

Robust Alternating-Time Temporal Logic

Aniello Murano¹[0000–0003–4876–3448], Daniel Neider^{2,3}[0000–0001–9276–6342], and
Martin Zimmermann⁴[0000–0002–8038–2453]

¹ Università degli Studi di Napoli “Federico II”, Naples, Italy

² TU Dortmund University, Germany

`daniel.neider@tu-dortmund.de`

³ Center for Trustworthy Data Science and Security, UA Ruhr, Germany

⁴ Aalborg University, Denmark

Abstract. In multi-agent system design, a crucial aspect is to ensure robustness, meaning that for a coalition of agents A , small violations of adversarial assumptions only lead to small violations of A 's goals. In this paper we introduce a logical framework for robust strategic reasoning about multi-agent systems. Specifically, inspired by recent works on robust temporal logics, we introduce and study rATL and rATL*, logics that extend the well-known Alternating-time Temporal Logic ATL and ATL* by means of an opportune multi-valued semantics for the strategy quantifiers and temporal operators. We study the model-checking and satisfiability problems for rATL and rATL* and show that dealing with robustness comes at no additional computational cost. Indeed, we show that these problems are PTIME-complete and EXPTIME-complete for rATL, respectively, while both are 2EXPTIME-complete for rATL*.

Keywords: Multi-Agents · Temporal Logic · Robustness.

1 Introduction

Multi-agent system verification has been receiving a lot of attention in recent years, thanks to the introduction of powerful logics for strategic reasoning [5,22,44,39,12]. Along this line of research, a story of success is *Alternating-Time Temporal Logic* (ATL) introduced by Alur, Henzinger, and Kupferman [5]. ATL is a generalization of Computation Tree Logic (CTL) [23], obtained by replacing the path quantifier \exists (and its dual \forall), with the modality $\langle\langle A \rangle\rangle$ (and its dual $\llbracket A \rrbracket$), where A is a set of agents. The interpretation of $\langle\langle A \rangle\rangle\varphi$ is that the coalition A has a strategy such that the outcome of this strategy satisfies φ , no matter how the coalition of the agents not in A behaves. ATL formulas are interpreted over concurrent game structures, which extend classical Kripke structures to represent the dynamism of the agents. The model-checking problem of ATL is PTIME-complete [5], while the satisfiability problem is EXPTIME-complete [54].

A crucial aspect in multi-agent system design is to ensure system *robustness*, which should reflect the ability of a coalition of agents to tolerate violations (possibly up to some extent) of adversarial assumptions [21]. Numerous studies have shown that reactive AI systems can be very sensitive to intentional or

unintentional external perturbations, posing huge risks to safety-critical applications [36]. Notably, the formal methods community has put large efforts in reasoning about system robustness in several specific settings, mainly concerning closed system verification or (two-player) reactive synthesis [19,31,17,25,27,28]. As far as we are aware of, there are no logic-based works dealing with robust strategic reasoning in multi-agent systems. To highlight the significance of this challenge, we describe a few specific contexts in which multi-agent systems act as the natural model and robustness plays a crucial role.

Scenario 1. Climate change threatens people with food and water scarcity, increased flooding, extreme heat, diseases, and economic loss. Human migration and conflict can be a result. The World Health Organization calls climate change the greatest threat to global health in the 21st century. Recently, researchers examining alternative policies to address the threat of climate change have become increasingly concerned about uncertainty and the fact that we cannot predict the future. This requires to develop mathematical models to properly represent the intricate interaction among all decision makers and the ability to define strategies that are robust against a wide range of plausible climate-change futures [40]. For risk-averse policy-makers, such strategies would perform reasonably well, at least compared to the alternatives, even if confronted with surprises or catastrophes. Robust strategies may also provide a more solid basis for consensus on political action among stakeholders with different views of the future, because it would provide reasonable outcomes no matter whose view proved correct.

Scenario 2. The fast-evolving domain of *autonomous vehicles* is one of the best examples of multi-agent modelling, where safety-critical decisions strongly rely on sensor observations (e.g., ultrasound, radar, GPS, Lidar, and camera signals) [53]. It is of primary importance that the resulting decisions are robust to perturbations, which often are treated as adversarial perturbations [43]. A careful evaluation of such adversarial behaviours is necessary to build and deploy safer autonomous vehicle systems.

Scenario 3. Power systems play an important role in all sectors of the national economy and in our daily lives. Ensuring a safe and reliable power supply from the power network is a fundamental requirement. As renewable energy-based smart grid and micro-grid systems rise in popularity, multi-agent system technology has been establishing itself as a useful paradigm of choice for modelling, analysis, control and optimization of power systems [33,48,51]. The model usually consists of several agents competing not only among themselves to get energy resources, but also playing against the unpredictable behaviour of nature. Then, a classical safety requirement amounts to ensuring system robustness, in the meaning that the power system has to keep operating, possibly by rationing resources, despite the loss of any single asset such as lines or power plants at any time [1,16]. This is usually enforced by following a simple guiding redundancy principle while developing the system: designers have to predict the effect of having any line disconnected in any moment and cope with it, in real time and even at larger scales [46]. This may also require the players to coordinate and/or play rational while keeping the system under equilibrium [14].

Our Contribution In this paper we introduce rATL, a robust version of the logic ATL. Our approach here follows and extends an approach originally introduced for robust Linear Temporal Logic (rLTL) [52] and later extended to robust Computation Tree Logic (rCTL and rCTL*) [45]. To illustrate the robust semantics, consider an invariant of the form $\Box p$ specifying that the proposition p always holds. There are several ways this invariant can be violated, with varying degrees of severity. For example, p failing to hold a finite number of times is less severe than p failing infinitely often. An even worse situation is p holding only finitely often while p not even holding once is the worst way to violate the invariant. The authors in [52] argue that these five degrees are canonical and use them as the basis of a five-valued robust semantics for temporal logics. The semantics of the Boolean operators are then defined to capture the intuition that there are different degrees of violation of a formula while the other temporal operators, e.g., next and eventually, are defined as usual. In particular, the definition of implications captures the idea that, in a specification of the form $\varphi \rightarrow \psi$, a “small” violation of an environment assumption φ must lead to only a “small” (proportional) violation of a system’s guarantee ψ .

Here, we devise a meaningful robust semantics for the strategy quantifiers to obtain a robust variant of ATL, and show that it is capable to reason about the robustness of multi-agent systems. More precisely, rATL allows to assess whether a strategy f of a coalition A is robust in the sense that, with respect to the outcome of f , small violations of the adversarial team assumptions only lead to small violations of A ’s goals. We study expressiveness of rATL and show that it strictly subsumes ATL, as rATL can express fairness. We also study the model-checking and satisfiability problems for rATL and show that dealing with robustness comes at no additional computational cost. Indeed, we show that these problems are PTIME-complete and EXPTIME-complete, respectively. This is in line with the results on rLTL and rCTL, for which model-checking and satisfiability are also not harder than for LTL [52] and CTL [45], respectively.

Finally, we also study rATL*, the robustification of ATL*, showing that also in this setting, robustness comes for free: model-checking and satisfiability for rATL* are 2EXPTIME-complete, as they are already for ATL* [5,49].

All proofs omitted due to space restrictions can be found in the appendix.

Related work There are several works done in formal strategic reasoning that have been used (or can be easily adapted) to develop robust systems. Besides those reported above, we recall the works dealing with strategy logics extended with probabilistic [34,50,7] and knowledge (imperfect information) aspects [26]. These works allow to reason about the unpredictable behaviour of the environment. Unfortunately, in both cases, the model-checking problem becomes highly undecidable, unless one restricts strategies to be memoryless. In the imperfect information case, memoryfull strategies with less severe restrictions have been also studied (e.g., hierarchical visibility [15] and public action [13]) although model-checking remains infeasible, i.e., non-elementary, in practice.

Other lines of research have considered quantitative aspects of the logic, in different directions. Bouyer et al. [18] considered a fuzzy extension of ATL*,

namely $\text{ATL}^*[\mathcal{F}]$. The satisfaction value of $\text{ATL}^*[\mathcal{F}]$ formulas is a real value in $[0, 1]$, reflecting “how much” or “how well” the strategic on-going objectives of the underlying agents are satisfied. In [18] a double exponential-time model-checking procedure for $\text{ATL}^*[\mathcal{F}]$ is presented. A careful inspection of that procedure yields, for the special case of $\text{ATL}[\mathcal{F}]$, an EXPTIME-completeness result by means of an exponential reduction to Büchi games. Faella, Napoli, and Parente [30] and Aminof et al. [8] considered a graded extension of the logics ATL and ATL^* with the ability of checking for the existence of redundant winning strategies.

Module checking is another example of a formal method to devise robust systems. Indeed, module checking amounts to checking whether a strategic behaviour of a coalition of agents satisfies a goal, irrespective to all possible nondeterministic behaviours of an external environment [38,35].

