Rational Synthesis in the Commons with Careless and Careful Agents

Rodica Condurache
A. I. Cuza University of Iaşi
Iaşi, Romania
rodica.condurache@info.uaic.ro

Catalin Dima
LACL – Université Paris-Est Créteil
Créteil, France
catalin.dima@u-pec.fr

Youssouf Oualhadj
LACL – Université Paris-Est Créteil
Créteil, France
youssouf.oualhadj@u-pec.fr

Nicolas Troquard
KRDB – Free University of Bozen-Bolzano
Bozen-Bolzano, Italy
nicolas.troquard@unibz.it

ABSTRACT

Turn-based games on graphs are games where the states are controlled by one and only one player who decides which edge to follow. Each player has a temporal objective that he tries to achieve. One player is the designated 'controller', whose objective captures the desirable outcomes of the whole system. Cooperative rational synthesis is the problem of computing a Nash equilibrium that satisfies the controller's objective. In this paper, we tackle this problem in the context of a commons, where each action has a cost or a benefit on one shared common pool energy resource. The paper investigates the problem of synthesising the controller in a commons such that there exists an individually rational behaviour of all the agents in the commons that satisfies the controller's objective and does not deplete the resource. We consider two types of agents: careless and careful. Careless agents only care for their temporal objective, while careful agents also pay attention not to deplete the system's resource. We solve the problem of cooperative rational synthesis in these games, focusing on parity objectives.

KEYWORDS

Games on graphs; rational synthesis; resources; Nash equilibria

ACM Reference Format:
Rodica Condurache, Catalin Dima, Youssouf Oualhadj, and Nicolas Troquard. 2021. Rational Synthesis in the Commons with Careless and Careful Agents. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

Common-pool resources are resources like water, air, coal, pastures, or fish stocks [21]. They are *non-excludable*: they are out there for the taking. They are *rivalrous*: one agent's consumption can limit or prevent another agent to consume it. *Energy* also is a kind of resource that typically can be framed as a commons [19]. In this paper, we will focus on the resource of energy, as it is understood in the literature in Computer Science. In fact, it is both more abstract, and much more than 'energy': it captures any common-pool resource that can be conveniently quantified by assigning a number to it, and where a bigger number indicates a greater amount of the resource.

For a long time, it was believed that common-pool resources were bound to collapse due to the actions of self-interested agents, causing the resources to be depleted or spoiled. The great contribution of E. Ostrom [21] was to evidence wide-ranging instances where this is not in fact the case, and to identify design principles of successful common-pool resource management. The models and algorithmic solutions presented here are contributions to the pursued efforts of engineering solutions for commons management. Specifically, we investigate the synthesis of a commons controller agent's strategy for which there exists a rational behaviour of all the agents in the commons, modelled as a Nash equilibrium, that satisfies the controller's objective. The algorithms presented in this paper can then serve to recommend a behaviour to all the agents in a commons, desirable from the point of view of the system, that the agents have individually no incentive to reject.

The setting. In this paper, the agents interact in turn-based fashion in an arena, which is a graph where in each state one agent decides the next edge to follow. To each edge is associated an integer which corresponds to the *energy cost* incurred by the whole system when it follows this edge. Each agent has a temporal objective, over the set of states. In this paper, we will concentrate our attention on parity objectives, which are canonical for representing all ω-regular properties. Büchi objectives are a particular case. Temporal objectives like reachability, safety, LTL, etc, can be translated into Büchi and parity objectives [3].

In addition to their temporal objective, the agents may be concerned with a quantitative objective as well. We consider two types of agents: *careless* and *careful*. Careless agents bother only about their temporal objective. Careful agents also pay attention to not deplete the system's resources.

Hence, a careless agent will want to satisfy his temporal objective, regardless of the level of the system's resources. On the other hand, a careful agent will want to satisfy her temporal objective in a run that never depletes the system's resources.

Decision problems. We study the *cooperative* synthesis problem [16]. In this problem, agent 1 holds the special role of *controller*, and is always careful. Cooperative synthesis consists in finding a strategy for agent 1 such that in the sub-game that it defines,

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

there exists a Nash Equilibrium whose outcome satisfies agent 1's objective, and never depletes the system's resources.

Main results. The results of computational complexity when the weights are represented in binary are summarized in Table 1. We also show that some of these problems are *not strongly complete* for their complexity class. Careful Büchi and Parity can be solved in NP when the weights are represented in unary.

	no energy	careless	careful
Büchi	PTIME-c [12]	PTIME-c	PSPACE-c
Parity	NP-c [12]	NP-c	PSPACE-c

Table 1: Complexity of the cooperative synthesis problem, when the weights are represented in binary. *careless*: only the controller is careful. *careful*: all players are careful.

Related work. The synthesis problem consists of automatically designing a controller for a system that will enforce a given specification. Introduced formally by Church [11], it puts two entities in opposition: a *system controller* and a *system environment.* A specification for the global system is chosen, and the goal of the synthesis is to automatically design a behavior (a strategy) for the system controller such that the specification holds against any behavior of the system environment. The first solution for this problem was presented in [6] for systems modeled by turn-based arenas and specifications described by ω-regular objectives.

One clear drawback of this mathematical framing, also known as zero-sum games, is to consider the controller and the system to be adversaries. This indeed entails rather pessimistic models and quite conservative solutions. In an effort to circumvent these forced antagonistic assumptions, Ummels [23] considered the synthesis problem in a setting where each agent is assigned an individual specification together with a list of the specification that 'must-hold'. In this case, one aims at constructing a global controller that ensures a rational behavior for the global system such that its outcome ensures the 'must-hold' specifications. Later Fisman et. al. [16] gave a logical characterization of this problem, stating it in terms of model checking with Strategy Logic. This problem is now known as the *cooperative rational synthesis*, as it assumes that the different part of the system will agree on a global behavior as long as some behavioural rationality is guaranteed, viz., no agent has an incentive to unilaterally defect from it. Finally Condurache et. al. [12] presented a complete picture of complexity bounds for a variety of ω-regular specifications.

In a different line of work Bouyer et. al. [24] studied the problem of existence of a Nash equilibrium in systems with qualitative specifications but the obtained results were of negative nature. The decidability was hard to obtain and quite strict restriction on the behaviour of the agents had to be made in order handle this case.

While mixing temporal and quantitative specifications is rather natural idea, it has been mostly studied in the setting of zero-sum games [7, 8, 10]. The model checking of resource-bounded logics of strategies has also been rather extensively investigated [1, 2, 20, 25]. A notable work presenting non-zero-sum games with temporal and quantitative specifications is [17], which considered a setting where each player is assigned a conjunction of temporal and quantitative goal. Still, the quantitative resources were private and not commons.

To prove the results of this paper, we are going to rely on more specific existing work. Two important extensions of LTL are useful in this work to obtain optimal algorithms for parity objectives. They play the role that plays vanilla LTL in [12]. Energy LTL [4] is used for the careless case, and Constrained LTL [14] is used for the careful case. In fact Energy LTL is a particular case of Constrained LTL. We use both to emphasise the crucial differences between the careless and the careful cases. Energy LTL permits to verify that an LTL property holds over a run whose energy does not drop below zero, which will be just enough for the careless case. In the careful case, one will also need to check as much, but one will also need to check that some states are not traversed with a level of energy too high to allow a profitable deviation from the agent controlling it. The extended language of Constrained LTL, allowing the comparison of the level of energy with any constant is providing the necessary formal machinery. Keeping the energy not only above zero, but below certain bounds in some states is thus crucial for the careful case. This will also be reflected by our use of bounded one-counter automata [15] to obtain optimal lower-bounds for our problems.

2 GAMES ON FINITE GRAPHS

For any set Q we denote by Q^* the set of finite sequences of elements in Q and Q^ω the set of infinite sequences of elements of Q. Let $w \in Q^* \cup Q^\omega$, and $i \geq 1$, we denote by $w[i]$ the i-th element in w; we denote by $w[..i]$ the prefix of w of size i and $w[i..]$ the suffix that starts at the i-th letter. For an element $q \in Q^*$, $\mathsf{lst}(q)$ is the last element in the sequence q.

