

Careful Autonomous Agents in Environments With Multiple Common Resources

Rodica Condurache

“Alexandru Ioan Cuza” University of Iași, Romania
rodica.condurache@info.uaic.ro

Catalin Dima

Univ. Paris Est Créteil, LACL, F-94010 Creteil, France
dima@u-pec.fr

Madalina Jitaru

Univ. Paris Est Créteil, LACL, F-94010 Creteil, France
madalina.jitaru@u-pec.fr

Youssef Oualhadj

Free University of Bozen-Bolzano, Italy
Univ. Paris Est Créteil, LACL, F-94010 Creteil, France
youssef.oualhadj@unibz.it

Nicolas Troquard

Free University of Bozen-Bolzano, Italy
nicolas.troquard@unibz.it

Careful rational synthesis was defined in [7] as a quantitative extension of Fisman et al.’s rational synthesis [11], as a model of multi-agent systems in which agents are interacting in a graph arena in a turn-based fashion. There is one common resource, and each action may decrease or increase the resource. Each agent has a temporal qualitative objective and wants to maintain the value of the resource positive. One must find a Nash equilibrium. This problem is decidable.

In more practical settings, the verification of the critical properties of multi-agent systems calls for models with many resources. Indeed, agents and robots consume and produce more than one type of resource: electric energy, fuel, raw material, manufactured goods, etc. We thus explore the problem of careful rational synthesis with several resources. We show that the problem is undecidable. We then propose a variant with bounded resources, motivated by the observation that in practical settings, the storage of resources is limited. We show that the problem becomes decidable, and is no harder than controller synthesis with Linear-time Temporal Logic objectives.

1 Introduction

The presence of autonomous agents in modern societies has become commonplace. We interact with them every day, and they may be of different levels of autonomy, e.g., self-checkout, chatbots, robot vacuum cleaners, or virtual assistants. A current tendency is that agents are intruding on the physical world, and robots are expanding their territory beyond their confined industrial environment.

The access to the resources necessary for an agent to accomplish his tasks could have been simply assumed in many application domains before: direct wire to an electricity source, a human operator providing raw material, etc. Nowadays, typical agents must be more autonomous than before in managing the multiple resources they need. They must carefully consume them, and in presence of competitors, they must also be careful in how they produce them.

Linear-time Temporal Logic (LTL) [16] has been a very popular logic for specifying temporal properties of systems. Planning with objectives expressed in some temporal logic has been well studied [3, 9, 4, 6]. Some logics have also been proposed to explicitly verify the properties of multiagent systems in presence of resource constraints [5, 15, 2]. When agents roam more freely the physical world,

they are more likely to compete with other agents, human or artificial, which may have conflicting goals. When planning in such environments an agent needs to adapt his behaviour to the capabilities and goals of others. A solution to a multiagent planning problem in this setting is a non-cooperative strategic equilibrium: a vector of strategies, one for each agent, such that no individual agent can be better off by unilaterally changing their strategy. This is what has come to be known as a Nash equilibrium [14].

This paper aims to contribute to the line of research interested in the formal verification of the existence of Nash equilibria in a multiagent system [18, 11, 8, 1]. When there is a solution Nash equilibrium, the techniques used can actually return a multiagent plan that satisfies the requisites. The paper has a special focus to consider agents that must be autonomous in an environment with multiple common resources, to bring the theory closer to the reality that engineers are working with.

In [7], the problem of careful rational synthesis is defined as a quantitative variant of rational synthesis [11]. Agents interact in a graph arena in a turn-based fashion. Each state is controlled by one and only one agent who decides which edge to follow. Each agent has a temporal objective that he tries to achieve. There is one integer common resource, and each action may decrease or increase the resource. The rational synthesis problem consists in computing a Nash equilibrium that satisfies a system objective. It is shown that in presence of one common resource, deciding the existence of a strategic equilibrium for careful autonomous agents (with *parity* objectives, a canonical representation of temporal properties on infinite traces [10]) can be solved in polynomial space. With LTL objectives, the problem can be solved in doubly exponential space.

But in real-case scenarios, physical agents are operating in a world where there is more than one resource. In this paper, we explore the problem of careful rational synthesis with several common resources.

Example 1. *Consider the game with 2 resources illustrated on Figure 1. Players 1, 2 and 3 control the states a , b , and c respectively. The other states are controlled by Player 1 (but note that the agent who controls them is irrelevant). Player 1 wants to reach a state with \circ , Player 2 wants to reach a state with \square , Player 3 wants to reach a state with \diamond . All of them want to keep the resources in check: they would be dissatisfied if any of the resources were to go below zero. The objective of the system is \circ . A solution to the synthesis problem is thus a Nash equilibrium that reaches the state (\circ, \square) , and never depletes the resources.*

One starts with the resources being $(0,0)$. Player 1 must pump thrice on a , which brings the resources to $(6,3)$. (Only he can increase resource two, and at least an amount of 3 is necessary to reach his objective and the objective of the system.) Player 1 can then go to b , which brings the resources to $(6,2)$.

At that point, Player 2 could go down. This would be the outcome of a Nash equilibrium, but it would not be a solution to our synthesis problem since we are seeking an equilibrium satisfying the system's objective. Instead, let Player 2 go to c ; this brings the resources vector to $(4,1)$.