Finally, robustness is also an active field of study in reinforcement learning [47], which treats environment mismatches as adversarial perturbations against a coalition of agents. In the simplest version, the underlying model is a two-player zero-sum simultaneous game between the protagonist who aims to find a robust strategy across environments and the adversary who exerts perturbations. Computational methods have been proposed to solve this game and to find a robust strategy for the protagonist (see Pinto et al. [47] and the references therein).

2 Preliminaries

We denote the nonnegative integers by \mathbb{N} , and the power set of a set S by 2^S . Throughout the paper, we fix a finite set AP of atomic propositions.

A concurrent game structure $\mathcal{S} = (St, Ag, Ac, \delta, \ell)$ consists of a finite set St of states, a finite set Ag of agents, a finite set Ac of actions, and a labeling function $\ell: St \rightarrow 2^{\text{AP}}$. An action vector for a subset $A \subseteq Ag$ is a mapping $v: A \rightarrow Ac$. Let AV denote the set of action vectors for the full set Ag of agents. The transition function $\delta: St \times AV \rightarrow St$ maps a state and an action vector to a state. The size of \mathcal{S} is defined as $|St \times AV|$.

We say that a state s' is a successor of a state s if there is an action vector $v \in AV$ such that $s' = \delta(s, v)$. A path of \mathcal{S} is an infinite sequence $\pi = s_0 s_1 s_2 \dots$ of states such that s_{n+1} is a successor of s_n for every $n \geq 0$. We write $\pi[n]$ for s_n .

A strategy for an agent is a function $f: St^+ \rightarrow Ac$. Given a set $F_A = \{f_a \mid a \in A\}$ of strategies, one for each agent in some set $A \subseteq Ag$, $out(s, F_A)$ denotes the set of paths starting in s that are consistent with F_A . Formally, a path $s_0 s_1 s_2 \dots$ is in $out(s, F_A)$ if $s_0 = s$ and for all $n \geq 0$, there is an action vector $v \in AV$ with $v(a) = f_a(s_0 \dots s_n)$ for all $a \in A$ and $s_{n+1} = \delta(s_n, v)$. Intuitively, $out(s, F_A)$ contains all paths that are obtained by the agents in A picking their actions according to their strategies and the other agents picking their actions arbitrarily.

3 rATL

The basic idea underlying our robust version of ATL, or *rATL* for short, is that a “small” violation of an environment assumption (along the outcome of

a strategy) must lead to only a “small” violation of a system’s guarantee. This is obtained by devising a robust semantics for the strategy quantifiers and by stating formally what it is meant for a “small” violations of a property. For the latter, we follow and adapt the approach by Tabuada and Neider [52], initially proposed for a robust version of Linear Temporal Logic (rLTL), and use five truth values: 1111, 0111, 0011, 0001, and 0000. Let \mathbb{B}_4 denote the set of these truth values. Our motivation for using the seemingly odd-looking truth values in \mathbb{B}_4 is that they represent five canonical ways how a system guarantee of the form “always p ” ($\Box p$ in LTL) can be satisfied or violated. Clearly, we prefer that p always holds, represented by the truth value 1111. However, if this is impossible, the following best situation is that p holds at least almost always, represented by 0111. Similarly, we would prefer p being satisfied at least infinitely often, represented by 0011, over p being satisfied at least once, represented by 0001. Finally, the worst situation is that p never holds, represented by 0000. Put slightly differently, the bits of each truth value represent (from left to right) the modalities “always” (\Box), “eventually always” ($\Diamond\Box$), “always eventually” ($\Box\Diamond$), and “eventually” (\Diamond). We refer the reader to Anevlavis et al. [11] for an in-depth explanation of why these five ways are canonical.

Following the intuition above, we order the truth values in \mathbb{B}_4 by

$$1111 \succ 0111 \succ 0011 \succ 0001 \succ 0000.$$

This order spans a spectrum of truth values ranging from 1111, corresponding to *true*, on one end, to 0000, corresponding to *false*, on the other end. Since we arrived at the set \mathbb{B}_4 by considering the canonical ways of how the invariant property $\Box p$ can fail, we interpret all truth values different from 1111 as *shades of false*. We return to this interpretation when we later define the semantics for the negation in rATL.

Having formally discussed how we “grade” the violation of a property along paths, we are now ready to define the syntax of rATL via the following grammar:

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \varphi \rightarrow \varphi \mid \langle\langle A \rangle\rangle\Phi \mid \llbracket A \rrbracket\Phi \\ \Phi &::= \odot\varphi \mid \Diamond\varphi \mid \Box\varphi \end{aligned}$$

where p ranges over atomic propositions and A ranges over subsets of agents. We distinguish between *state formulas* (those derivable from φ) and *path formulas* (those derivable from Φ). If not specified, an rATL formula is a state formula.

Various critical remarks should be made concerning the syntax of rATL. First, we add “dots” to temporal operators (following the notation by Tabuada and Neider [52]) to distinguish between the original operators in ATL and their robustified counterparts in rATL—otherwise, the syntax stays the same. Second, many operators of rATL, most notably the negation and implication, can no longer be derived via De Morgan’s law or simple logical equivalencies due to rATL’s many-valued nature. Hence, they need to be added explicitly. Third, we omit the until and release operators here to avoid cluttering our presentation too much. Both can be added easily, as in rLTL [10,9].

We define the semantics of rATL by an *evaluation function* V that maps a state formula and a state or a path formula and a path to a truth value in \mathbb{B}_4 . To simplify our presentation, we use $b[k]$ as a shorthand notation for addressing the k -th bit, $k \in \{1, 2, 3, 4\}$, of a truth value $b = b_1b_2b_3b_4 \in \mathbb{B}_4$ (i.e., $b[k] = b_k$). It is worth emphasizing that our semantics for rATL is a natural extension of the Boolean semantics of ATL and is deliberately designed to generalize the original Boolean semantics of ATL (see Subsection 3.3).

Turning to the definition of rATL's semantics, let us begin with state formulas. For atomic propositions $p \in \text{AP}$, we define the valuation function by

$$V(s, p) = \begin{cases} 1111 & \text{if } p \in \ell(s); \text{ and} \\ 0000 & \text{if } p \notin \ell(s). \end{cases}$$

Note that this definition mimics the semantics of ATL in that propositions get mapped to one of the two truth values *true* (1111) or *false* (0000). As a consequence, the notion of robustness in rATL does not arise from atomic propositions (e.g., as in LTL[\mathcal{F}] by Almagor, Boker, and Kupferman [4] or fuzzy logics) but from the evolution of the temporal operators (see the semantics of path formulas). This design choice is motivated by the observation that assigning meaningful (robustness) values to atomic propositions is often highly challenging in practice—if not impossible.

The semantics of conjunctions and disjunctions are defined as usual for many-valued logics in terms of the functions min and max:

$$\begin{aligned} V(s, \varphi_1 \vee \varphi_2) &= \max(V(s, \varphi_1), V(s, \varphi_2)) \\ V(s, \varphi_1 \wedge \varphi_2) &= \min(V(s, \varphi_1), V(s, \varphi_2)) \end{aligned}$$

To define the semantics of negation, remember our interpretation of the truth values in \mathbb{B}_4 : 1111 corresponds to *true* and all other truth values correspond to different shades of *false*. Consequently, we map 1111 to 0000 and all other truth values to 1111. This idea is formalized by

$$V(s, \neg\varphi) = \begin{cases} 0000 & \text{if } V(s, \varphi) = 1111; \text{ and} \\ 1111 & \text{if } V(s, \varphi) \prec 1111. \end{cases}$$

Note that the definition of $V(s, \neg\varphi)$ is not symmetric, which is in contrast to other many-valued logics, such as LTL[\mathcal{F}]. However, it degenerates to the standard Boolean negation if one considers only two truth values.

Since our negation is defined in a non-standard way, we cannot recover implication from negation and disjunction. Instead, we define the implication $a \rightarrow b$ by requiring that $c \prec a \rightarrow b$ if and only if $\min\{a, c\} \prec b$ for every $c \in \mathbb{B}_4$. This notion leads to

$$V(s, \varphi_1 \rightarrow \varphi_2) = \begin{cases} 1111 & \text{if } V(s, \varphi_1) \preceq V(s, \varphi_2); \text{ and} \\ V(s, \varphi_2) & \text{if } V(s, \varphi_1) \succ V(s, \varphi_2). \end{cases}$$

Again, this definition collapses to the usual Boolean definition in case one considers only two truth values.