2.1 Arenas, Strategies, and Profiles

Multi-player arenas. A *multi-player arena* is a tuple $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \ldots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$, where S is a finite set of states, $(\mathsf{S}_1 \uplus \ldots \uplus \mathsf{S}_n)$ is a partition of S, s_{ini} is an initial state, $\mathsf{P} = \{1, \ldots, n\}$ is the set of players, and E is in an edge relation in $\mathsf{S} \times \mathsf{S}$. For every edge $e = (s, t)$, $\mathsf{Src}(e)$ is s and $\mathsf{Trgt}(e)$ is t.

Plays and strategies. For an arena $\mathcal{G}$, we denote by $\mathsf{Plys}(\mathcal{G})$ the set of elements $s_{\mathsf{ini}} s_1 s_2 \ldots$ in S^ω such that for all $i \geq 0$, (s_i, s_{i+1}) is in E. The set $\mathsf{Hst}(\mathcal{G})$ is the set of finite and proper prefixes of elements in $\mathsf{Plys}(\mathcal{G})$. Moreover $\mathsf{Hst}_i(\mathcal{G})$ for i in P is the set of elements in $\mathsf{Hst}(\mathcal{G})$ whose last element is in S_i i.e., $\mathsf{Hst}_i(\mathcal{G}) = \{h \in \mathsf{Hst}(\mathcal{G}) \mid \mathsf{lst}(h) \in \mathsf{S}_i\}$. A *strategy* for player i is a function $\sigma_i \colon \mathsf{Hst}_i(\mathcal{G}) \to \mathsf{S}$ mapping a history whose last element is s to a state s' such that $(s, s') \in \mathsf{E}$. For a strategy σ_i for player i, we define the set $\langle \sigma_i \rangle$ as the set of plays that are *compatible* with σ_i i.e.,

$$\{\pi \in \mathsf{Plys}(\mathcal{G}) \mid \forall j \geq 0,\ \pi[..j] \in \mathsf{Hst}_i(\mathcal{G}) \implies \sigma_i(\pi[..j]) = \pi[j+1]\}$$

Profile of strategies. Once a strategy σ_i for each player i is chosen, we obtain a strategy profile $\overline{\sigma} = \langle \sigma_1, \ldots, \sigma_n \rangle$. $\overline{\sigma}_{\text{-}i}$ is the corresponding partial profile without the strategy for player i. For a strategy σ_i' for a player i, we write $\langle \overline{\sigma}_{\text{-}i}, \sigma_i' \rangle$ the profile $\langle \sigma_1, \ldots, \sigma_i', \ldots, \sigma_n \rangle$. We denote by $\langle \overline{\sigma} \rangle$ the unique outcome of the strategy profile $\overline{\sigma}$.

2.2 Objectives and Payoffs

An objective Obj is a subset of $\mathsf{Plys}(\mathcal{G})$. We write Obj_i to specify that it is the objective of player i. We define the payoff $\mathsf{Payoff}_i(\overline{\sigma})$

of player i wrt. the profile $\overline{\sigma}$ as follows: $\mathsf{Payoff}_i(\overline{\sigma}) = 1$ if $\langle\overline{\sigma}\rangle$ is in Obj_i and 0 otherwise. In the case where the arena consists of only two players, we can define zero-sum objectives, i.e. objectives that oppose for the players, i.e., $\forall i \in \{1, 2\},\ \mathsf{Obj}_{3-i} = \mathsf{Plys}(\mathcal{G}) \setminus \mathsf{Obj}_i$.

Once an arena $\mathcal{G}$ is equipped with an objective Obj_i for each player i, we will often call *game* the tuple $\langle \mathcal{G}, \mathsf{Obj}_1, \dots, \mathsf{Obj}_n\rangle$. When the objective is clear from the context we will simply write $\mathcal{G}$.

Thanks to the zero-sum nature, we can define the notion of a *winning* strategy for player i, i.e., a strategy σ_i s.t. $\langle\sigma_i\rangle$ is a subset of Obj_i.

Given a a multiplayer arena $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \dots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$, we write $\mathcal{G}^{i,\text{-}i}$ for the zero-sum game where player 1 is i, and player 2 is the coalition of the rest of players seen as one entity. Formally $\mathsf{S}_1 = \mathsf{S}_i$, $\mathsf{S}_2 = \bigcup_{j\neq i} \mathsf{S}_j$, $\mathsf{Obj}_1 = \mathsf{Obj}_i$, and $\mathsf{Obj}_2 = \mathsf{Plys}(\mathcal{G}) \setminus \mathsf{Obj}_1$.

Parity objectives. Let π in $\mathsf{Plys}(\mathcal{G})$, we denote by $\mathsf{Inf}(\pi)$ the set of states occurring infinitely often along π. Let C be a finite subset of $\mathbb{N}$, and let $\mathsf{prty}\colon \mathsf{S} \to C$ be a priority function. The parity objective for a game $\mathcal{G}$ equipped with the priority function prty is given by the set Parity defined as follows

$$\mathsf{Parity}(\mathcal{G}) = \big\{\pi \in \mathsf{Plys}(\mathcal{G}) \mid \min\{\mathsf{prty}(s) \mid s \in \mathsf{Inf}(\pi)\} \text{ is even}\big\}$$

Theorem 1 ([26]). *Deciding if player* 1 *has a winning strategy in a zero-sum parity game is in* NP ∩ co-NP.

The special case where $C = \{0, 1\}$ is called Büchi objective. In this case we denote by F the set of priority 0 states, the Büchi objectives requires that states in F are visited infinitely often.

Remark 2. *An important feature of parity objectives is the one of prefix-independence. The play π being in $\mathsf{Parity}(\mathcal{G})$ depends only on the infinite suffix. Formally, if $\pi = uv$ with $u \in \mathsf{S}^*$ and $v \in \mathsf{S}^\omega$, then $uv \in \mathsf{Parity}(\mathcal{G})$ iff $v \in \mathsf{Parity}(\mathcal{G})$.*

Theorem 3 ([9]). *We can decide in polynomial time whether player 1 has a winning strategy in a zero-sum Büchi game.*

Energy objectives. Let $\mathsf{cst}\colon \mathsf{E} \to \mathbb{Z}$ be a cost function. To lighten the notation, we write $\mathsf{cst}(s, t)$ instead of $\mathsf{cst}((s, t))$. Let $h = s_{\mathsf{ini}} \dots s_n$ be a history in $\mathsf{Hst}(\mathcal{G})$; we abusively write $\mathsf{cst}(h)$ to mean the extension of cst to histories that is: $\mathsf{cst}(h) = \mathsf{cst}(s_{\mathsf{ini}}, s_1) + \sum_{i=1}^{n-1} \mathsf{cst}(s_i, s_{i+1})$. The energy objective for a game $\mathcal{G}$ equipped with a cost function cst is given by the set Energy described as follows:

$$\mathsf{Energy}(\mathcal{G}) = \{\pi \in \mathsf{Plys}(\mathcal{G}) \mid \forall i \geq 1,\ \mathsf{cst}(\pi[..i]) \geq 0\}$$

We denote by W the largest absolute value that appears in cst, i.e. $\mathsf{W} = \max\{|c| \in \mathbb{Z} \mid \exists e \in \mathsf{E},\ \mathsf{cst}(e) = c\}$. Throughout the paper, values of cst are encoded in binary, thus W is exponential in its encoding which is $\log(\mathsf{W})$, with the exception of Section 4.3 where the weights are encoded in unary.

Energy-parity objectives. Let $\mathcal{G}$ be a zero-sum game equipped with both a priority function prty and a cost function cst, the energy parity objective EnergyParity for this game is given by the set $\mathsf{EnergyParity}(\mathcal{G}) = \mathsf{Energy}(\mathcal{G}) \cap \mathsf{Parity}(\mathcal{G})$. Given an energy-parity game and a state, *the initial credit* problem asks whether there exists an initial value for the energy such that the first player has a strategy to ensure both objectives. This problem was solved by Chatterjee et. al. [7].