At that point, Player 3 can go down. Once again this is the outcome of a Nash equilibrium, but this would not be solution. Instead, Player 3 could go right, and the run so obtained would satisfy the objective of the system and keep the resources in check. However this is not the outcome of a Nash equilibrium since Player 3 can deviate and increase his payoff by going down.

In fact, there is no solution to the synthesis problem.

It is unfortunately a negative result that we must report in Section 3. Deciding the existence of a strategic equilibrium for careful autonomous agents in environments with multiple common resources is indeed undecidable.

We then propose in Section 4 a variant with bounded resources. In this setting, every resource has a maximum capacity.




Figure 1: A 2-resource 3-player game.

Example 2. Suppose now that both resources are bounded with bounds $(3, 3)$.

As before, Player 1 must pump thrice on a , with the resources values being successively $(2, 1)$, $(3, 2)$, and $(3, 3)$. As before, Player 2 could win in b by going down, and again this would be the outcome of a Nash equilibrium but would not be a solution. Instead, let Player 2 go to c , which brings the resources to $(1, 2)$.

To achieve his goal, in c , Player 3 must go down. To keep the first resource above zero, he must pump on c twice, thus bringing the value of the first resource to 3. But doing so he would deplete the second resource. If Player 3 instead carefully moves to the right, Player 1 and Player 2 meet their objectives, and so does the system. Hence this outcome results in a Nash equilibrium, that is, a solution.

To summarize, when the resources are bounded with bounds $(3, 3)$, the strategies of Player 1 taking the self-loop thrice, then going to b , Player 2 going to c , and Player 3 going to (\circ, \square) , form a Nash equilibrium which is a solution to the synthesis problem.

This variant with bounded resource storage capacity is of interest for the practical engineering of autonomous multiagent systems for two reasons. The first reason is conceptual. In many real-case scenarios, resources are bounded: e.g., in a community, a shared tank of water can only contain a predetermined amount of water, a shared microgrid powerpack can only contain a predetermined amount of energy, etc. The second reason is algorithmic. We will show that unlike in the setting with unbounded resources, the problem of rational synthesis in this bounded setting becomes decidable. Even better, with objectives expressed in LTL, it is not harder than the plain reactive synthesis problem, which is 2EXPTIME-complete [17].

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 $q \in Q^*$, $\text{lst}(q)$ is the last element in the sequence q .

2.1 Arenas, strategies and profiles

Multi-player arenas A multi-player arena is a tuple $\mathcal{G} = \langle S, (S_1 \uplus \dots \uplus S_n), s_0, P, E, AP, \ell \rangle$, where S is a finite set of states, $(S_1 \uplus \dots \uplus S_n)$ is a partition of S , s_0 is an initial state, $P = \{1, \dots, n\}$ is the set of players, E is an edge relation in $S \times S$, AP is the set of labels (atomic propositions), and $\ell: S \rightarrow 2^{AP}$ is the labeling function. For every edge $e = (s, t)$, $\text{Src}(e)$ is s and $\text{Trgt}(e)$ is t .

Plays and strategies For an arena \mathcal{G} , we denote the plays of this game by $\text{Plys}(\mathcal{G})$ that is the set of elements $s_0s_1s_2\dots$ in S^ω such that for all $i \geq 0$, (s_i, s_{i+1}) is in E . The set $\text{Hst}(\mathcal{G})$ is the set of prefixes of elements in $\text{Plys}(\mathcal{G})$. Moreover $\text{Hst}_i(\mathcal{G})$ for i in P is the set of elements in $\text{Hst}(\mathcal{G})$ whose last element is in S_i :

$$\text{Hst}_i(\mathcal{G}) = \{h \in \text{Hst}(\mathcal{G}) \mid \text{lst}(h) \in S_i\}$$

A *strategy* for player i is a function

$$\sigma_i: \text{Hst}_i(\mathcal{G}) \rightarrow S$$

mapping a history whose last element is s to a state s' such that $(s, s') \in E$. 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.,

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

Profile of strategies Once a strategy σ_i for each player i is chosen, we obtain a strategy profile $\bar{\sigma} = \langle \sigma_1, \dots, \sigma_n \rangle$. Note that a strategy profile has a unique play in its outcome. $\bar{\sigma}_{-i}$ is the corresponding partial profile without the strategy for player i . For a strategy σ'_i for a player i , we write $\langle \bar{\sigma}_{-i}, \sigma'_i \rangle$ the profile $\langle \sigma_1, \dots, \sigma'_i, \dots, \sigma_n \rangle$. We denote by $\langle \bar{\sigma} \rangle$ the unique outcome of the strategy profile $\bar{\sigma}$.

2.2 Objectives and payoffs

An objective Obj is a subset of $\text{Plys}(\mathcal{G})$. We write Obj_i to specify that it is the objective of player i . We define the payoff $\text{Payoff}_i(\bar{\sigma})$ of player i wrt. the profile $\bar{\sigma}$ as follows:

$$\text{Payoff}_i(\bar{\sigma}) = \begin{cases} 1 & \text{if } \langle \bar{\sigma} \rangle \in \text{Obj}_i \\ 0 & \text{otherwise} \end{cases}$$

LTL objective We describe specifications using the Linear-time Temporal Logic (LTL). An LTL specification is a formula ϕ defined using the following grammar:

$$\phi ::= \alpha \mid \neg\phi \mid \phi \vee \psi \mid X\phi \mid \phi \text{ U } \psi .$$

where α is in AP. As usual we denote with \diamond the ‘‘finally’’ operator, defined as $\diamond\phi = \text{true} \text{ U } \phi$.