We now provide the robust semantics for the strategy quantifiers, which are the key ingredient in rATL. First, notice that the strategy quantifiers $\langle\langle\cdot\rangle\rangle$ and $\llbracket\cdot\rrbracket$ are not dual in our robustified version of ATL and require their individual definitions. Intuitively, $\langle\langle A \rangle\rangle\Phi$ is the largest truth value that the coalition A of agents can enforce for the path formula Φ , while $\llbracket A \rrbracket\Phi$ is the largest truth value that $Ag \setminus A$ can enforce against A . Formally, we have the following:

- $V(s, \langle\langle A \rangle\rangle\Phi)$ is the maximal truth value $b \in \mathbb{B}_4$ such that there is a set F_A of strategies, one for each agent in A , such that for all paths $\pi \in out(s, F_A)$ we have $V(\pi, \Phi) \succeq b$.
- $V(s, \llbracket A \rrbracket\Phi)$ is the maximal truth value $b \in \mathbb{B}_4$ such that for all sets F_A of strategies, one for each agent in A , there exists a path $\pi \in out(s, F_A)$ with $V(\pi, \Phi) \succeq b$.

Let us now turn to the semantics of path formulas. We begin with the \square -operator. This operator captures the five canonical ways an invariant property “always p ” can be satisfied or violated, thereby implementing the intuition we have presented at the beginning of this section. Formally, the valuation function $V(\pi, \square\varphi)$ is given by $V(\pi, \square\varphi) = b_1b_2b_3b_4$ where

$$\begin{aligned} b_1 &= \min_{i \geq 0} V(\pi[i], \varphi)[1], & b_3 &= \min_{i \geq 0} \max_{j \geq i} V(\pi[j], \varphi)[3], \\ b_2 &= \max_{i \geq 0} \min_{j \geq i} V(\pi[j], \varphi)[2], & b_4 &= \max_{i \geq 0} V(\pi[i], \varphi)[4]. \end{aligned}$$

Note that for $p \in AP$ and a path π , the semantics of the formula $\square p$ on π amounts to the four-tuple $(\square p, \diamond\square p, \square\diamond p, \diamond p)$ because $V(s, p)$ is either 0000 or 1111 on every state s along π (i.e., all bits are either 0 or 1). However, the interpretation of $V(\pi, \square\varphi)$ becomes more involved once the formula φ is nested since the semantics of the \square -operator refers to individual bits of $V(\pi, \varphi)$.

Finally, the semantics for the \diamond -operator and \odot -operator are straightforward as there are only two possible outcomes: either the property is satisfied, or it is violated. Consequently, we define the valuation function by

- $V(\pi, \diamond\varphi) = b_1b_2b_3b_4$ with $b_k = \max_{i \geq 0} V(\pi[i], \varphi)[k]$; and
- $V(\pi, \odot\varphi) = b_1b_2b_3b_4$ with $b_k = V(\pi[1], \varphi)[k]$.

Again, note that both $V(\pi, \diamond\varphi)$ and $V(\pi, \odot\varphi)$ refer to individual bits of $V(\pi, \varphi)$.

Example 1. Consider the formula $\varphi = \langle\langle A \rangle\rangle\square p$. We have

- $V(s, \varphi) = 1111$ if the coalition A has a (joint) strategy to ensure that p holds at every position of every outcome.
- $V(s, \varphi) = 0111$ if the coalition A has strategy to ensure that p holds at all but finitely many positions of every outcome.
- $V(s, \varphi) = 0011$ if the coalition A has strategy to ensure that p holds at infinitely many positions of every outcome.
- $V(s, \varphi) = 0001$ if the coalition A has strategy to ensure that p holds at least once on every outcome.

3.1 rATL Model-Checking

The model-checking problem for rATL is as follows: Given a concurrent game structure \mathcal{S} , a state s , an rATL formula φ , and a truth value $t \in \mathbb{B}_4$, is $V(s, \varphi) \succeq t$?

Theorem 1. *rATL model-checking is PTIME-complete.*

The proof is based on capturing the semantics of the strategy quantifiers $\langle\langle A \rangle\rangle$ and $\llbracket A \rrbracket$ by sequential two-player games, one player representing the agents in A and the other representing the agents in the complement of A . We begin by introducing the necessary background on such games.

A (sequential) two-player game structure $\mathcal{S} = (St, St_1, St_2, Ac_1, Ac_2, \delta)$ consists of a set St of states partitioned into the states $St_p \subseteq St$ of Player $p \in \{1, 2\}$, an action set Ac_p for Player $p \in \{1, 2\}$, and a transition function $\delta: St_1 \times Ac_1 \cup St_2 \times Ac_2 \rightarrow St$. The size of \mathcal{S} is $|St_1 \times Ac_1 \cup St_2 \times Ac_2|$. A path of \mathcal{S} is an infinite sequence $s_0 s_1 s_2 \dots$ of states such that $s_{n+1} = \delta(s_n, \alpha)$ for some action α . A strategy for Player 1 is a mapping $f: St^* St_1 \rightarrow Ac_1$. A path $s_0 s_1 s_2 \dots$ is an outcome of f starting in s , if $s_0 = s$ and $s_{n+1} = \delta(s_n, f(s_0 \dots s_n))$ for all $n \geq 0$ such that $s_n \in St_1$. A two player game $\mathcal{G} = (\mathcal{S}, \text{Win})$ consists of a two-player game structure \mathcal{S} and a winning condition $\text{Win} \subseteq St^\omega$, where St is the set of states of \mathcal{S} . We say that a strategy f for Player 1 is a winning strategy for \mathcal{G} from a state s , if every outcome of f starting in s is in Win .

Given a concurrent game structure $\mathcal{S} = (St, Ag, Ac, \delta, \ell)$ and $A \subseteq Ag$, we define the two-player game structure $\mathcal{S}_A = (St_1 \cup St_2, St_1, St_2, Ac_1, Ac_2, \delta')$ where $St_1 = St$ and $St_2 = St \times Ac_1$, Ac_1 is the set of action vectors for A , Ac_2 is the set of action vectors for $Ag \setminus A$, $\delta'(s, v) = (s, v)$ for $s \in St_1$ and $v \in Ac_1$, and $\delta'((s, v), v') = \delta(s, v \oplus v')$ for $(s, v) \in St_2$ and $v' \in Ac_2$, where $v \oplus v'$ is the unique action vector for Ag induced by v and v' . Note that the size of \mathcal{S}_A is at most linear in the size of \mathcal{S} .

A path in \mathcal{S}_A alternates between states of \mathcal{S} and auxiliary states (those in $St \times Ac_1$), i.e., it is in $(St \cdot (St \times Ac_1))^\omega$. Thus, when translating paths between \mathcal{S} and \mathcal{S}_A , only states at even positions are relevant (assuming we start the path in \mathcal{S}_A in St). Hence, given a property $P \subseteq St^\omega$ of paths in \mathcal{S} , we extend it to the corresponding winning condition $P' = \{s_0 s_1 s_2 \dots \in (St \cdot (St \times Ac_1))^\omega \mid s_0 s_2 s_4 \dots \in P\}$ of paths in \mathcal{S}_A .

The next lemma reduces the (non-) existence of strategies that allow a set A of agents to enforce a property in \mathcal{S} (which formalize the semantics of $\langle\langle A \rangle\rangle$ and $\llbracket A \rrbracket$) to the (non-) existence of winning strategies for Player 1 in \mathcal{S}_A . It derives from results of de Alfaro and Henzinger [2] for concurrent ω -regular games.

Lemma 1. *Let \mathcal{S} be a concurrent game structure with set St of states containing s , let A be a subset of its agents, and let $P \subseteq St^\omega$.*

1. *There is a set F_A of strategies, one for each agent $a \in A$, such that $\text{out}(s, F_A) \subseteq P$ iff Player 1 has a winning strategy for (\mathcal{S}_A, P') from s .*
2. *For all sets F_A of strategies, one for each agent $a \in A$, $\text{out}(s, F_A) \cap P \neq \emptyset$ iff Player 1 does not have a winning strategy for $(\mathcal{S}_A, (St^\omega \setminus P)')$ from s .*

In the following, we consider the following winning conditions for a two-player game played in \mathcal{S}_A , all induced by a set $F \subseteq St$ of states:

$$\begin{aligned} \text{Next}(F) &= \{s_0 s_1 s_2 \cdots \in (St \cdot (St \times Ac_1))^\omega \mid s_2 \in F\} \\ \text{Reach}(F) &= \{s_0 s_1 s_2 \cdots \in (St \cdot (St \times Ac_1))^\omega \mid s_n \in F \text{ for some even } n\} \\ \text{Safety}(F) &= \{s_0 s_1 s_2 \cdots \in (St \cdot (St \times Ac_1))^\omega \mid s_n \in F \text{ for all even } n\} \\ \text{Büchi}(F) &= \{s_0 s_1 s_2 \cdots \in (St \cdot (St \times Ac_1))^\omega \mid \\ &\quad s_n \in F \text{ for infinitely many even } n\} \\ \text{coBüchi}(F) &= \{s_0 s_1 s_2 \cdots \in (St \cdot (St \times Ac_1))^\omega \mid \\ &\quad s_n \in F \text{ for all but finitely many even } n\} \end{aligned}$$

Again, note that these conditions only refer to even positions, as they will be used to capture a property of paths in \mathcal{S} , i.e., the auxiliary states are irrelevant.