Theorem 4 ([7]). *The initial credit problem can be solved in $O(|E|\ .\ D\ .\ |\mathsf{S}|^{D+2}\ .\ \mathsf{W})$ where D is the highest priority in the game.*

2.3 Solution Concept

We define in our setting the notion of equilibrium introduced by Nash. A Nash equilibrium is a profile of strategies in which no player could do better by unilaterally changing his strategy, provided that the other players keep their strategies unchanged. Here, player 1 has the distinguished role of controller. So as in [12], we will define the cooperative rational synthesis problem in terms of the *fixed* Nash Equilibrium, where we ignore player 1's deviations. A strategy profile $\overline{\sigma}$ is a fixed Nash equilibrium (f-NE) if for any strategy σ_i' for any player i in $\mathsf{P} \setminus \{1\}$ we have: $\mathsf{Payoff}_i(\overline{\sigma}) \geq \mathsf{Payoff}_i(\langle\overline{\sigma}_{\text{-}i}, \sigma_i'\rangle)$. We write $\mathsf{f\text{-}NE}(\mathcal{G})$ for the set of all the profiles that are fixed Nash equilibria in $\mathcal{G}$.

2.4 Rational Synthesis in the Commons

Cooperative rational synthesis. Given an arena $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \dots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ and objectives $\mathsf{Obj}_1, \dots, \mathsf{Obj}_n$ the *cooperative rational synthesis* problem is to decide whether there exists a profile $\overline{\sigma}$ such that $\overline{\sigma}$ is a f-NE and $\overline{\sigma}$ is in Obj_1.

We propose two quantitative extensions to the rational synthesis problem. In a careless context, only player 1 is concerned with the cost of a play, while in a careful while, all the players are.

Careless cooperative rational synthesis. Let $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \dots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ be a game, $\mathsf{cst}\colon \mathsf{E} \to \mathbb{Z}$ be a cost function, objectives $\mathsf{Obj}_1, \dots, \mathsf{Obj}_n$, and let $\overline{\sigma}$ be a strategy profile. Then $\overline{\sigma}$ is a solution to the careless cooperative rational synthesis problem if: $\langle\overline{\sigma}\rangle \in \mathsf{Energy}(\mathcal{G})$, and $\overline{\sigma} \in$ f-NE. We denote the set of all the solutions by $\underline{\mathsf{f\text{-}NE}}(\mathcal{G})$.

Careful cooperative rational synthesis. Let $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \dots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ be a game, $\mathsf{cst}\colon \mathsf{E} \to \mathbb{Z}$ be a cost function, objectives $\mathsf{Obj}_1, \dots, \mathsf{Obj}_n$, and let $\overline{\sigma}$ be a profile. Then $\overline{\sigma}$ is a solution to the careful cooperative rational problem if

$$\langle\overline{\sigma}\rangle \in \mathsf{Energy}(\mathcal{G}), \text{and } \forall \sigma_i' \text{ a strategy for player } i > 1,$$
$$\langle\sigma_{-i}, \sigma_i'\rangle \in \mathsf{Obj}_i(\mathcal{G}) \cap \mathsf{Energy}(\mathcal{G}) \implies \langle\overline{\sigma}\rangle \in \mathsf{Obj}_i(\mathcal{G})$$

We denote the set of all the solutions by $\overline{\mathsf{f\text{-}NE}}(\mathcal{G})$.

Example 5. *Consider the arena depicted in Fig. 1. Player* 1 *(circle) controls state a, his objective is given by the set of all the plays that ultimately reach state* $(\circ, \square)$. *Player* 2 *(square) controls state b, his objective is given by the set of all the plays that ultimately reach either state* $(\square)$ *or* $(\circ, \square)$. *Player* 3 *(diamond) controls state c, his objective is given by the set of all the plays that ultimately reach state* $(\diamond)$. *Clearly from state a, player* 1 *has to move the play to state b, but since the cost of this edge is* -1 *he has to take the self-loop in a at least once. Suppose player* 1 *takes the self-loop in a 3 times then goes to state b. Suppose further player* 2 *chooses to advance to state c. The current energy level is then* 1. *Now it is up to player* 3 *to decide where the play will end. Despite the fact that he meets his objective by going to state* $(\diamond)$, *in the careful setting player* 3 *cannot go to this state since the cost of that transition is* -2 *and this would bring the energy level below 0. Therefore, the only possible move is towards* $(\circ, \square)$. *The resulting strategy profile is a fixed Nash equilibrium. Indeed, player* 2 *has no incentive to deviate, and since player* 1 *meets his objective,* *the described strategy is a solution to the careful rational synthesis problem.*

Note that not all Nash equilibria for players 2 and 3 are solutions for the cooperative rational synthesis. For instance, if in state b, player 2 moves the play to ($\Box$)*, the resulting play is a Nash equilibrium but it is not in player 1's objective.*

One can also see that the careless synthesis problem has no solution. The idea is that any path which is in player 1's objective would have to take the transition from c to ($\circ, \Box$)*. But this transition is player 3's decision, who is unsatisfied by doing this and could deviate to state* ($\Diamond$) *which is beneficial since the new play is in his objective. The subtle difference with the careful setting is that this deviation does not agree with the careful nature of the players.*

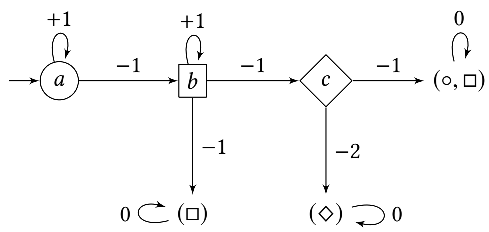

Figure 1: Game where there exists a careful solution, but no careless solution exists.

3 CARELESS SYNTHESIS

We provide a characterization of the careless cooperative synthesis in terms of Energy LTL model checking. We show that the problem is in NP for parity objectives and PTIME for Büchi objectives.

3.1 Logical Characterization

Our goal here is to provide a formal tool to recognize "good" outcomes, i.e. plays that are generated by profiles that are solutions to our problem. Our formal tool will heavily rely on a logical characterization using a quantitative version of the Linear Temporal Logic (LTL).

Linear Temporal Logic and its model checking. We assume the reader familiar with LTL and introduce the minimal setting that we will be using, the interested reader may find a detailed introduction in [3]. For our needs, we use the following shorthands: $\Diamond\phi = \top$ Until ϕ and $\Box\phi = \neg\Diamond\neg\phi$. The first one is satisfied by runs where ϕ is eventually true, the second one is satisfied by runs where ϕ always holds true.

eLTL. Since we are considering quantitative extension for the rational synthesis problem, the following extension of LTL will come in handy. Introduced in [4], it is evaluated over plays as follows: $\rho \models_{\mathsf{Energy}} \phi$ iff $\rho \models \phi$ and $\rho \in \mathsf{Energy}(\mathcal{G})$. Essentially, this logic is defined similarly to LTL but is evaluated over weighted runs and run ρ satisfies a formula ϕ if and only if it as a model for ϕ and $\mathsf{cst}(\rho[..n]) \geq 0$ for any $n \geq 1$.

Problem 6 (Energy LTL Model checking). *Given an arena $\mathcal{G}$, a cost function, and an LTL formula ϕ, decide whether there exists a play π in $\mathsf{Plys}(\mathcal{G})$ such that $\pi \models_{\mathsf{Energy}} \phi$.*

Theorem 7 ([5, 14]). *Problem 6 is* PSPACE*-complete.*

3.1.1 Reduction to eLTL model-checking. Let $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \ldots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ be an arena where each player i is given a parity objective Obj_i induced by a priority function prty_i and assume that $\mathcal{G}$ is endowed with a cost function $\mathsf{cst}\colon \mathsf{E} \to \mathbb{Z}$.

For each player i, we denote by $\underline{\mathsf{Win}}[\mathsf{Obj}_i]$ the set of states from where i has a winning strategy for his objective Obj_i against the coalition -i. This set is described by the propositional formula $\bigvee_{s \in \underline{\mathsf{Win}}[\mathsf{Obj}_i]} s$. For each player, the set $\underline{\mathsf{Win}}[\mathsf{Obj}_i]$ can be computed by any classical algorithm for the parity objectives, in non-deterministic polynomial time (Theorem 1). For each player i, we denote by Φ_{prty_i} the LTL formula that defines the set of plays in $\mathsf{Parity}_i(\mathcal{G})$. This formula is defined as follows:

$$\Phi_{\mathsf{prty}_i} \equiv \bigwedge_{\substack{s \in \mathsf{S} \\ \mathsf{prty}_i(s) \text{ is odd}}} \left(\Box\Diamond s \to \bigvee_{\substack{s' \in \mathsf{S} \\ \mathsf{prty}_i(s') < \mathsf{prty}_i(s) \\ \mathsf{prty}_i(s') \text{ is even}}} \Box\Diamond s'\right)$$

Consider the following formula:

$$\Phi_{\mathsf{f\text{-}NE}} \equiv \bigwedge_{i>1} \left(\neg\Phi_{\mathsf{prty}_i} \to \Box\neg\left(\bigvee_{s \in \underline{\mathsf{Win}}[\mathsf{Obj}_i]} s\right)\right)$$

The intuition behind the above formula is that along any play π a player that did not achieve his parity goal and never visited his winning region cannot deviate and improve his payoff. In [12], it is shown that any play satisfying the above formula is the outcome of some profile in $\mathsf{f\text{-}NE}(\mathcal{G})$.