LTL formulas are evaluated over plays as follows:

$$\begin{aligned} \rho \models \alpha & \text{ iff } \alpha \in \ell(\rho[0]) & \rho \models \neg\phi & \text{ iff } \rho \not\models \phi & \rho \models \phi \vee \psi & \text{ iff } \rho \models \phi \text{ or } \rho \models \psi, \\ \rho \models X\phi & \text{ iff } \rho[1..] \models \phi & \rho \models \phi \text{ U } \psi & \text{ iff } \exists i \geq 0, \rho[i..] \models \psi \text{ and } \forall 0 \leq j < i, \rho[j..] \models \phi, \end{aligned}$$

where $\rho \in S^\omega$, $\alpha \in \text{AP}$, $\phi \in \text{LTL}$, and $\psi \in \text{LTL}$.

For an LTL formula ϕ , we define the set $\langle \phi \rangle$ as the set of plays satisfying ϕ , i.e.,

$$\langle \phi \rangle = \{\rho \in S^\omega \mid \rho \models \phi\}$$

When the objectives are described as LTL formulas, a play ρ satisfies the objective of player i if $\rho \in \langle \text{Obj}_i \rangle$. Similarly we will sometimes write Obj_i to denote the set $\langle \text{Obj}_i \rangle$.

In the sequel, we will use the temporal modality $\diamond\phi$ as a shortcut for the formula $(\text{True} \text{ U } \phi)$.

Energy objectives Let $\text{cost}: E \rightarrow \mathbb{Z}$ be a cost function. To lighten the notation, we write $\text{cost}(s, t)$ instead of $\text{cost}((s, t))$. Let $h = s_0 s_1 \dots s_n$ be a history in $\text{Hst}(\mathcal{G})$; we abusively write $\text{cost}(h)$ to mean the extension of cost to histories that is

$$\text{cost}(h) = \sum_{i=0}^{n-1} \text{cost}(s_i, s_{i+1})$$

The energy objective for a game \mathcal{G} equipped with a cost function cost is given by the set Energy described as follows:

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

Throughout the paper, values of cost are encoded in binary.

Multi-energy objectives Let $\overline{\text{cost}}: E \rightarrow \mathbb{Z}^d$ be a multi dimensional cost function. We extend $\overline{\text{cost}}$ as expected to histories.

The multi-energy objective for a game \mathcal{G} equipped with a multi dimensional cost function $\overline{\text{cost}}$ is given by the set MultiEnergy described as follows:

$$\text{MultiEnergy}(\mathcal{G}) = \{\pi \in \text{Plys}(\mathcal{G}) \mid \forall i \geq 0, \overline{\text{cost}}(\pi[..i]) \geq (0, \dots, 0)\}$$

We will denote by cost_i the function obtained by projecting $\overline{\text{cost}}$ over the i -th dimension.

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. The set of all the Nash equilibria in a game is denoted NE.

Nash equilibria For a multi-player game $\mathcal{G} = \langle S, (S_1 \uplus \dots \uplus S_n), s_0, P, E, AP, \ell \rangle$ with objectives $\text{Obj}_1, \dots, \text{Obj}_n$ for each player, a profile $\bar{\sigma} = \langle \sigma_1, \dots, \sigma_n \rangle$ is a *Nash equilibrium* (NE) if for every player i and every strategy σ'_i for i the following holds true:

$$\text{Payoff}_i(\bar{\sigma}) \geq \text{Payoff}_i(\langle \bar{\sigma}_{-i}, \sigma'_i \rangle)$$

Equivalently for each player i and for each strategy σ'_i , if $\langle \bar{\sigma}_{-i}, \sigma'_i \rangle \in \text{Obj}_i$ then $\langle \bar{\sigma} \rangle \in \text{Obj}_i$.

2.4 Rational synthesis in the commons

Careful cooperative rational synthesis Let $\mathcal{G} = \langle S, (S_1 \uplus \dots \uplus S_n), s_0, P, E, AP, \ell \rangle$ be a game, $\overline{\text{cost}}: E \rightarrow \mathbb{Z}^d$ be a multi dimensional cost function, objectives $\text{Obj}_1, \dots, \text{Obj}_n$, Obj a global specification and let $\bar{\sigma}$ be a strategy profile. Then $\bar{\sigma}$ is a solution to the careful cooperative rational synthesis problem if:

$$\begin{aligned} &\langle \bar{\sigma} \rangle \in \text{MultiEnergy}(\mathcal{G}) \cap \text{Obj}, \text{ and } \forall \sigma'_i \text{ a strategy for player } i, \\ &\langle \bar{\sigma}_{-i}, \sigma'_i \rangle \in \text{Obj}_i \cap \text{MultiEnergy}(\mathcal{G}) \implies \langle \bar{\sigma} \rangle \in \text{Obj}_i \end{aligned}$$

3 Undecidability

We present multi-counter automata and the problem of reachability which is undecidable. We reduce it to the problem of careful cooperative rational synthesis.

