulle Endriss, and Ann Nowé, editors, *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, pages 1272–1280. ACM, 2021. URL: <https://dl.acm.org/doi/10.5555/3463952.3464099>.
- [UW11] 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.

- [VCD⁺15] Y. Verner, 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.
- [ZP96] 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.