Fact 1 ([12]). *Let $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \ldots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ be an arena where each player i is given an objective Obj_i induced by a priority function prty_i then $\forall\overline{\sigma}$, $\langle\overline{\sigma}\rangle \models \Phi_{\mathsf{f\text{-}NE}} \iff \overline{\sigma} \in \mathsf{f\text{-}NE}(\mathcal{G})$.*

We now take advantage of both Fact 1 and Theorem 7 and prove the following proposition.

Proposition 8. *Let $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \ldots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ be an arena where each player i is given a parity objective Parity_i. There exists a solution to the cooperative careless problem iff there exists π in $\mathsf{Plys}(\mathcal{G})$ such that: $\pi \models_{\mathsf{Energy}} \Phi_{\mathsf{prty}_1} \wedge \Phi_{\mathsf{f\text{-}NE}}$.*

Proof. Left to right; Let $\overline{\sigma}$ be a solution profile, let $\pi = \langle\overline{\sigma}\rangle$, then clearly enough, π is in $\mathsf{Obj}_1 \cap \mathsf{Energy}(\mathcal{G})$. The fact that $\pi \models \Phi_{\mathsf{f\text{-}NE}}$ is a direct consequence of the Fact 1.

Right to left; Let $\pi \in \mathsf{Plys}(\mathcal{G})$ s.t. $\pi \models_{\mathsf{Energy}} \Phi_{\mathsf{prty}_1} \wedge \Phi_{\mathsf{f\text{-}NE}}$. We will construct a profile $\overline{\sigma}$ that is a solution. In order to achieve this we introduce (recall) some useful notation. For all $h \in \mathsf{Hst}(\mathcal{G})$ s.t. h is not a prefix of π, let :

- $\bar{h}$ is the longest common prefix shared by h and π,
- p is the player s.t. $\mathsf{lst}(\bar{h}) \in \mathsf{S}_p$,
- $\sigma_{\text{-}p}$ is a winning strategy for the coalition -p in the zero-sum game played between p and -p with the objectives Obj_p, and $\mathsf{Plys}(\mathcal{G}) \setminus \mathsf{Obj}_p$. Note that this strategy must exist since π is always outside $\underline{\mathsf{Win}}[\mathsf{Obj}_p]$.

The profile $\overline{\sigma}$ is obtained as follows:

$$\sigma_j(s_{\mathsf{ini}}) = \bar{h}[1],\ \text{if } s_{\mathsf{ini}} \in S_j \tag{1}$$

$$\forall 1 \leq i < |\bar{h}|,\ \sigma_j(\bar{h}[..i]) = \bar{h}[i+1] \text{ if } \bar{h}[i] \in S_j \tag{2}$$

$$\forall i \geq |\bar{h}|, \forall j \neq p,\ \sigma_j(\pi[..i]) = \sigma_{\text{-}p}(\pi[..i]) \text{ if } \pi[i] \in S_j \tag{3}$$

$$\forall i \geq |\bar{h}|,\ \sigma_p(\pi[..i]) = * \text{ if } \pi[i] \in S_p \tag{4}$$

where $*$ means that player p can chose any available successor.

Let us show that $\overline{\sigma}$ is in $\underline{\mathsf{f\text{-}NE}}(\mathcal{G})$. If $\langle\overline{\sigma}\rangle = \pi$, then thanks to Fact 1, and the definition of Problem 6, π satisfies the objectives of player 1, and it is in $\mathsf{Energy}(\mathcal{G})$. Assume that along π there exists a player p whose payoff is 0. We will show that he cannot improve his payoff by deviating from π. Assume that p chooses to deviate after some prefix h, since π is a model for $\Phi_{\mathsf{f\text{-}NE}}$, this deviation happens from a state outside $\underline{\mathsf{Win}}[\mathsf{Obj}_p]$ according to a strategy τ, since the coalition -p uses a winning strategy $\sigma_{\text{-}p}$, the outcome of the couple $h' = \langle(\tau, \sigma_{\text{-}p})\rangle$ that start after the prefix h will have a payoff 0, now since parity objectives are prefix-independent, the full outcome that is hh' will also have a payoff 0 showing that τ is not a profitable deviation. □

3.2 The Parity Case is NP-Complete

We now present the main result of this section, namely a non deterministic polynomial procedure for the careless cooperative rational synthesis problem when all the players have a parity objective.

Let $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \ldots \uplus \mathsf{S}_n), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E}\rangle$ be an arena where each player i is given a parity objective Obj_i induced by a priority function prty_i and assume that $\mathcal{G}$ is endowed with a cost function $\mathsf{cst}\colon \mathsf{E} \longrightarrow \mathbb{Z}$. The rest of this section is dedicated to the proof of the following result:

Proposition 9. *Careless cooperative rational synthesis is in* NP *for Parity objectives.*

Proof. To decide if there is a strategy profile $\bar{\sigma}$ that is a solution to the careless cooperative synthesis problem, the non-deterministic algorithm proceeds as follows:

- Guess a payoff $\bar{z} \in \{0,1\}^n$ for all players and with $z_1 = 1$. We denote $I = \{i \leq n \mid z_i = 0\}$.
- For each player $i \in I$, guess a positional strategy $\tau_{\text{-}i}$ for the zero-sum game $\mathcal{G}^{i,-i}$.
- Guess a set $U \subseteq \bigcap_{i\in I}(S \setminus \underline{\mathsf{Win}}[\mathsf{Parity}_i])$. This set will serve for satisfying the parity conditions while maintaining the energy level.
- Guess the followings:
  i. A subset V s.t. $V \subseteq U$.
  ii. A cycle h_{par} of size $\leq 2\cdot|I|\cdot|S|$ in V that visits all the states of V. h_{par} is meant to satisfy the parity conditions.
  iii. A cycle h_{ch} of size $\leq |S|$ in U. h_{ch} is meant to be a "charging cycle", useful when h_{par} has a negative cost.
  iv. Two paths h^1_{lnk} and h^2_{lnk} with $h^1_{lnk}[1] = h^2_{lnk}[|h^2_{lnk}|] = h_{par}[1]$ and $h^2_{lnk}[1] = h^1_{lnk}[|h^1_{lnk}|] = h_{ch}[1]$ in U. These two paths are meant to connect h_{ch} with h_{par}.
  v. A cycle h_{acc} and two paths h_1 and h_2 of size $\leq |S|$ s.t. $h_1[1] = s_{ini}$, $h_1[|h_1|] = h_{acc}[1] = h_2[1]$ and $h_2[|h_2|] = h_{par}[1]$. These three items are needed for reaching h_{par} with sufficient energy stores and ensuring that the energy level is still non-negative.

Then the algorithm performs the following steps – in which any "check" step which does not succeed blocks the algorithm:

1. Check whether $\forall i$, $\tau_{\text{-}i}$ is winning for -i in the game $\mathcal{G}^{i,\text{-}i}$.
2. Check whether the set V is a subset of U, is strongly connected, and $\forall i \leq n$, $min\{\mathsf{prty}_i(s) \mid s \in U\}$ is odd if and only if $i \in I$.
3. Check that h_{par} satisfies the desired parity conditions: for each $i \notin I$, $min\{\mathsf{prty}_i(h_{par}[j]) \mid 1 \leq j \leq |h_{par}|\}$ is even.
4. Compute the costs $\mathsf{cst}(h_{par})$ and $\mathsf{cst}(h_{ch})$.
5. If $\mathsf{cst}(h_{par}) \geq 0$, then check that at least one of the following properties hold:
   (a) For each $j \leq |h_1h_2h_{par}|$ check that $\mathsf{cst}(h_1h_2h_{par}[..j]) \geq 0$.
   (b) For each $j \leq |h_1h_{acc}|$ check whether $\mathsf{cst}(h_1h_{acc}[..j]) \geq 0$.
6. If $\mathsf{cst}(h_{par}) < 0$, check whether $\mathsf{cst}(h_{ch}) > 0$ and further perform the iterative checks (5.a) and (5.b) above. Further check that h_{lnk} has a nonempty intersection with both h_{par} and h_{ch}.