3.1 Multi-counter automata

A *n-counter automaton* Γ is a tuple (L, δ, 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 $(l, \vec{w}, \vec{g}, l')$ where l and l' are locations in L , $\vec{w} \in \mathbb{Z}^n$ represents the weights of the transition, and $\vec{g} \in (\mathbb{N} \times (\mathbb{N} \cup \{\omega\}))^n$ represents the guards of the transitions. Given a transition δ , we note $g_i[l_0]$ the lower-bound for counter i and $g_i[up]$ the upper-bound.

A finite run in a *n-counter automaton* is a triple (k, μ_1, μ_2) , where $\mu_1 : \{0, \dots, k\} \rightarrow L$ and $\mu_2 : \{0, \dots, k\} \rightarrow \mathbb{Z}^n$, and such that:

- $\mu_1(0) = l_0$ and $\mu_2(0) = (0, \dots, 0)$
- for every $i < k$, if $\mu_1(i) = l$ and $\mu_2(i) = (c_1, \dots, c_n)$, and $\mu_1(i+1) = l'$ and $\mu_2(i+1) = (c'_1, \dots, c'_n)$, then there is $(l, \vec{w}, \vec{g}, l') \in \delta$, such that for all $0 \leq i \leq n$, we have $g_i[l_0] \leq c_i \leq g_i[up]$ and $c'_i = c_i + w_i$.

The *reachability problem* (Γ, t) in *n-counter automata* asks, given a *n-counter automaton* $\Gamma = (L, \delta, l_0)$ and a location $t \in L$, whether there is a finite run (k, μ_1, μ_2) such that $\mu_1(k) = t$, and $\mu_2(k) = (0, \dots, 0)$.

The following lemma can be easily proved using a reduction from 2-counter machines [13].

Lemma 3. *The reachability problem in 2-counter automata is undecidable.*

3.2 Undecidability of multi-resources careful cooperative rational synthesis

We reduce the reachability problem in 2-counter automata into the problem of careful cooperative rational synthesis with two resources and two players.

From a 2-counter automaton $\Gamma = (L, \delta, l_0)$ and a target location t , we are going to build a game $\mathcal{G}_{\Gamma, t} = \langle S_{\Gamma, t}, (S_1 \uplus S_2), s_{\Gamma, t}, \{1, 2\}, E_{\Gamma, t}, AP_{\Gamma, t}, \ell_{\Gamma, t} \rangle$ with costs $\overline{\text{cost}}_{\Gamma, t}$ and objectives $\text{Obj}_1, \text{Obj}_2$, $\text{Obj} = \text{Obj}_1$, in such a way that a solution to the reachability problem exists iff a solution to the careful cooperative rational synthesis exists.

Construction We are going to use two players in this construction. Player 1's role will be to build a solution, choosing the transitions to follow. Player 2's role will be to "check" that the transitions are legitimate, making Player 1 fail in his tasks if a transition that does not respect the guards is taken.

In $\mathcal{G}_{\Gamma, t}$, we first add the three states: W_1 representing the winning state of Player 1, and W_2 , and W'_2 representing the winning states of Player 2. They will be sink states, and it does not matter who controls them.

Each location in Γ is also a state in $\mathcal{G}_{\Gamma, t}$, controlled by Player 1.

For each transition $\tau = (l, \vec{w}, \vec{g}, l')$ in δ we introduce two states $\tau_<$ and $\tau_>$, both controlled by Player 2, and a few transitions. Intuitively, the state $\tau_>$ will serve as a state in which Player 2 will "check" that the upper-guard is satisfied. (Player 2 will have the opportunity to win if it does not.) The state $\tau_<$ will serve for the system to "check" that the lower guard is satisfied. (The value of one of the resources will go below zero if it is not the case.) There are four cases to consider; They are illustrated on Figure 2: (2a) $g_1[up] \neq \omega$ and $g_2[up] \neq \omega$; (2b) $g_1[up] \neq \omega$ and $g_2[up] = \omega$; (2c) $g_1[up] = \omega$ and $g_2[up] \neq \omega$; (2d) $g_1[up] = \omega$ and $g_2[up] = \omega$.