Collectively, we refer to games with any of the above winning conditions as NRSBC games. The following result is a generalization of standard results on infinite games (see, e.g., Grädel, Thomas, and Wilke [32]) that accounts for the fact that only states at even positions are relevant.

Proposition 1. *The following problem is in PTIME: Given an NRSBC game \mathcal{G} and a state s , does Player 1 have a winning strategy for \mathcal{G} from s ?*

Proof of Theorem 1. Consider a concurrent game structure \mathcal{S} with set St of states and an rATL formula φ . We show how to inductively compute the satisfaction sets $\text{Sat}(\varphi', t) = \{s \in St \mid V(s, \varphi') \succeq t\}$ for all (state) subformulas φ' of φ and all truth values $t \in \mathbb{B}_4$. Note that $\text{Sat}(\varphi', 0000) = St$ for all formulas φ' , so these sets can be computed trivially.

The cases of atomic propositions and Boolean connectives follow straightforwardly from the definition of their semantics (cp. the semantics of rCTL [45]), so we focus on the case of formulas of the form $\langle\langle A \rangle\rangle \Phi$ or $\llbracket A \rrbracket \Phi$. Note that we only have to consider three cases for Φ , e.g., $\Phi = \odot \varphi'$, $\Phi = \diamond \varphi'$, and $\Phi = \square \varphi'$ for some state formula φ' . The following characterizations are consequences of Lemma 1:

- $s \in \text{Sat}(\langle\langle A \rangle\rangle \odot \varphi', t)$ if and only if Player 1 has a winning strategy for $(\mathcal{S}_A, \text{Next}(\text{Sat}(\varphi', t)))$ from s .
- $s \in \text{Sat}(\langle\langle A \rangle\rangle \diamond \varphi', t)$ if and only if Player 1 has a winning strategy for $(\mathcal{S}_A, \text{Reach}(\text{Sat}(\varphi', t)))$ from s .
- $s \in \text{Sat}(\langle\langle A \rangle\rangle \square \varphi', 1111)$ if and only if Player 1 has a winning strategy for $(\mathcal{S}_A, \text{Safety}(\text{Sat}(\varphi', 1111)))$ from s .
- $s \in \text{Sat}(\langle\langle A \rangle\rangle \square \varphi', 0111)$ if and only if Player 1 has a winning strategy for $(\mathcal{S}_A, \text{coBüchi}(\text{Sat}(\varphi', 0111)))$ from s .
- $s \in \text{Sat}(\langle\langle A \rangle\rangle \square \varphi', 0011)$ if and only if Player 1 has a winning strategy for $(\mathcal{S}_A, \text{Büchi}(\text{Sat}(\varphi', 0011)))$ from s .
- $s \in \text{Sat}(\langle\langle A \rangle\rangle \square \varphi', 0001)$ if and only if Player 1 has a winning strategy for $(\mathcal{S}_A, \text{Reach}(\text{Sat}(\varphi', 0001)))$ from s .

Analogously, the satisfaction of formulas $\llbracket A \rrbracket \Phi$ can be characterized by the non-existence of winning strategies for Player 1, relying on the duality of the reachability (Büchi) and safety (coBüchi) winning conditions and the self-duality of the winning condition capturing the next operator. For example, we have $s \in \text{Sat}(\llbracket A \rrbracket \odot \varphi', t)$ if and only if Player 1 does not have a winning strategy for $(\mathcal{S}_A, \text{Next}(St \setminus \text{Sat}(\varphi', t)))$ from s .

Now, to solve the model-checking problem with inputs \mathcal{S} , φ , s and t , we inductively compute all satisfaction sets $\text{Sat}(\varphi', t')$ and check whether s is in $\text{Sat}(\varphi, t)$. Using Proposition 1 and the fact that each NRSBC game we have to solve during the computation is of linear size (in $|\mathcal{S}|$), these $\mathcal{O}(|\varphi| \cdot |\mathcal{S}|)$ many sets can be computed in polynomial time, where $|\varphi|$ is the number of state subformulas of φ .

Finally, the lower bound follows from the PTIME-hardness of CTL model-checking [24], which is a fragment of rATL (see Subsection 3.3). Furthermore, let us note that the PTIME lower bound for CTL model-checking already holds for fragment without until and release [37] (recall that we do not include until and release in rATL for the sake of simplicity). \square

3.2 rATL Satisfiability

This subsection considers the satisfiability problem for rATL, which is stated as follows: Given an rATL formula φ and a truth value $t \in \mathbb{B}_4$, is there a concurrent game structure \mathcal{S} with a state s such that $V(s, \varphi) \succeq t$?

Theorem 2. *rATL satisfiability is EXPTIME-complete.*

Proof sketch. The upper bound is proven by embedding rATL into the alternating μ -calculus while the lower bound already holds for CTL, a fragment of rATL. \square

3.3 Expressiveness

The main impetus for introducing rATL is to devise a robust generalization of ATL as a powerful formalism to deal with robust strategic reasoning in multi-agent systems. A natural question is to state the expressive power of rATL with respect to ATL and the robust version of CTL (rCTL) [45]. In this subsection, we show that both ATL and rCTL can be embedded into rATL, i.e., rATL generalizes both of these logics. Furthermore, we show that rATL is strictly more expressive than both of them. We begin by comparing rATL and ATL, and show first that rATL is at least as expressive as ATL, witnessing that our robust extension is set up correctly. This fact is formalized in the lemma below, intuitively stating that the first bit of the evaluation function captures the semantics of ATL.

Lemma 2. *Let φ be an ATL formula. Then, there exists an rATL formula φ^* such that for every concurrent game structure \mathcal{S} and all states s of \mathcal{S} : $V(s, \varphi^*) = 1111$ if and only if $\mathcal{S}, s \models \varphi$.*

Proof sketch. We obtain the rATL formula φ^* as follows: First, we eliminate every implication $\varphi_1 \rightarrow \varphi_2$ in the ATL formula φ by replacing it with the expression $\neg\varphi_1 \vee \varphi_2$. Second, we bring the formula into negation normal form by pushing all negations inwards to the level of atomic propositions. Finally, we dot all the temporal operators to obtain the rATL formula φ^* . The claim of Lemma 2 can then be shown by induction over the structure of φ . \square

As we have observed above with Example 1, rATL is able to express basic forms of fairness such as "for a given structure \mathcal{S} there exists a strategy for a coalition of agents A such that a certain property p holds infinitely often". Formally this corresponds to the formula $\varphi = \langle\langle A \rangle\rangle \Box p$ with $V(s, \varphi) \succeq 0011$. As shown by Alur, Henzinger, and Kupferman [5], such a property cannot be expressed in ATL, but rather requires the more expressive logic ATL*. Indeed, it corresponds to the ATL* formula $\varphi = \langle\langle A \rangle\rangle \Box \Diamond p$. So, by using the result reported in Lemma 2, the following holds.

Theorem 3. *rATL is strictly more expressive than ATL.*

Now, we compare rATL and rCTL: The latter logic is obtained by robustifying CTL along the same lines as described in Section 3 (see [45] for detailed definitions). Let us just remark that rCTL formulas, as CTL formulas, are evaluated over Kripke structures by means of a valuation function V_{rCTL} . Thus, to compare the expressiveness of both logics, as usual, we have to interpret a Kripke structure as a (one-agent) concurrent game structure. We start by showing that rATL is at least as expressive as rCTL, just as ATL is at least as expressive as CTL.

Lemma 3. *Let φ be an rCTL formula. Then, there exists an rATL formula φ^* such that for every Kripke structure \mathcal{K} the following holds for all states s of \mathcal{K} : $V(s, \varphi^*) = V_{rCTL}(s, \varphi)$.*

Proof sketch. Our construction proceeds as follows: First, we turn a Kripke structure \mathcal{K} into a concurrent game structure with one agent a , having the same states and state labels, a suitable set of actions, and a transition function δ such that there is a transition in \mathcal{K} from s to s' if and only if $s' = \delta(s, \alpha)$ for some action α . Second, we replace each existential path quantifier \exists in φ by $\langle\langle \{a\} \rangle\rangle$ and each universal path quantifier \forall by $\langle\langle \emptyset \rangle\rangle$, obtaining the rATL formula φ^* . The claim of Lemma 3 can then be shown by induction over the structure of φ . \square

Now, we recall that Alur, Henzinger, and Kupferman [5] have observed that in ATL there are formulas that cannot be expressed in CTL. The reason is that, given a concurrent game structure, CTL can only reason about a single path (with the existential modality) or all paths (with the universal modality). Conversely, ATL can reason about an arbitrary number of paths by means of strategies. The same argument can be extended to rATL and rCTL. Thus, by putting together this observation with the statement of Lemma 3, the following holds.