First, note that each of the above witnesses are poly-size and the checks can be performed in polynomial time. We now argue the correctness of the algorithm. From the set $S \setminus \underline{\mathsf{Win}}[\mathsf{Parity}_i]$, no player $i \in I$ can deviate and win, because the others will play the strategy τ_{-i} against him. Therefore, on any path which stays in the set U, the players in I (i.e., with payoff 0) cannot deviate and improve their outcome.

Consider a play ρ which is generated by the following ω-regular expression:

$$h_1(h_{acc})^* h_2(h_{par})^\omega + h_1 h^*_{acc} h_2 \big(h_{par} h^1_{lnk} (h_{ch})^* h^2_{lnk}\big)^\omega$$

Note first that ρ satisfies the parity condition for each $i \notin I$. Some of these plays might not satisfy the energy conditions, but:

- When the checks at points (5) and (5.a) are satisfied, we have the guarantee that the path $h_1h_2(h_{par})^\omega$ satisfies the energy level.
- When the checks at points (5) and (5.b) are satisfied, we have the guarantee that there exists some $k \geq 0$ such that the path $h_1(h_{acc})^k h_2(h_{par})^\omega$ satisfies the energy level.
- When the checks at points (6) and (5.a) are satisfied, we have the guarantee that there exists some $l \geq 0$ such that the path $h_1h_2\big(h_{par}h^1_{lnk}(h_{ch})^l h^2_{lnk}\big)^\omega$ satisfies the energy level.
- When the checks at points (6) and (5.b) are satisfied, we have the guarantee that there exists some $k, l \geq 0$ such that the path $h_1(h_{acc})^k h_2\big(h_{par}h^1_{lnk}(h_{ch})^l h^2_{lnk}\big)^\omega$ satisfies the energy level.

Hence, in each of the above cases, the path is a model of the formula from Proposition 8. □

Proposition 10. *Careless cooperative rational synthesis is* NP-*hard for Parity objectives.*

Proof. The claim holds due to the fact that the cooperative setting from [12] is a particular case of our problem and it is NP-hard for parity objectives. □

3.3 A Polynomial Time Solution for the Büchi Case

In this section, we show that when the objectives of player are given by a Büchi condition, the complexity of the careless cooperative rational synthesis becomes more tractable. We recall that Büchi objective is the special case where the priorities are in $\{0, 1\}$. For each player i, F_i will denote the set of states $\{s \in \mathsf{S} \mid \mathsf{prty}_i(s) = 0\}$. Our result for Büchi objectives are:

Proposition 11. *Careless cooperative rational synthesis is in* PTIME *for Büchi objectives.*

Our proof will rely on an algorithm that computes Nash equilibria in muli-player games with Büchi objective [23]. We will also use an algorithm from [7] for solving one-player energy Büchi games.

We sketch a procedure to decide in polynomial time whether there is a strategy profile $\bar{\sigma}$ that is a solution to the careless cooperative synthesis with Büchi objectives. First the procedure decomposes the arena into strongly connected components (SCCs) and for each SCC C which is good for player 1, i.e., $F_1 \cap C \neq \emptyset$ we compute a winning strategy for the 1-player energy Büchi game induced by C and Obj_1. Since this strategy is with finite memory, it induces a lasso-shape run. We check that along this run, any player that loses cannot deviate, i.e., for any player i that loses, $\underline{\mathsf{Win}}[\mathsf{Obj}_i]$ is not visited. If this is the case then the run induced by this strategy is a solution. Indeed, this run satisfies the objective of player 1 and satisfies the energy constraint. The fact that it avoids $\underline{\mathsf{Win}}[\mathsf{Obj}_i]$ ensures for the players with payoff 0 ensures that it is a fixed Nash equilibrium. In case some states in $\underline{\mathsf{Win}}[\mathsf{Obj}_i]$ are visited, we remove them and repeat the procedure.

The polynomial time upper-bound follows from the fact that the decomposition into SCCs can be performed in polynomial time, that thanks to [7] we can compute a winning strategy for one-player energy Büchi games in polynomial time, that we can compute $\underline{\mathsf{Win}}[\mathsf{Obj}_i]$ for any $i > 1$ in polynomial time (c.f. [9]), and that after each recursive call, at least one state is removed.

4 CAREFUL SYNTHESIS

We now consider the case where all agents are careful about not depleting the resource. We propose a logical characterization in term of Constrained LTL model checking. This characterization is used to provide a PSPACE solution to the cooperative synthesis problem. The PSPACE lower bound is obtained by a reduction from the reachability problem in bounded one-counter automata. Finally, we show that the complexity drops to NP when the weights are represented in unary.

4.1 A Logical Characterization

1-CLTL and its model checking. Let x be an integer variable, we will present an extension of LTL that interprets x as a counter value and allows to write formulas constrained by x. These constraints will be described using the following grammar:

$$\alpha ::= x \sim d \mid \neg\alpha \mid \alpha \wedge \alpha \tag{5}$$

where x is the counter variable, $d \in \mathbb{Z}$, and $\sim \in \{<, >, \leq, \geq, =\}$.

The formulas of 1-CLTL are then:

$$\phi ::= \alpha \mid p \mid \neg\phi \mid \phi \vee \phi \mid \mathsf{X}\,\phi \mid \phi \;\mathsf{Until}\; \phi$$

where α is a counter constraint. Note that if we remove α from the above grammar the resulting logic is plain LTL.

Fix a set of atomic propositions PROP. The models of 1-CLTL(PROP) are pairs of mappings $\langle \mu_1, \mu_2 \rangle$. The mapping $\mu_1 : \mathbb{N} \to 2^{\mathsf{PROP}}$ indicates which are the atomic propositions true in every instant. The mapping $\mu_2 : \mathbb{N} \to \mathbb{Z}$ indicates the counter value at all instants.

In these models, propositional variables are evaluated wrt. μ_1 and counter constraints are evaluated wrt. μ_2 as follows:

$$\langle \mu_1, \mu_2 \rangle, i \models p \text{ iff } p \in \mu_1(i), \quad \langle \mu_1, \mu_2 \rangle, i \models x \sim d \text{ iff } \mu_2(i) \sim d\ .$$

The temporal operators are evaluated as in LTL.

One-counter automata. Is a tuple $\Gamma = (L, \delta, l_0)$, where L is a finite set of locations, δ is a set of transitions, and $l_0 \in L$ is the initial location. A transition in δ is a tuple (l, p, g, l'), where l and l' are locations, $p \in \mathbb{Z}$ is the weight of the transition, and g is a guard generated by the grammar of Equation (5). A run in a one-counter automaton is a pair $\langle \mu_1, \mu_2 \rangle$, where $\mu_1 : \mathbb{N} \to L$ and $\mu_2 : \mathbb{N} \to \mathbb{Z}$, and for every $i \geq 0$, if $\mu_1(i) = l$ and $\mu_2(i) = c$, $\mu_1(i+1) = l'$, $\mu_2(i+1) = c'$, then there is $(l, p, g, l') \in \delta$ such that $c' = c + p$ and $\langle \mu_1, \mu_2 \rangle, i \models g$.

Problem 12 (1-CLTL Model checking). *Given a one-counter automaton* Γ *and a 1-CLTL formula* ϕ*, decide whether there exists a run* $\langle \mu_1, \mu_2 \rangle$ *such that* $\langle \mu_1, \mu_2 \rangle \models \phi$.

1-CLTL is a particular case of Constrained LTL with propositional variables, one integer variable, and one-step look-ahead, whose model checking is PSPACE-complete [14].