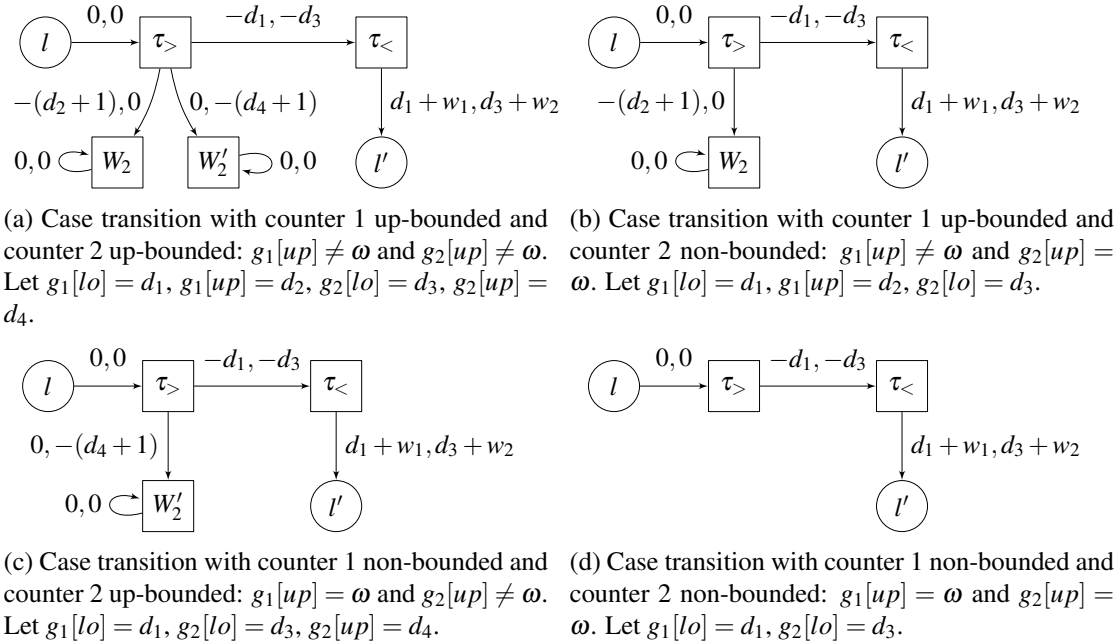



Figure 2: Gadgets to encode the transitions.


 Figure 3: Gadget to encode the target location t with counter values $(0,0)$.

We also need a gadget to “check” that solutions are runs that reach the target location with the values of the counters being zero. We introduce a state $t_?$ in $\mathcal{G}_{\Gamma,t}$, controlled by Player 2. See Figure 3.

The five gadgets also completely specify the cost function $\overline{\text{cost}}_{\Gamma,t}$.

The game $\mathcal{G}_{\Gamma,t}$ does not contain any other state or transition.

The initial state is $s_{\Gamma,t}$. The set of propositions $\text{AP}_{\Gamma,t}$ is $S_{\Gamma,t}$, and the labeling function $\ell_{\Gamma,t}$ is the identity.

In the game so obtained, the objective of Player 1 is to reach the state W_1 and the objective of Player 2 is to reach state W_2 or W_2' .

We now prove that the construction above can serve as a reduction from the reachability problem in 2-counter automata into the problem of careful rational synthesis (with two common resources).

Proposition 4. *Let $\Gamma = (L, \delta, l_0)$ be a 2-counter automaton and let t be a location in L . The reachability problem (Γ, t) has a positive answer iff there is a solution to the careful cooperative rational synthesis in the game $\mathcal{G}_{\Gamma,t} = \langle S_{\Gamma,t}, (S_1 \uplus S_2), s_{\Gamma,t}, \{1, 2\}, E_{\Gamma,t}, \text{AP}_{\Gamma,t}, \ell_{\Gamma,t} \rangle$.*

Proof. Left to right. Suppose there is a solution to the reachability problem. There is a run (k, μ_1, μ_2) such that $\mu_1(k) = t$, and $\mu_2(k) = (0,0)$. For every i , let $\tau^i = (l^i, (w_1^i, w_2^i), (g_1^i, g_2^i), l'^i)$ be the transition

between $\mu_1(i)$ and $\mu_1(i+1)$. By definition of the reachability problem, we know that the guards of all of them are satisfied. Since it reaches t at step k , the following sequence is a run in $\mathcal{G}_{\Gamma,t}$: $\pi_\rho := \mu_1(0) \cdot \tau_{>}^0 \cdot \tau_{<}^0 \cdot \mu_1(1) \cdot \tau_{>}^1 \cdot \tau_{<}^1 \cdots \mu_1(k) \cdot t_\gamma \cdot (W_1)^\omega$. We argue that π_ρ satisfies the objective, it never depletes the common resources, it is the play of a Nash equilibrium.

- Since the play π_ρ enters W_1 , it is in Obj.
- Since (k, μ_1, μ_2) is a solution to the reachability problem, by construction of $\mathcal{G}_{\Gamma,t}$, the values of the counters along the play π_ρ never go below 0.
- Along the play π_ρ , in every $\tau_{<}^i$, Player 2 chooses to go $\tau_{>}^i$. Since the guards are respected along the run (k, μ_1, μ_2) in Γ , by construction of $\mathcal{G}_{\Gamma,t}$, Player 2 never has an opportunity in any state $\tau_{>}^i$ to deviate carefully (and profitably) to W_2 or W_2' .

In state t_γ , Player 2 chooses to go to W_1 . Since $\mu_2(k) = (0, 0)$, the play π_ρ enters the state t_γ with both counters being 0. Thus, Player 2 cannot deviate carefully (and profitably) to W_2 or W_2' .

Since the play π_ρ enters W_1 , it is winning for Player 1, who has no incentive to deviate.

Hence, π_ρ is the outcome of a Nash equilibrium.

So there is a solution to the careful cooperative rational synthesis.

Right to left. Suppose there is a solution to the problem of careful cooperative rational synthesis.