Theorem 4. *rATL is strictly more expressive than rCTL.*

Notice that the argument that rATL formulas expressing fairness properties such as "infinitely often" cannot be expressed in ATL (used in Theorem 3 for the strict containment of ATL in rATL) can also be applied to rCTL. Similarly, the argument used above to show that rATL formulas cannot be translated into rCTL (used in Theorem 4 for the strict containment of rCTL in rATL) can also be applied to ATL. This leads to the following corollary.

Corollary 1. *ATL and rCTL are incomparable.*

4 Robust ATL*

Just as one generalizes CTL, rCTL, and ATL by allowing nesting of temporal operators in the scope of a single path/strategy quantifier (obtaining CTL*, rCTL*, and ATL*, respectively), we now study rATL*, the analogous generalization of rATL. Again, we will prove that adding robustness comes for free.

The formulas of rATL* are given by the grammar

$$\begin{aligned}\varphi &::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \varphi \rightarrow \varphi \mid \langle\langle A \rangle\rangle\Phi \mid \llbracket A \rrbracket\Phi \\ \Phi &::= \varphi \mid \neg\Phi \mid \Phi \vee \Phi \mid \Phi \wedge \Phi \mid \Phi \rightarrow \Phi \mid \odot\Phi \mid \diamond\Phi \mid \square\Phi\end{aligned}$$

where p ranges over atomic propositions and A over subsets of agents. Again, we distinguish between *state formulas* (those derivable from φ) and *path formulas* (those derivable from Φ). If not specified, an rATL* formula is a state formula.

The semantics of rATL* are again defined via an evaluation function V that maps a state formula and a state or a path formula and a path to a truth value in \mathbb{B}_4 . The cases for state formulas are defined as for rATL and we define for every path π , every state formula φ , and all path formulas Φ_1 and Φ_2

$$\begin{aligned}- & V(\pi, \varphi) = V(\pi[0], \varphi), \\ - & V(\pi, \neg\Phi) = \begin{cases} 0000 & \text{if } V(\pi, \Phi) = 1111, \\ 1111 & \text{if } V(\pi, \Phi) \prec 1111, \end{cases} \\ - & V(\pi, \Phi_1 \vee \Phi_2) = \max(V(\pi, \Phi_1), V(\pi, \Phi_2)), \\ - & V(\pi, \Phi_1 \wedge \Phi_2) = \min(V(\pi, \Phi_1), V(\pi, \Phi_2)), \\ - & V(\pi, \Phi_1 \rightarrow \Phi_2) = \begin{cases} 1111 & \text{if } V(\pi, \Phi_1) \preceq V(\pi, \Phi_2), \\ V(\pi, \Phi_2) & \text{if } V(\pi, \Phi_1) \succ V(\pi, \Phi_2), \end{cases} \\ - & V(\pi, \odot\Phi) = b_1 b_2 b_3 b_4 \text{ with } b_k = V(\pi[1], \Phi)[k], \\ - & V(\pi, \diamond\Phi) = b_1 b_2 b_3 b_4 \text{ with } b_k = \max_{i \geq 0} V(\pi[i], \Phi)[k], \text{ and} \\ - & V(\pi, \square\Phi) = b_1 b_2 b_3 b_4 \text{ where}\end{aligned}$$

$$\begin{aligned}b_1 &= \min_{i \geq 0} V(\pi[i], \Phi)[1], & b_3 &= \min_{i \geq 0} \max_{j \geq i} V(\pi[j], \Phi)[3], \\ b_2 &= \max_{i \geq 0} \min_{j \geq i} V(\pi[j], \Phi)[2], & b_4 &= \max_{i \geq 0} V(\pi[i], \Phi)[4].\end{aligned}$$

We show that every rATL* formula (w.r.t. a fixed truth value) can be translated into an equivalent ATL* formula of polynomial size. This allows us to settle the complexity of rATL* model-checking and satisfiability as well as the expressiveness of rATL*. Below, \models denotes the ATL* satisfaction relation [5].

Lemma 4. *For every $rATL^*$ formula φ and every truth value $t \in \mathbb{B}_4$, there is an ATL^* formula φ_t such that $V(s, \varphi) \succeq t$ if and only if $\mathcal{S}, s \models \varphi_t$. Furthermore, the function mapping φ and t to φ_t is polynomial-time computable.*

The $rATL^*$ model-checking and satisfiability problems are defined as their counterparts for $rATL$. Both model-checking and satisfiability for ATL^* are $2EXPTIME$ -complete [5,49]. Due to Lemma 4, we obtain the same results for $rATL^*$, thereby showing that adding robustness comes indeed for free.

Theorem 5. *The $rATL^*$ model-checking problem and the $rATL^*$ satisfiability problem are both $2EXPTIME$ -complete.*

Another consequence of the translation from $rATL^*$ to ATL^* and the fact that ATL^* is a fragment of $rATL^*$ is that both logics are equally expressive.

Corollary 2. *$rATL^*$ and ATL^* are equally expressive.*

5 A Practical Example

Let us consider a smart grid with a set U of utility companies and a set C of consumers. Assume that for every consumer $c \in C$ there is a proposition ℓ_c indicating that c 's energy consumption is within the pre-agreed limit. Conversely, c 's consumption is higher than the limit if ℓ_c is violated. Furthermore, there is a proposition “stable” that holds true if and only if the grid is stable (i.e., the utility companies coordinate to provide the right amount of electricity).

Let us now consider the ATL^* formula

$$\langle\langle U \rangle\rangle [C] (\Box \bigwedge_{c \in C} \ell_c) \rightarrow \Box \text{stable}.$$

This formula expresses that the utility companies U have a strategy such that no matter how the consumers behave, the following is satisfied: if each consumer's consumption always stays within their limit, then the utility companies keep the grid always stable. However, this specification is not robust and provides only limited information when satisfied: even if a single consumer exceeds their limit once, there is no further obligation on the utility companies, and the formula is satisfied independently of whether the grid is always stable or not.

So, let us illustrate how the $rATL^*$ formula

$$\varphi = \langle\langle U \rangle\rangle [C] (\Box \bigwedge_{c \in C} \ell_c) \rightarrow \Box \text{stable}$$

does capture robustness. To this end, assume for now that φ evaluates to 1111. Then, there is a strategy for U such that for all outcomes π that are consistent with that strategy, the following holds:

- If $\bigwedge_{c \in C} \ell_c$ holds in every position of π , i.e., $\Box \bigwedge_{c \in C} \ell_c$ evaluates to 1111 then by the semantics of \rightarrow the formula $\Box \text{stable}$ also evaluates to 1111. This means the proposition “stable” also holds in every position. Therefore, the grid supply is always stable. Hence, the desired goal is retained when the assumption regarding the consumers holds with no violation. Note that this is equivalent to what the original ATL^* formula above expresses.

- Assume now that the consumer assumption $\bigwedge_{c \in C} \ell_c$ is violated finitely many times, i.e., finitely often some consumer violates their consumption limit. This means that the formula $\Box \bigwedge_{c \in C} \ell_c$ evaluates to 0111. Then, by the semantics of rATL*, \Box stable evaluates to 0111 or higher, which means that “stable” holds at every state, except for a finite number of times. So, the degree of violation of the guarantee required by U is at most the degree of violation of the assumptions on the consumers.
- Similarly, if $\Box \bigwedge_{c \in C} \ell_c$ holds infinitely (finitely) often, then \Box stable holds infinitely (finitely) often.

If the formula φ evaluates to 1111, then U has a strategy that does not behave arbitrarily in case the assumption $\Box \bigwedge_{c \in C} \ell_c$ fails, but instead satisfies the guarantee \Box stable to at least the same degree that the guarantee holds.

Finally, even if φ evaluates to a truth value $t < 1111$, this reveals crucial information about U 's ability to guarantee a stable grid, i.e., the premise $\Box \bigwedge_{c \in C} \ell_c$ evaluates to some truth value $t' > t$ while the conclusion “stable” evaluates to t .

6 Discussion and Future Work

This paper introduces rATL and rATL*, the first logic formalisms able to deal with robust strategic reasoning in multi-agent systems. As we have shown along the paper, rATL results to be very expressive, useful in practice, and not more costly than the subsumed logics ATL and rCTL. Similarly, rATL* is not more costly than the subsumed logic ATL*.