Theorem 13. *Problem 12 is* PSPACE*-complete.*

Reduction to 1-CLTL model checking. Similarly to the careless case, we will express the existence of a solution to the careful synthesis problem as a model checking question. First we define a mapping $\mathsf{Credit} : \mathsf{S} \times \mathsf{P} \to \mathbb{N} \cup \{\omega\}$, that associates with each couple of state and player the minimal initial credit to meet its objective against the coalition with a positive energy, ω means that the player cannot win with any initial credit. The minimal initial credit is computed thanks to the algorithm used in [7] that computes for each state a minimal credit and a winning strategy in an energy-parity game for the first player (protagonist).

Let $\overline{\mathsf{Win}}[\mathsf{Obj}_i] = \{s \in \mathsf{S} \mid \mathsf{Credit}(s, i) \neq \omega\}$.

Extended play. Let π be a play in $\mathsf{Plys}(\mathcal{G})$, define the extended play in the one-counter automaton:

$$\mu_{\pi_1}(i) = \pi[i+1], \quad \mu_{\pi_2}(0) = 0, \quad \mu_{\pi_2}(i+1) = \mathsf{cst}(\pi[..(i+1)])\ .$$

We also define the following 1-CLTL formula:

$$\Phi_{\overline{\text{f-NE}}} \equiv \bigwedge_{i>1} \left(\neg\Phi_{\mathsf{prty}_i} \to \Box\neg \left(\bigvee_{s \in \overline{\mathsf{Win}}[\mathsf{Obj}_i]} s \wedge (x \geq \mathsf{Credit}(s, i)) \right) \right)$$

Now using a construction analogous to the one presented in the proof of Proposition 8, we can state the two following propositions.

Proposition 14. *Let* $\langle \mu_{\pi_1}, \mu_{\pi_2} \rangle$ *be the extended play of some play* π *in* $\mathcal{G}$ *such that* $\langle \mu_{\pi_1}, \mu_{\pi_2} \rangle \models \Phi_{\overline{\text{f-NE}}}$*, then* π *is the outcome of an* $\overline{\text{f-NE}}$.

Proposition 15. *There exists a solution to the cooperative careful rational synthesis problem iff there exists a play π such that:*

$$\langle \mu_{\pi_1}, \mu_{\pi_2} \rangle \models \Phi_{\mathsf{prty}_1} \wedge \Phi_{\overline{\mathsf{f\text{-}NE}}} \wedge \Box(x \geq 0)$$

They are directly instrumental to establish the complexity of the cooperative synthesis problem.

4.2 PSPACE-Completeness

Proposition 16. *Careful cooperative rational synthesis is* PSPACE-*complete for Büchi and Parity objectives when the weights are encoded in binary.*

The PSPACE membership follows immediately from the poly-size logical characterization of solutions using 1-CLTL and Theorem 13. On the other hand, the hardness will be established by reducing the problem of reachability in bounded one-counter automata [15]. A *bounded one-counter automaton* Γ is a tuple (L, b, δ, l_0), where (L, δ, l_0) is a one-counter automaton, and $b \in \mathbb{N}$ is a counter bound. A transition in δ is of the form $(l, p, x \geq d_1 \wedge x \leq d_2, l')$ where, $p \in \{-b, \ldots, b\}$, and $d_1, d_2 \in \{0, \ldots, b\}$. The *reachability problem in bounded one-counter automata* asks, given a bounded one-counter automaton $\Gamma = (L, b, \delta, l_0)$ and a location $t \in L$, whether there is a run $\langle \mu_1, \mu_2 \rangle$ and an $i \in \mathbb{N}$ such that $\mu_1(i) = t$, and $\mu_2(i) = 0$. [15, Corollary 12] states that the reachability problem in bounded one-counter automata is PSPACE-complete.

Proof of hardness. Let $\Gamma = (L, b, \delta, l_0)$ be a bounded one-counter automaton, and let $t \in L$ be a target location in Γ.

We construct a multi-player reachability game with the following intuitions. Every location of Γ is a state controlled by Player 1. We create the states w_1, w_2, with only outgoing edge a self-loop of cost 0. The qualitative objective of Player 1 (resp. 2) is to reach the state w_1 (resp. w_2).

To ensure that the solutions stay within the bounds, for every transition $(l, p, x \geq d_1 \wedge x \leq d_2, l') \in \delta$, we create two fresh states $l_<$ and $l_>$ both controlled by Player 2, and four transitions. One transition from l to $l_>$, with cost p, serves to update the counter/energy level. One transition from $l_>$ to w_2, with cost $-(d_2 + 1)$. Player 2 taking this transition makes sure he meets his qualitative objective, but respects his carefulness iff the energy level is above the guard's upper bound d_2. One transition from $l_>$ to $l_<$, with cost $-(d_1 + 1)$, that takes the energy level below zero if it is already below the guard's lower bound d_1. One transition from $l_<$ to l', with cost $+(d_1 + 1)$, serves to reset the energy level to what it was before entering in $l_<$ (that is, the energy level out of l). The main gadgets of the above reduction are depicted in Figure 2. To ensure that the

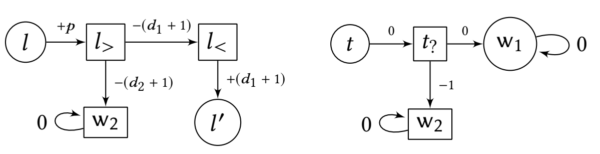

Figure 2: Gadgets for the reduction

solutions end with a zero-value counter, we also create a transition from the target t to a fresh state $t_?$ controlled by Player 2, and a transition from $t_?$ to w_1, both of cost 0. From the newly created state $t_?$ we also create a transition to w_2, of cost -1. Player 2 taking this transition meets his qualitative objective, but respects his carefulness iff the energy level is not 0. Player 2 has no incentive to choose this transition over going to w_1 iff the energy level at t is 0. Formally, we define $\mathcal{G} = \langle \mathsf{S}, (\mathsf{S}_1 \uplus \mathsf{S}_2), s_{\mathsf{ini}}, \mathsf{P}, \mathsf{E} \rangle$, where:

- $\mathsf{S} = L \cup \{l^\tau_<, l^\tau_> \mid \tau \in \delta\} \cup \{t_?\} \cup \{\mathsf{w}_1, \mathsf{w}_2\}$
- $\mathsf{S}_1 = L \cup \{\mathsf{w}_1\}$
- $\mathsf{S}_2 = \{l^\tau_<, l^\tau_> \mid \tau \in \delta\} \cup \{t_?\} \cup \{\mathsf{w}_2\}$
- $s_{\mathsf{ini}} = l_0$
- $\mathsf{P} = \{1, 2\}$
- $\mathsf{E} = \{(l, l^\tau_>), (l^\tau_>, l^\tau_<), (l^\tau_<, l'), (l^\tau_>, \mathsf{w}_2) \mid \tau \in \delta\} \cup \{(\mathsf{w}_1, \mathsf{w}_1), (\mathsf{w}_2, \mathsf{w}_2)\} \cup \{(t, t_?), (t_?, \mathsf{w}_2), (t_?, \mathsf{w}_1)\}$

The cost function of the edges in $\mathcal{G}$ is as follows:

- For every transition $\tau = (l, p, x \geq d_1 \wedge x \leq d_2, l') \in \delta$:
  - $\mathsf{cst}(l, l^\tau_>) = p$
  - $\mathsf{cst}(l^\tau_>, l^\tau_<) = -(d_1 + 1)$
  - $\mathsf{cst}(l^\tau_<, l') = +(d_1 + 1)$
  - $\mathsf{cst}(l^\tau_>, \mathsf{w}_2) = +(d_2 + 1)$
- $\mathsf{cst}(\mathsf{w}_1, \mathsf{w}_1) = \mathsf{cst}(\mathsf{w}_2, \mathsf{w}_2) = 0$
- $\mathsf{cst}(t, t_?) = \mathsf{cst}(t_?, \mathsf{w}_1) = 0$
- $\mathsf{cst}(t_?, \mathsf{w}_2) = -1$

Let the qualitative objectives of the players be defined as reachability objectives as follows:

- $\mathsf{Obj}_1 = \{\rho \mid \rho \in \mathsf{Plys}(\mathcal{G}), \mathsf{w}_1 \in \rho\}$
- $\mathsf{Obj}_2 = \{\rho \mid \rho \in \mathsf{Plys}(\mathcal{G}), \mathsf{w}_2 \in \rho\}$

The reachability problem for the bounded one-counter automata $\Gamma = (L, b, \delta, l_0)$ and a location $t \in L$ is true iff there exists a solution to the cooperative careful rational synthesis problem in $\mathcal{G}$ (with initial credit 0).