By definition of the problem of careful cooperative rational synthesis, there is a strategy profile $\bar{\sigma}$ such that:

1. The profile $\bar{\sigma}$ is a Nash equilibrium.
2. The value of each resource never goes below 0.
3. The play $\langle \bar{\sigma} \rangle$ reaches the state W_1 .

We argue that the play $\langle \bar{\sigma} \rangle$ of $\mathcal{G}_{\Gamma,t}$ reaches t with the value of the resources being $(0, 0)$.

4. Since $\langle \bar{\sigma} \rangle$ reaches the state W_1 , by construction of the game, $\langle \bar{\sigma} \rangle$ is losing for Player 2. By construction also, $\langle \bar{\sigma} \rangle$ goes through the state t_γ . Since $\bar{\sigma}$ is a Nash equilibrium, it must be that the value of the resources when it goes through the state t_γ are $(0, 0)$, otherwise, Player 2 would profitably choose to go to W_2 or W_2' instead of W_1 . Suppose $\langle \bar{\sigma} \rangle$ reaches t at index k_t . I.e., $\langle \bar{\sigma} \rangle[k_t] = t$, and $\overline{\text{cost}}(\langle \bar{\sigma} \rangle[..k_t]) = (0, 0)$.

Let $\rho_{\bar{\sigma}} = (k_t, \mu_1, \mu_2)$ be the finite run in Γ that is the projection of $\langle \bar{\sigma} \rangle[..k_t]$ onto L . We argue that $\rho_{\bar{\sigma}}$ is a solution to the reachability problem (Γ, t) .

- From item 4, $\langle \bar{\sigma} \rangle$ reaches t at index k_t with resource values $(0, 0)$. So $\rho_{\bar{\sigma}}$ reaches the state $(t, (0, 0))$.
- We argue that the upper guards are always respected along $\rho_{\bar{\sigma}}$. By item 3 and by construction, we know that $\langle \bar{\sigma} \rangle$ is not winning for Player 2. But since $\bar{\sigma}$ is a Nash equilibrium (item 1), Player 2 never has an opportunity to carefully (and profitably) deviate to W_2 or W_2' . So when going through a transition $\tau = (l, \vec{w}, \vec{g}, l')$, say at step j along $\rho_{\bar{\sigma}}$, if $\mu_2(j) = (c_1, c_2)$ then $c_1 \leq g_1[up]$ and $c_2 \leq g_2[up]$.
- By construction, it follows from item 2 that the lower guards are always respected along $\rho_{\bar{\sigma}}$.

□

The next result follows at once.

Theorem 5. *The problem of careful cooperative rational synthesis is undecidable, even with two players and two resources, and reachability objectives.*

4 Bounded resources and decidability

In this section we consider the so-called bounded setting. Here each counter will be bounded from above by some bound B and cannot store more than B . Intuitively, one can continue to charge a battery, but the energy exceeding the capacity will be lost in heat; just like one can continue to fill up a tank of water but it will spill over when the capacity is reached.

The main result of this section will be the decidability of the synthesis problem in this case.

Given a game $\mathcal{G} = \langle S, (S_1 \uplus \dots \uplus S_n), s_0, P, E, AP, \ell \rangle$, and multidimensional cost function $\overline{\text{cost}}: E \rightarrow \mathbb{Z}^d$, we fix a vector in \mathbb{N}^d representing the maximal capacity of each counter. We use $\overline{\text{cost}}(s, s')[i]$ to denote the i -th component of the tuple $\overline{\text{cost}}(s, s')$. Along a run when a counter is at capacity its value cannot increase. Formally, Assume that the capacity is given by the following vector $\vec{B} = (B_1, \dots, B_d)$. We define the operator $\oplus_{\vec{B}}$ over vectors in \mathbb{Z}^d as follows:

$$(c_1, \dots, c_d) \oplus_{\vec{B}} (c'_1, \dots, c'_d) = (x_1, \dots, x_d) \text{ where } \forall 1 \leq i \leq d, x_i = \min(c_i + \overline{\text{cost}}(s, s')[i], B_i)$$

We can now define the cost vector along a history $h = s_0 s_1 \dots s_n$ inductively as follows

$$\overline{\text{cost}}(h) = \overline{\text{cost}}(s_0 s_1 \dots s_{n-1}) \oplus_{\vec{B}} \overline{\text{cost}}(s_n, s_{n+1})$$

The decidability result is obtained through an unfolding of the arena. This unfolding constructs a multiplayer game without costs $\widehat{\mathcal{G}}$ where the set of states is $\widehat{S} = S \times \{0, \dots, B_1\} \times \dots \times \{0, \dots, B_d\} \cup \{\perp\}$. The set of edges is \widehat{E} in $(\widehat{S} \times \widehat{S}) \cup (\widehat{S} \times \{\perp\}) \cup \{(\perp, \perp)\}$ and is defined as follows:

$$((s, c_1, \dots, c_d), (s', c'_1, \dots, c'_d)) \in \widehat{E}$$

if

- (s, s') in E
- $(c_1, \dots, c_d) \oplus_{\vec{B}} \overline{\text{cost}}(s, s') = (c'_1, \dots, c'_d)$
- $(c'_1, \dots, c'_d) \geq (0, \dots, 0)$

and

$$((s, c_1, \dots, c_d), \perp) \in \widehat{E}$$

if

- there exists s such that (s, s') in E
- for some counter value c_i we have $c_i + \text{cost}_i(s, s') < 0$.