The positive results about rATL represent the foundation for a number of useful extensions, mainly by extending robustness to logics for strategic reasoning that are more expressive than ATL and ATL* such as Strategy Logic [44] and the like. Notably, Strategy Logic is much more expressive than ATL* [5]. Indeed it can express several game-theoretic concepts including Nash Equilibria over LTL goals. Interestingly, the formula expressing Nash Equilibria uses an implication. In words the formula says that n agents' strategies $\sigma_1, \dots, \sigma_n$ form an equilibrium if, for every agent, it holds that whenever by unilaterally changing her strategy the goal is also satisfied, then it *implies* that the goal is satisfied with the original tuple of strategies as well. Robustness in Strategy Logic (by means of rLTL goals in place of LTL) then allows to define a stronger notion of Nash Equilibrium.

Another interesting direction for future work is to come up with an implementation of the model-checking procedure for rATL, possibly by extending existing tools such as MCMAS [41,20].

Acknowledgments. This research has been supported by the PRIN project RIPER (No. 20203FFYLK), the PNRR MUR project PE0000013-FAIR, the InDAM project “Strategic Reasoning in Mechanism Design”, and DIREC - Digital Research Centre Denmark. Furthermore, this work has been financially supported by Deutsche Forschungsgemeinschaft, DFG Project numbers 434592664 and 459419731, and the Research Center Trustworthy Data Science and Security (<https://rc-trust.ai>), one of the Research Alliance centers within the UA Ruhr (<https://uaruhr.de>).

References

1. Afzal, S., Mokhlis, H., Illias, H.A., Mansor, N.N., Shareef, H.: State-of-the-art review on power system resilience and assessment techniques. *IET Generation, Transmission & Distribution* **14**(25), 6107–6121 (2020)
2. de Alfaro, L., Henzinger, T.A.: Concurrent omega-regular games. In: *LICS 2000*. pp. 141–154. IEEE Computer Society (2000). <https://doi.org/10.1109/LICS.2000.855763>
3. de Alfaro, L., Henzinger, T.A., Majumdar, R.: From verification to control: Dynamic programs for omega-regular objectives. In: *LICS 2001*. pp. 279–290. IEEE (2001)
4. Almagor, S., Boker, U., Kupferman, O.: Formally reasoning about quality. *J. ACM* **63**(3), 24:1–24:56 (2016). <https://doi.org/10.1145/2875421>
5. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *J. ACM* **49**(5), 672–713 (2002). <https://doi.org/10.1145/585265.585270>
6. Alur, R., Torre, S.L., Madhusudan, P.: Playing games with boxes and diamonds. In: Amadio, R.M., Lugiez, D. (eds.) *CONCUR 2003*. LNCS, vol. 2761, pp. 127–141. Springer (2003). https://doi.org/10.1007/978-3-540-45187-7_8
7. Aminof, B., Kwiatkowska, M., Maubert, B., Murano, A., Rubin, S.: Probabilistic strategy logic. In: Kraus, S. (ed.) *IJCAI 2019*. pp. 32–38. *ijcai.org* (2019). <https://doi.org/10.24963/ijcai.2019/5>
8. Aminof, B., Malvone, V., Murano, A., Rubin, S.: Graded modalities in strategy logic. *Information and Computation* **261**, 634–649 (2018)
9. Anevclavis, T., Neider, D., Philippe, M., Tabuada, P.: Evrostos: the rLTL verifier. In: Ozay, N., Prabhakar, P. (eds.) *HSCC 2019*. pp. 218–223. ACM (2019). <https://doi.org/10.1145/3302504.3311812>
10. Anevclavis, T., Philippe, M., Neider, D., Tabuada, P.: Verifying rLTL formulas: now faster than ever before! In: *CDC 2018*. pp. 1556–1561. IEEE (2018). <https://doi.org/10.1109/CDC.2018.8619014>
11. Anevclavis, T., Philippe, M., Neider, D., Tabuada, P.: Being correct is not enough: Efficient verification using robust linear temporal logic. *ACM Trans. Comput. Log.* **23**(2), 8:1–8:39 (2022). <https://doi.org/10.1145/3491216>
12. Belardinelli, F., Jamroga, W., Kurpiewski, D., Malvone, V., Murano, A.: Strategy logic with simple goals: Tractable reasoning about strategies. In: Kraus, S. (ed.) *IJCAI 2019*. pp. 88–94. *ijcai.org* (2019). <https://doi.org/10.24963/ijcai.2019/13>
13. Belardinelli, F., Lomuscio, A., Murano, A., Rubin, S.: Verification of multi-agent systems with public actions against strategy logic. *Artificial Intelligence* **285**, 103302 (2020)
14. Belhaiza, S., Baroudi, U.: A game theoretic model for smart grids demand management. *IEEE Transactions on Smart Grid* **6**(3), 1386–1393 (2014)
15. Berthon, R., Maubert, B., Murano, A., Rubin, S., Vardi, M.Y.: Strategy logic with imperfect information. *ACM Transactions on Computational Logic (TOCL)* **22**(1), 1–51 (2021)
16. Bevrani, H.: *Robust power system frequency control*, vol. 4. Springer (2014)
17. Bloem, R., Chatterjee, K., Greimel, K., Henzinger, T.A., Jobstmann, B.: Robustness in the presence of liveness. In: *CAV 2010*. pp. 410–424. Springer (2010)
18. Bouyer, P., Kupferman, O., Markey, N., Maubert, B., Murano, A., Perelli, G.: Reasoning about quality and fuzziness of strategic behaviors. *ACM Trans. Comput. Log.* **24**(3), 21:1–21:38 (2023)

19. Bouyer, P., Markey, N., Reynier, P.A.: Robust analysis of timed automata via channel machines. In: FOSSCAS 2008. pp. 157–171. Springer (2008)
20. Cermák, P., Lomuscio, A., Murano, A.: Verifying and synthesising multi-agent systems against one-goal strategy logic specifications. In: Bonet, B., Koenig, S. (eds.) Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25–30, 2015, Austin, Texas, USA. pp. 2038–2044. AAAI Press (2015)
21. Chaaban, Y., Müller-Schloer, C.: A survey of robustness in multi-agent systems. In: Cognitive13, Fifth International Conference on Advanced Cognitive Technologies and Applications. pp. 7–13 (2013)
22. Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy logic. *Information and Computation* **208**(6), 677–693 (2010). <https://doi.org/10.1016/j.ic.2009.07.004>
23. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching-time temporal logic. In: Kozen, D. (ed.) Logics of Programs, Workshop, Yorktown Heights, New York, USA, May 1981. pp. 52–71. Springer (1981). <https://doi.org/10.1007/BFb0025774>
24. Clarke, E.M., Emerson, E.A., Sistla, A.P.: Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Trans. Program. Lang. Syst.* **8**(2), 244–263 (1986). <https://doi.org/10.1145/5397.5399>
25. Dallal, E., Neider, D., Tabuada, P.: Synthesis of safety controllers robust to unmodeled intermittent disturbances. In: CDC 2016. pp. 7425–7430. IEEE (2016)
26. Dima, C., Tiplea, F.L.: Model-checking ATL under imperfect information and perfect recall semantics is undecidable. arXiv preprint arXiv:1102.4225 (2011)
27. Donzé, A., Maler, O.: Robust satisfaction of temporal logic over real-valued signals. In: FORMATS 2010. pp. 92–106. Springer (2010)
28. Doyen, L., Henzinger, T.A., Legay, A., Nickovic, D.: Robustness of sequential circuits. In: ACSD 2010. pp. 77–84. IEEE (2010)
29. Emerson, E.A., Halpern, J.Y.: Decision procedures and expressiveness in the temporal logic of branching time. *J. Comput. Syst. Sci.* **30**(1), 1–24 (1985). [https://doi.org/10.1016/0022-0000\(85\)90001-7](https://doi.org/10.1016/0022-0000(85)90001-7)
30. Faella, M., Napoli, M., Parente, M.: Graded alternating-time temporal logic. *Fundamenta Informaticae* **105**(1-2), 189–210 (2010)
31. French, T., Cabe-Dansted, M., John, C., Reynolds, M.: A temporal logic of robustness. In: International Symposium on Frontiers of Combining Systems. pp. 193–205. Springer (2007)
32. Grädel, E., Thomas, W., Wilke, T. (eds.): Automata, Logics, and Infinite Games: A Guide to Current Research, LNCS, vol. 2500. Springer (2002). <https://doi.org/10.1007/3-540-36387-4>
33. Hassan, S.R.M., Hasan, N., Siddique, M.A., Fahim, K.S., Rahman, R., Iftekhar, L.: Incorporating multi-agent systems technology in power and energy systems of bangladesh: A feasibility study. In: ICREST 2021. pp. 342–347. IEEE (2021)
34. Huang, X., Luo, C.: A logic of probabilistic knowledge and strategy. In: AAMAS. pp. 845–852. Citeseer (2013)
35. Jamroga, W., Murano, A.: On module checking and strategies. In: AAMAS 2014. pp. 701–708 (2014)
36. Kaur, D., Uslu, S., Rittichier, K.J., Durresi, A.: Trustworthy artificial intelligence: a review. *ACM Computing Surveys (CSUR)* **55**(2), 1–38 (2022)
37. Krebs, A., Meier, A., Mundhenk, M.: The model checking fingerprints of CTL operators. *Acta Informatica* **56**(6), 487–519 (2019). <https://doi.org/10.1007/s00236-018-0326-9>