For Parity and Büchi, let $\mathsf{prty}_1(\mathsf{w}_1) = 0$ and $\mathsf{prty}_1(s) = 1$ for every $s \in \mathsf{S} \setminus \{\mathsf{w}_1\}$, and let $\mathsf{prty}_2(\mathsf{w}_2) = 0$ and $\mathsf{prty}_2(s) = 1$ for every $s \in \mathsf{S} \setminus \{\mathsf{w}_2\}$. Player i reaches W_i iff he meets his parity (Büchi) objective. □

4.3 The Unary Case is in NP

Proposition 17. *Careful cooperative rational synthesis is in* NP *for Parity objectives when weights are encoded in unary.*

Proof. We show how to decide, in NP, whether there is a strategy profile $\bar{\sigma}$ that is a solution to the careful cooperative synthesis for a given arena $\mathcal{G}$, a cost function cst and a set of parity objectives $(\mathsf{Parity}_i)_{1 \leq i \leq n}$. We assume, w.l.o.g., that the cost of each edge is in $\{-1, 0, 1\}$, by eventually transforming each edge $(s, t) \in E$ with cost $\mathsf{cst}(s, t) > 1$ in the initial arena into a sequence of edges with cost +1 which connect s with t and pass through new intermediate states owned by the same agent as s, and a similar transformation for edges with negative costs. Note that this transformation produces an arena whose size is $W \cdot |S|$ where W is the largest cost of an edge in the initial arena.

Recall that, for each player i and state $s \in S$ we denote by $\mathcal{G}^{i,-i}_s$ the underlying 2-player zero-sum parity game with the same game graph as $\mathcal{G}$, in which i is the protagonist (who's winning condition is Parity_i) and the initial state is s. Also we denote D_i the maximal priority for agent i.

In the sequel we will utilize the following theorem which summarizes Lemma 6 and Theorem 1 from [7]:

Theorem 18. *Assume that $\mathcal{G}$ is an energy-parity game in which the cost of each edge is in $\{-1, 0, 1\}$. Then:*

(1) *If player 1 has a winning strategy then he has a winning strategy with memory of size $|S| \cdot D$ and initial credit $(|S| - 1)$, where S is the set of states and D is the maximal priority.*
(2) *If player 2 has a winning strategy then he has a memoryless winning strategy.*

We also utilize here the notion of *simple 1-counter multi-parity automata* which are tuples $\mathcal{A} = (L, \theta, l_0, (\mathsf{prty}_i, \mathsf{Parity}_i)_{i \leq n})$ in which (L, θ, l_0) is a 1-counter automaton, $n \in \mathbb{N}$ and for each for each $i \leq n$, $\mathsf{prty}_i : S \to C$ is a priority mapping and Parity_i is the parity objective induced by prty_i. Note that these counter automata are not bounded in the sense of subsection 4.2.

Proposition 19. *The emptiness problem for simple 1-counter multi-parity automata is equivalent with the problem of finding a lasso path ρ of size $\leq 2 \cdot (|S|^3 + (n + 2) \cdot |S|^2)$ which satisfies each parity objective.*

This result is a generalization of Theorem 6 from [14], which requires adapting the proofs of Lemmas 13–19 of the same paper (and the complexity analysis provided there) in order to handle multi-parity acceptance conditions, instead of just Büchi acceptance as in [14]. The proof goes by showing that, if there exists an accepting run with infinitely many zero-tests, then there exists such a run in which the counter values never exceed $2(|S|^2 + n \cdot |S|)$. On the other hand, if there exists an accepting run having only finitely many zero-tests, then there exists a lasso-type accepting run $\rho = \rho_1 \cdot \rho_2^\omega$ such that in ρ_1 the counter values never exceed $|S|^2$, ρ_2 has no zero-tests and $|\rho_2| \leq (|S|^3 + n \cdot |S|^2)$. The details are omitted due to lack of space.

The non-deterministic algorithm for solving the careful cooperative rational synthesis is the following:

A. We guess a payoff vector $\bar{z} \in \{0, 1\}^n$, where $z_1 = 1$, with the intended meaning that $z_i = 1$ iff Player i wins in $\bar{\sigma}$. We further denote $I = \{i \leq n \mid z_i = 0\}$ the set of agents which are guessed as "losing".

B. For each player $i \in I$ and state $s \in S$ we guess the following:
 (1) An integer $0 \leq w_{s,i} \leq |S| + 1$,
 (2) If $w_{s,i} \leq |S|$, a strategy τ_i with memory $D_i \cdot |S|$ for the protagonist in $\mathcal{G}_s^{i,-i}$.
 (3) If $w_{s,i} > 0$, a positional strategy τ_{-i} for the antagonist in $\mathcal{G}_s^{i,-i}$.

 Then check the following properties:
 i. τ_i is a winning strategy for the protagonist in $\mathcal{G}_s^{i,-i}$ for the initial energy level $w_{s,i} - 1$. This check is not performed when $w_{s,i} = 0$. Note that, in case this check succeeds, $w_{s,i} \geq \mathsf{Credit}(s, i)$.
 ii. τ_{-i} is a wining strategy for the antagonist in $\mathcal{G}_s^{i,-i}$, for the initial energy level $w_{s,i}$. This check is not performed when $w_{s,i} = |S| + 1$. Note that, in case this check succeeds, $w_{s,i} < \mathsf{Credit}(s, i)$.

 Theorem 18 ensures that there exists at least one choice at step (B.) that satisfies these properties.

C. Guess in $\mathcal{G}$ a lasso path ρ with size smaller than $2(|S|^3 + (|I| + 2)|S|^2)$, which satisfies each parity condition for each $i \in I$ and such that, whenever ρ passes through a state s with $w_{s,i} \leq |S|$, then the energy level for that passage is $\leq |S|$.

To do this, we build the simple one-counter multi-parity automaton $\mathcal{A}_I = (S', (s_0, 0), \delta, (\mathsf{Parity}'_i)_{i \in I})$ where:

- $S' = (S \cup E) \times \{0, \ldots, |S|\}$.
- For each $(s, t) \in E$ with $\mathsf{cst}(s, t) = +1$ we add $((s, t), 0) \xrightarrow{+1} (t, 0) \in \delta$.
- For each $(s, t) \in E$ with $\mathsf{cst}(s, t) = -1$ we add $((s, t), 0) \xrightarrow{\geq 0?, -1} (t, 0) \in \delta$.
- For each state $s \in S$, denote $w_s = \min\{w_{s,i} \mid w_{s,i} \leq |S|, i \in I\}$ Then:

(S1) If $w_s \neq \infty$, then for each edge $(s, t) \in E$ and $0 \leq j < w_s$, we add transitions $(s, j) \xrightarrow{\geq 0?, -1} (s, j + 1) \in \delta$, $(s, w_s) \xrightarrow{=0?} ((s, t), w_s) \in \delta$ and $((s, t), j + 1) \xrightarrow{+1} ((s, t), j) \in \delta$.
(S2) Otherwise for each edge $(s, t) \in E$ append transition $(s, 0) \xrightarrow{\geq 0?} ((s, t), 0) \in \delta$.

- $\mathsf{Parity}'_i(s, j) = \mathsf{Parity}'_i((t, s), j) = \mathsf{Parity}_i(s)$ for any $s \in S$, $(t, s) \in E$, $i \in I$ and $0 \leq j \leq |S|$.

Then we guess a lasso path ρ' in $\mathcal{A}_I$ of size smaller than $2(|S|^3 + (|I| + 2)|S|^2)$ which satisfies each parity condition. According with Proposition 19, building such a lasso path is equivalent with checking emptiness of the automaton $\mathcal{A}_I$, which, in our case, means the existence of a play in $\mathcal{G}$ which satisfies the formula $\Phi_{\overline{\mathsf{f\text{-}NE}}}$. Then, from ρ' we build a lasso path ρ in $\mathcal{G}_s^{i,-i}$ by simply skipping the transitions of type (S1), which serve for checking whether the energy level is smaller than w_s while passing through state s. □

5 DISCUSSION AND CONCLUSION

We modelled commons as a multi-player game on graphs where each transition incurs a cost or a benefit on a unique shared resource. We studied the problem of cooperative rational synthesis in this setting. It aims at finding a Nash equilibrium that satisfies the controller's objective that can be played forever without depleting the resource.