Also $(\perp, \perp) \in \widehat{E}$.

In this new game a state (s, c_1, \dots, c_d) belongs to player i if s belongs to player i . Player 1 controls also the fresh state \perp . The objective of each player is the same and the global specification in $\widehat{\mathcal{G}}$ is $\text{Obj} \wedge \neg \diamond \perp$. Plays in this unfolding are infinite sequences of \widehat{S} . In order to relate plays in \mathcal{G} with plays in $\widehat{\mathcal{G}}$ we use the following projection π defined over the set of histories as follows: first,

$$\pi(s_0) = (s_0, 0, \dots, 0)$$

and for a history $h = s_0 s_1 \dots s_l$ in \mathcal{G} :

$$\pi(s_0 s_1 \dots s_l) = \pi(s_0 s_1 \dots s_{l-1})(s, \overline{\text{cost}}(s_0 s_1 \dots s_l))$$

We extend π over the plays as expected and denote by π^{-1} the inverse mapping.

For a play $\hat{\rho}$, we say that $\hat{\rho}$ satisfies the objective of player i if $\pi^{-1}(\hat{\rho})$ satisfies Obj_i . We will say that a play $\hat{\rho}$ satisfies $\text{Obj} \wedge \neg \diamond \perp$ if $\pi(\hat{\rho})$ satisfies Obj and $\hat{\rho}$ satisfies $\neg \diamond \perp$.

Proposition 6. *There exists a solution to the careful synthesis if and only if there exists a Nash equilibrium in the unfolding whose outcome satisfies $\text{Obj} \wedge \neg \diamond \perp$.*

Proof. Let $\bar{\sigma}$ be a solution in \mathcal{G} , we construct $\widehat{\sigma}$ as follows:

$$\widehat{\sigma} = (\widehat{\sigma}_1, \dots, \widehat{\sigma}_n)$$

such that each $\widehat{\sigma}_i$ is defined as follows:

$$\widehat{\sigma}_i(\hat{h}) = \sigma_i(\pi(\hat{h}))$$

where \hat{h} is a history of $\widehat{\mathcal{G}}$, and σ_i is the strategy of player i in the profile $\bar{\sigma}$. We argue that $\widehat{\sigma}$ is also a solution thanks to the following fact:

- $\widehat{\sigma}$ is a Nash equilibrium since each $\widehat{\sigma}_i$ ensures the same payoff as σ_i .
- $\bar{\sigma}$ is solution, hence it ensures that the energy along its outcome never drops below 0 for all the counters, hence by construction \perp is never visited.

Let $\widehat{\sigma}$ be a solution of $\widehat{\mathcal{G}}$, then we construct $\bar{\sigma}$ as follows:

$$\bar{\sigma} = (\sigma_1, \dots, \sigma_n)$$

where for each history h ,

$$\sigma_i = \widehat{\sigma}_i(\pi(h))$$

We argue that $\bar{\sigma}$ is also a solution thanks to the following fact:

- $\bar{\sigma}$ is a Nash equilibrium since each σ_i ensures the same payoff as $\widehat{\sigma}_i$.
- $\widehat{\sigma}$ is solution, hence it ensures that \perp is never visited, therefore by construction the energy along the outcome of $\bar{\sigma}$ never drops below 0 for all the counters,

□

By Proposition 6, we know that we can solve careful synthesis in the original arena by reducing it to plain rational synthesis in the unfolding. It is readily seen that the size of the unfolding of the arena defined above is exponential in the size of the original arena. On the other hand, solving the cooperative rational synthesis with LTL objectives is in 2EXPTIME in the size of the objectives formulas, and polynomial in the size of the arena [11, 12]. It follows that our problem is 2EXPTIME when the counters are bounded.

Theorem 7. *The careful cooperative rational synthesis is 2EXPTIME-complete when the counters are bounded.*

5 Conclusion

As agents and robots are always more likely to roam the physical world, the formal tools to engineer them need to take into account the resource-sensitiveness of their activities.

We presented a model for autonomous and rational agents interacting in environments with multiple resources. We focused on a problem of rational planning, rational synthesis, that consists in finding a

non-cooperative equilibrium (a Nash equilibrium) that satisfies a system objective, and never depletes the resources. We showed that this problem is undecidable.

We then proposed a variant where the storage capacity is bounded for all resources. We claim that this is promising for the applicability to real-world settings. The storage of resources is indeed generally limited (energy power capacity of a battery, volume of a water tank, etc). Moreover, we proved that the problem of rational synthesis with LTL objectives becomes decidable in double-exponential time, which is no harder than plain controller synthesis for LTL specifications.

In the future, we are interested in the study of problems to elaborate tools that better equip the engineers of agents and robots in resource-sensitive environments. In particular, we will investigate problems of *parameterized* synthesis allowing an engineer to partially model a system, leaving some quantities unspecified (for example, we can leave unspecified some weights of transitions or the bounds of the resources), and with the aim of automatically completing the system in a way that it admits a solution to the synthesis problem.