38. Kupferman, O., Vardi, M.Y., Wolper, P.: Module checking. *Information and Computation* **164**(2), 322–344 (2001)
39. Laroussinie, F., Markey, N.: Augmenting ATL with strategy contexts. *Information and Computation* **245**, 98–123 (2015). <https://doi.org/10.1016/j.ic.2014.12.020>
40. Lempert, R.J., Schlesinger, M.E.: Robust strategies for abating climate change. *Climatic Change* **45**(3-4), 387–401 (2000)
41. Lomuscio, A., Qu, H., Raimondi, F.: MCMAS: an open-source model checker for the verification of multi-agent systems. *Int. J. Softw. Tools Technol. Transf.* **19**(1), 9–30 (2017). <https://doi.org/10.1007/s10009-015-0378-x>
42. Meier, A., Mundhenk, M., Thomas, M., Vollmer, H.: The complexity of satisfiability for fragments of CTL and CTL^{*}. In: Halava, V., Potapov, I. (eds.) RP 2008. ENTCS, vol. 223, pp. 201–213. Elsevier (2008). <https://doi.org/10.1016/j.entcs.2008.12.040>
43. Modas, A., Sanchez-Matilla, R., Frossard, P., Cavallaro, A.: Toward robust sensing for autonomous vehicles: An adversarial perspective. *IEEE Signal Processing Magazine* **37**(4), 14–23 (2020)
44. Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.* **15**(4), 34:1–34:47 (2014). <https://doi.org/10.1145/2631917>
45. Nayak, S.P., Neider, D., Roy, R., Zimmermann, M.: Robust computation tree logic. In: NFM 2022. pp. 538–556. Springer (2022). https://doi.org/10.1007/978-3-031-06773-0_29
46. Omnes, L., Marot, A., Donnot, B.: Adversarial training for a continuous robustness control problem in power systems. In: 2021 IEEE Madrid PowerTech. pp. 1–6. IEEE (2021)
47. Pinto, L., Davidson, J., Sukthankar, R., Gupta, A.: Robust adversarial reinforcement learning. In: International Conference on Machine Learning. pp. 2817–2826. PMLR (2017)
48. Sampaio, R.F., Melo, L.S., Leão, R.P., Barroso, G.C., Bezerra, J.R.: Automatic restoration system for power distribution networks based on multi-agent systems. *IET Generation, Transmission & Distribution* **11**(2), 475–484 (2017)
49. Schewe, S.: ATL^{*} satisfiability is 2EXPTIME-complete. In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfssdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part II. LNCS, vol. 5126, pp. 373–385. Springer (2008). https://doi.org/10.1007/978-3-540-70583-3_31
50. Schnoor, H.: Epistemic and probabilistic ATL with quantification and explicit strategies. In: ICAART. pp. 131–148. Springer (2013)
51. Singh, V.P., Kishor, N., Samuel, P.: Distributed multi-agent system-based load frequency control for multi-area power system in smart grid. *IEEE Transactions on Industrial Electronics* **64**(6), 5151–5160 (2017)
52. Tabuada, P., Neider, D.: Robust linear temporal logic. In: CSL 2016. LIPIcs, vol. 62, pp. 10:1–10:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016). <https://doi.org/10.4230/LIPIcs.CSL.2016.10>
53. Veres, S.M., Molnar, L., Lincoln, N.K., Morice, C.P.: Autonomous vehicle control systems—a review of decision making. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering* **225**(2), 155–195 (2011)
54. Walther, D., Lutz, C., Wolter, F., Wooldridge, M.: ATL satisfiability is indeed EXPTIME-complete. *Journal of Logic and Computation* **16**(6), 765–787 (2006)

A Proof of Theorem 2

Recall that Theorem 2 states that rATL satisfiability is EXPTIME-complete.

Proof of Theorem 2. For the upper bound we follow a reasoning similar to the one used for the satisfiability for rCTL [45], opportunely extended to deal with ATL. Precisely, rATL satisfiability can be solved by translating a given rATL formula and a given truth value into an equivalent alternating μ -calculus formula (see de Alfaro and Henzinger [2] and de Alfaro, Henzinger, and Majumdar [3]) of linear size and then checking the resulting formula for satisfiability. The procedure to which we refer is used to translate an ATL* formula φ into an equivalent alternating μ -calculus formula φ' , and relies on the NRSBC games introduced in Subsection 3.1.

Note that the translation from ATL* to alternating μ -calculus is in general exponential in the maximal number of nested temporal operators. More precisely, let $\Sigma = \{\circlearrowleft, U, R, \diamond, \square\}$ the set of all classical temporal operators used in ATL*. The translation is exponential in the length i of the longest path in the syntax tree labeled by operators in Σ . In our specific case of φ being an rATL formula, we have $i = 2$ since the interpretations of the temporal operators \circlearrowleft , \diamond , and \square correspond to a nesting of at most two classical temporal operators in Σ . Since the satisfiability problem for the alternating μ -calculus is EXPTIME-complete [5], rATL satisfiability is in EXPTIME.

A matching lower bound already holds for CTL [29], and thus also for rATL. Aain, the lower bound for CTL already holds for the fragment without until and release [42]. \square

B Proof of Lemma 2

In order to prove Lemma 2, we have to introduce the logic ATL first [5]. We then show the statement by induction.

The logic ATL The syntax of ATL is identical to that of rATL as introduced in Section 3, except for “un-dotted” temporal operators. More precisely, its syntax is given by the grammar

$$\begin{aligned} \varphi &::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\Phi \mid \llbracket A \rrbracket\Phi \\ \Phi &::= \circlearrowleft\varphi \mid \diamond\varphi \mid \square\varphi, \end{aligned}$$

where p ranges over atomic propositions and A ranges over subsets of agents. As usual, ATL distinguishes between state formulas (those derivable from φ) and path formulas (those derivable from Φ). If not specified otherwise, an ATL formula is a state formula.

Note that we have included conjunction and the $\llbracket \]$ -operator explicitly in the syntax. We do this because we will later only consider ATL formulas in negation normal form. Furthermore, let us mention that we do not consider the until-operator since we have defined rATL without it.

The semantics of ATL is defined in terms of a satisfaction relation \models that relates the states/paths of a concurrent game structure with all ATL formulas that are satisfied at that state. In what follows, let $\mathcal{S} = (St, Ag, Ac, \delta, \ell)$ be a concurrent game structure, s a state of \mathcal{S} , π a path of \mathcal{S} , and $S \subseteq Ag$ a set of agents.

$$\begin{aligned}
\mathcal{S}, s &\models p \Leftrightarrow p \in \ell(s) \\
\mathcal{S}, s &\models \neg\varphi \Leftrightarrow \mathcal{S}, s \not\models \varphi, \\
\mathcal{S}, s &\models \varphi_1 \vee \varphi_2 \Leftrightarrow \mathcal{S}, s \models \varphi_1 \text{ or } \mathcal{S}, s \models \varphi_2, \\
\mathcal{S}, s &\models \varphi_1 \wedge \varphi_2 \Leftrightarrow \mathcal{S}, s \models \varphi_1 \text{ and } \mathcal{S}, s \models \varphi_2, \\
\mathcal{S}, s &\models \langle\langle A \rangle\rangle\Phi \Leftrightarrow \text{there is a set } F_A \text{ of strategies, one for} \\
&\quad \text{each agent in } A, \text{ such that for all paths} \\
&\quad \pi \in \text{out}(s, F_A) \text{ we have } \mathcal{S}, \pi \models \Phi, \\
\mathcal{S}, s &\models \llbracket A \rrbracket\Phi \Leftrightarrow \text{for every set } F_A \text{ of strategies, one for} \\
&\quad \text{each agent in } A, \text{ there exists a path} \\
&\quad \pi \in \text{out}(s, F_A) \text{ such that } \mathcal{S}, \pi \models \Phi.
\end{aligned}$$

Moreover, we have

$$\begin{aligned}
\mathcal{S}, \pi &\models \bigcirc\varphi \Leftrightarrow \mathcal{S}, \pi[1] \models \varphi \\
\mathcal{S}, \pi &\models \diamond\varphi \Leftrightarrow \mathcal{S}, \pi[i] \models \varphi \text{ for an } i \geq 0, \\
\mathcal{S}, \pi &\models \square\varphi \Leftrightarrow \mathcal{S}, \pi[i] \models \varphi \text{ for all } i \geq 0.
\end{aligned}$$

We are now ready to prove Lemma 2.