We considered two cases, where agents are either careless or careful about the resource. We showed that in the careless case, the problem is NP-complete for parity objectives, and showed that it is tractable (PTIME-complete) for Büchi objectives. In the careful case, we showed that the problem is PSPACE-complete for parity and Büchi objectives. We also showed that it is in NP when the amount of the resource is represented as a bag of units.

Other results can be derived. In the careful case, the proof of Proposition 16 already establishes that the problem is PSPACE-hard for reachability objectives, and it can be straightforwardly extended to show that the problem is PSPACE-hard for co-Büchi and Streett objectives.

With LTL objectives, a corollary of our results is the 3EXPTIME membership for both settings. We are currently working on bridging the gap with the 2EXPTIME lower bound which is easily obtained from the original LTL synthesis problem [22].

We also plan to investigate the non-cooperative synthesis problem [12, 13, 18] in the commons.

REFERENCES

[1] N. Alechina, N. Bulling, S. Demri, and B. Logan. 2018. On the complexity of resource-bounded logics. *Theoretical Computer Science* 750 (2018), 69 – 100. https://doi.org/10.1016/j.tcs.2018.01.019 Reachability Problems: Special Issue.

[2] Natasha Alechina, Stéphane Demri, and Brian Logan. 2020. Parameterised Resource-Bounded ATL. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 7040–7046. https://aaai.org/ojs/index.php/AAAI/article/view/6189

[3] Christel Baier and Joost-Pieter Katoen. 2008. *Principles of model checking*. MIT Press.

[4] Aaron Bohy, Véronique Bruyère, Emmanuel Filiot, and Jean-François Raskin. 2013. Synthesis from LTL Specifications with Mean-Payoff Objectives. In *Tools and Algorithms for the Construction and Analysis of Systems - 19th International Conference, TACAS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7795)*, Nir Piterman and Scott A. Smolka (Eds.). Springer, 169–184. https://doi.org/10.1007/978-3-642-36742-7_12

[5] Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, and Orna Kupferman. 2014. Temporal Specifications with Accumulative Values. *ACM Trans. Comput. Log.* 15, 4 (2014), 27:1–27:25. https://doi.org/10.1145/2629686

[6] J. Richard Büchi and Lawrence H. Landweber. 1969. Definability in the Monadic Second-Order Theory of Successor. *J. Symb. Log.* 34, 2 (1969), 166–170.

[7] Krishnendu Chatterjee and Laurent Doyen. 2012. Energy parity games. *Theor. Comput. Sci.* 458 (2012), 49–60. https://doi.org/10.1016/j.tcs.2012.07.038

[8] Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski. 2005. Mean-Payoff Parity Games. In *20th IEEE Symposium on Logic in Computer Science (LICS 2005), 26-29 June 2005, Chicago, IL, USA, Proceedings*. IEEE Computer Society, 178–187. https://doi.org/10.1109/LICS.2005.26

[9] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. 2008. Algorithms for Büchi Games. *CoRR* abs/0805.2620 (2008). arXiv:0805.2620 http://arxiv.org/abs/0805.2620

[10] Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. 2012. Strategy Synthesis for Multi-Dimensional Quantitative Objectives. In *CONCUR 2012 - Concurrency Theory - 23rd International Conference, CONCUR 2012, Newcastle upon Tyne, UK, September 4-7, 2012. Proceedings (Lecture Notes in Computer Science, Vol. 7454)*, Maciej Koutny and Irek Ulidowski (Eds.). Springer, 115–131. https://doi.org/10.1007/978-3-642-32940-1_10

[11] Alonzo Church. 1957. Applications of recursive arithmetic to the problem of circuit synthesis. *Summaries of the Summer Institute of Symbolic Logic* 1 (1957), 3–50.

[12] Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. 2016. The Complexity of Rational Synthesis. In *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy (LIPIcs, Vol. 55)*, Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi (Eds.). Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 121:1–121:15. https://doi.org/10.4230/LIPIcs.ICALP.2016.121

[13] Rodica Condurache, Youssouf Oualhadj, and Nicolas Troquard. 2018. The Complexity of Rational Synthesis for Concurrent Games. In *29th International Conference on Concurrency Theory, CONCUR 2018, September 4-7, 2018, Beijing, China (LIPIcs, Vol. 118)*, Sven Schewe and Lijun Zhang (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 38:1–38:15. https://doi.org/10.4230/LIPIcs.CONCUR.2018.38

[14] Stéphane Demri and Régis Gascon. 2009. The Effects of Bounding Syntactic Resources on Presburger LTL. *J. Log. Comput.* 19, 6 (2009), 1541–1575. https://doi.org/10.1093/logcom/exp037

[15] John Fearnley and Marcin Jurdzinski. 2015. Reachability in two-clock timed automata is PSPACE-complete. *Inf. Comput.* 243 (2015), 26–36. https://doi.org/10.1016/j.ic.2014.12.004

[16] Dana Fisman, Orna Kupferman, and Yoad Lustig. 2010. Rational Synthesis. In *Tools and Algorithms for the Construction and Analysis of Systems, 16th International Conference, TACAS 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6015)*, Javier Esparza and Rupak Majumdar (Eds.). Springer, 190–204. https://doi.org/10.1007/978-3-642-12002-2_16

[17] Julian Gutierrez, Aniello Murano, Giuseppe Perelli, Sasha Rubin, and Michael Wooldridge. 2017. Nash Equilibria in Concurrent Games with Lexicographic Preferences. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 1067–1073. https://doi.org/10.24963/ijcai.2017/148

[18] Orna Kupferman, Giuseppe Perelli, and Moshe Y. Vardi. 2016. Synthesis with rational environments. *Ann. Math. Artif. Intell.* 78, 1 (2016), 3–20. https://doi.org/10.1007/s10472-016-9508-8

[19] Emilia Melville, Ian Christie, Kate Burningham, Celia Way, and Phil Hampshire. 2017. The electric commons: A qualitative study of community accountability. *Energy Policy* 106 (2017), 12–21. https://doi.org/10.1016/j.enpol.2017.03.035

[20] Hoang Nga Nguyen, Natasha Alechina, Brian Logan, and Abdur Rakib. 2018. Alternating-time temporal logic with resource bounds. *J. Log. Comput.* 28, 4 (2018), 631–663. https://doi.org/10.1093/logcom/exv034

[21] Elinor Ostrom. 1990. *Governing the Commons: The evolution of institutions for collective action*. Cambridge University Press.

[22] Amir Pnueli and Roni Rosner. 1989. On the Synthesis of a Reactive Module. In *Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 11-13, 1989*. ACM Press, 179–190. https://doi.org/10.1145/75277.75293

[23] Michael Ummels. 2008. The Complexity of Nash Equilibria in Infinite Multiplayer Games. In *Foundations of Software Science and Computational Structures*, Roberto Amadio (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 20–34.

[24] Michael Ummels and Dominik Wojtczak. 2011. The Complexity of Nash Equilibria in Limit-Average Games. In *CONCUR 2011 - Concurrency Theory - 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings (Lecture Notes in Computer Science, Vol. 6901)*, Joost-Pieter Katoen and Barbara König (Eds.). Springer, 482–496. https://doi.org/10.1007/978-3-642-23217-6_32

[25] Steen Vester. 2015. On the Complexity of Model-Checking Branching and Alternating-Time Temporal Logics in One-Counter Systems. In *Automated Technology for Verification and Analysis - 13th International Symposium, ATVA 2015, Shanghai, China, October 12-15, 2015, Proceedings (Lecture Notes in Computer Science, Vol. 9364)*, Bernd Finkbeiner, Geguang Pu, and Lijun Zhang (Eds.). Springer, 361–377. https://doi.org/10.1007/978-3-319-24953-7_27

[26] Wieslaw Zielonka. 1998. Infinite Games on Finitely Coloured Graphs with Applications to Automata on Infinite Trees. *Theor. Comput. Sci.* 200, 1-2 (1998), 135–183. https://doi.org/10.1016/S0304-3975(98)00009-7