References

- [1] Alessandro Abate, Julian Gutierrez, Lewis Hammond, Paul Harrenstein, Marta Kwiatkowska, Muhammad Najib, Giuseppe Perelli, Thomas Steeples & Michael J. Wooldridge (2021): *Rational verification: game-theoretic verification of multi-agent systems*. *Appl. Intell.* 51(9), pp. 6569–6584, doi:10.1007/s10489-021-02658-y.
- [2] Natasha Alechina & Brian Logan (2020): *State of the Art in Logics for Verification of Resource-Bounded Multi-Agent Systems*. In Andreas Blass, Patrick Cégielski, Nachum Dershowitz, Manfred Droste & Bernd Finkbeiner, editors: *Fields of Logic and Computation III - Essays Dedicated to Yuri Gurevich on the Occasion of His 80th Birthday, Lecture Notes in Computer Science* 12180, Springer, pp. 9–29, doi:10.1007/978-3-030-48006-6_2.
- [3] Fahiem Bacchus & Froduald Kabanza (2000): *Using temporal logics to express search control knowledge for planning*. *Artificial Intelligence* 116(1), pp. 123–191, doi:10.1016/S0004-3702(99)00071-5.
- [4] Meghyn Bienvenu, Christian Fritz & Sheila A. McIlraith (2006): *Planning with Qualitative Temporal Preferences*. In Patrick Doherty, John Mylopoulos & Christopher A. Welty, editors: *Proceedings, Tenth International Conference on Principles of Knowledge Representation and Reasoning, Lake District of the United Kingdom, June 2-5, 2006*, AAAI Press, pp. 134–144.
- [5] Nils Bulling & Valentin Goranko (2013): *How to Be Both Rich and Happy: Combining Quantitative and Qualitative Strategic Reasoning about Multi-Player Games (Extended Abstract)*. In Fabio Mogavero, Aniello Murano & Moshe Y. Vardi, editors: *Proceedings 1st International Workshop on Strategic Reasoning, SR 2013, Rome, Italy, March 16-17, 2013, EPTCS* 112, pp. 33–41, doi:10.4204/EPTCS.112.8.
- [6] Alessandro Cimatti, Marco Pistore & Paolo Traverso (2008): *Chapter 22 Automated Planning*. In Frank van Harmelen, Vladimir Lifschitz & Bruce Porter, editors: *Handbook of Knowledge Representation, Foundations of Artificial Intelligence* 3, Elsevier, pp. 841–867, doi:10.1016/S1574-6526(07)03022-2.
- [7] Rodica Condurache, Catalin Dima, Youssouf Oualhadj & Nicolas Troquard (2021): *Rational Synthesis in the Commons with Careless and Careful Agents*. In Frank Dignum, Alessio Lomuscio, Ulle Endriss & Ann Nowé, editors: *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, ACM, pp. 368–376, doi:10.5555/3463952.3464000.
- [8] Rodica Condurache, Emmanuel Filiot, Raffaella Gentilini & 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*, pp. 121:1–121:15, doi:10.4230/LIPIcs.ICALP.2016.121.
- [9] Patrick Doherty & Jonas Kvarnstram (2001): *TALplanner: A Temporal Logic-Based Planner*. *AI Magazine* 22(3), pp. 95–102, doi:10.1609/aimag.v22i3.1581.

- [10] Berndt Farwer (2002): *ω -Automata*. In Erich Grädel, Wolfgang Thomas & Thomas Wilke, editors: *Automata Logics, and Infinite Games: A Guide to Current Research*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 3–21, doi:10.1007/3-540-36387-4_1.
- [11] Dana Fisman, Orna Kupferman & 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*, pp. 190–204, doi:10.1007/978-3-642-12002-2_16.
- [12] Orna Kupferman & Noam Shenwald (2022): *The Complexity of LTL Rational Synthesis*. In Dana Fisman & Grigore Rosu, editors: *Tools and Algorithms for the Construction and Analysis of Systems*, Springer International Publishing, Cham, pp. 25–45, doi:10.1007/978-3-030-99524-9_2.
- [13] Marvin L. Minsky (1961): *Recursive Unsolvability of Post's Problem of "Tag" and other Topics in Theory of Turing Machines*. *Annals of Mathematics* 74(3), pp. 437–455, doi:10.2307/1970290.
- [14] John F. Nash (1951): *Non-Cooperative Games*. *The Annals of Mathematics* 54, pp. 286–295, doi:10.2307/1969529.
- [15] Hoang Nga Nguyen, Natasha Alechina, Brian Logan & Abdur Rakib (2018): *Alternating-time temporal logic with resource bounds*. *J. Log. Comput.* 28(4), pp. 631–663, doi:10.1093/logcom/exv034.
- [16] Amir Pnueli (1977): *The Temporal Logic of Programs*. In: *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, IEEE Computer Society, pp. 46–57, doi:10.1109/SFCS.1977.32.
- [17] Amir Pnueli & 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*, pp. 179–190, doi:10.1145/75277.75293.
- [18] Michael Ummels (2008): *The Complexity of Nash Equilibria in Infinite Multiplayer Games*. In Roberto Amadio, editor: *Foundations of Software Science and Computational Structures*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 20–34, doi:10.1007/978-3-540-78499-9_3.