Proof Recall the statement of Lemma 2: for every ATL formula φ , there exists an rATL formula φ^* such that for every concurrent game structure \mathcal{S} and every state s of \mathcal{S} we have $V(s, \varphi^*) = 1111$ if and only if $\mathcal{S}, s \models \varphi$.

Without loss of generality, we assume in the remainder that the ATL formula φ is in negation normal form (i.e., the only Boolean operators are negation, disjunction, and conjunction, and negation does only appear at atomic propositions). We make this assumption because the semantics of the implication does not easily map to the semantics of its non-robust counterpart. Moreover, pushing negations to the atomic propositions slightly simplifies the proof of Lemma 2. It is worth noting, however, that this is a minor deviation from the literature (e.g., Nayak et al. [45] only removed implications but did not go to negation normal form).

If φ is not in negation normal form, we can easily transform it as mentioned in the proof sketch of Lemma 2 in Subsection 3.3:

1. We eliminate every implication $\varphi_1 \rightarrow \varphi_2$ in φ by replacing it with the expression $\neg\varphi_1 \vee \varphi_2$.
2. We push all negations inwards to the level of atomic propositions.

It is not hard to verify that these two steps indeed transform any ATL formula into an equivalent ATL formula in negation normal form.

To construct the rATL formula φ^* , we apply the following recursive transformation t , which essentially just “dots” the temporal operators. For state formulas, we define

$$\begin{aligned} t(p) &:= p, \\ t(\neg p) &:= \neg p, \\ t(\varphi_1 \vee \varphi_2) &:= t(\varphi_1) \vee t(\varphi_2), \\ t(\varphi_1 \wedge \varphi_2) &:= t(\varphi_1) \wedge t(\varphi_2), \\ t(\langle\langle A \rangle\rangle\Phi) &:= \langle\langle A \rangle\rangle t(\Phi), \\ t(\llbracket A \rrbracket\Phi) &:= \llbracket A \rrbracket t(\Phi). \end{aligned}$$

For path formulas, we define

$$\begin{aligned} t(\bigcirc \varphi) &:= \bigcirc t(\varphi), \\ t(\diamond \varphi) &:= \diamond t(\varphi), \\ t(\square \varphi) &:= \square t(\varphi). \end{aligned}$$

Given an ATL formula φ in negation normal form, we then simply set $\varphi^* := t(\varphi)$. Note that φ and φ^* are structurally identical, except for the “dotting” of temporal operators.

To prove Lemma 2, we show the following, slightly more technical statement.

Lemma 5. *For every ATL state formula φ in negation normal form, every ATL path formula Φ in negation normal form, every concurrent game structure \mathcal{S} , every state s of \mathcal{S} , and every path π of \mathcal{S} , the transformed formulas $t(\varphi)$ and $t(\Phi)$ satisfy*

1. $V(s, t(\varphi)) = 1111$ if and only if $\mathcal{S}, s \models \varphi$; and
2. $V(\pi, t(\Phi)) = 1111$ if and only if $\mathcal{S}, \pi \models \Phi$.

Proof of Lemma 5. We prove the statement by induction over the structure of φ .

For the base case, let $\varphi = p$. We first observe that $t(\varphi) = p$ as well. Hence, we have $\mathcal{S}, s \models p$ and $V(s, p) = 1111$ if $p \in \ell(s)$ by the definitions of the ATL and rATL semantics (recall that ℓ is the labeling function of the concurrent game structure). Conversely, $\mathcal{S}, s \not\models p$ and $V(s, p) = 0000$ if $p \notin \ell(s)$. Thus, the claim holds.

For the induction step, we make a case distinction:

- Let $\varphi = \neg\varphi_1$. Since φ is in negation normal form, we know that negation can only appear at the level of atomic propositions (i.e., $\varphi_1 = p$). Therefore, $\varphi = \neg p$ and $t(\varphi) = \neg p$. We now distinguish between $p \in \ell(s)$ and $p \notin \ell(s)$. If $p \notin \ell(s)$, then $V(s, p) = 0000$ and $V(s, \neg p) = 1111$ by definition of the rATL semantics. Similarly, we have $\mathcal{S}, s \not\models p$ and $\mathcal{S}, s \models \neg p$ by definition of the ATL semantics.

On the other hand, if $p \in \ell(s)$, then $V(s, p) = 1111$ and $V(s, \neg p) = 0000$. Similarly, we have $\mathcal{S}, s \models p$ and $\mathcal{S}, s \not\models \neg p$.

- Let $\varphi = \varphi_1 \vee \varphi_2$ and, hence, $t(\varphi) = t(\varphi_1) \vee t(\varphi_2)$. By induction hypothesis, we know that $V(s, t(\varphi_1)) = 1111$ if and only if $\mathcal{S}, s \models \varphi_1$ and $V(s, t(\varphi_2)) = 1111$ if and only if $\mathcal{S}, s \models \varphi_2$.

We now distinguish the two cases $V(s, t(\varphi)) = 1111$ and $V(s, t(\varphi)) \neq 1111$. If $V(s, t(\varphi)) = 1111$, then $V(s, t(\varphi_1)) = 1111$ or $V(s, t(\varphi_2)) = 1111$ because $V(s, t(\varphi)) = \max\{V(s, t(\varphi_1)), V(s, t(\varphi_2))\}$ by definition of the rATL semantics. Applying the induction hypothesis then yields $\mathcal{S}, s \models \varphi_1$ or $\mathcal{S}, s \models \varphi_2$. Thus, $\mathcal{S}, s \models \varphi$.

Similarly, if $V(s, t(\varphi)) \neq 1111$, then $V(s, t(\varphi)) < 1111$ because 1111 is the largest truth value in \mathbb{B}_4 . Thus, $V(s, t(\varphi_1)) < 1111$ and $V(s, t(\varphi_2)) < 1111$ by definition of the rATL semantics, showing that $V(s, t(\varphi_1)) \neq 1111$ and $V(s, t(\varphi_2)) \neq 1111$. Applying the induction hypothesis then yields $\mathcal{S}, s \not\models \varphi_1$ and $\mathcal{S}, s \not\models \varphi_2$, implying $\mathcal{S}, s \not\models \varphi$ since $\varphi = \varphi_1 \vee \varphi_2$.

- Let $\varphi = \varphi_1 \wedge \varphi_2$ and, hence, $t(\varphi) = t(\varphi_1) \wedge t(\varphi_2)$. Since this case is analogous to $\varphi = \varphi_1 \vee \varphi_2$ (with min used for max), we skip it here.
- For a set $A \subseteq Ag$ of agents, let $\varphi = \langle\langle A \rangle\rangle \Phi$ and, hence, $t(\varphi) = \langle\langle A \rangle\rangle t(\Phi)$.

If $V(s, t(\varphi)) = 1111$, then there exists a set F_A of strategies, one for each agent in A , such that for all paths $\pi \in out(s, F_A)$ we have $V(\pi, t(\Phi)) \geq 1111$. In particular, this means that $V(\pi, t(\Phi)) = 1111$ holds for all such paths π since 1111 is the largest truth value in \mathbb{B}_4 . By applying the induction hypothesis, we also know that the strategies in F_A ensure $\mathcal{S}, \pi \models \Phi$ for all paths $\pi \in out(s, F_A)$. Thus, $\mathcal{S}, s \models \varphi$ since $\varphi = \langle\langle A \rangle\rangle \Phi$.

On the other hand, if $V(s, t(\varphi)) \neq 1111$, then $V(s, t(\varphi)) < 1111$. This means that for every set F_A of strategies, there exists a path $\pi \in out(s, F_A)$ with $V(\pi, t(\Phi)) < 1111$ (in particular, $V(\pi, t(\Phi)) \neq 1111$). Applying the induction hypothesis then yields that for every set F_A of strategies, there exists a path $\pi \in out(s, F_A)$ with $\mathcal{S}, \pi \not\models \Phi$. Hence, $\mathcal{S}, s \not\models \varphi$.

- For a set $A \subseteq Ag$ of agents, let $\varphi = \llbracket A \rrbracket \Phi$ and, hence, $t(\varphi) = \llbracket A \rrbracket t(\Phi)$. Since this case is analogous to $\varphi = \langle\langle A \rangle\rangle \Phi$ (with the case distinction for $V(s, t(\varphi)) = 1111$ and $V(s, t(\varphi)) \neq 1111$ swapped), we skip it here.

- Let $\varphi = \bigcirc \varphi_1$ and, hence, $t(\varphi) = \bigcirc t(\varphi_1)$. Since $V(\pi, \bigcirc t(\varphi_1)) = b_1 b_2 b_3 b_4$ with $b_k = V(\pi[1], t(\varphi_1))[k]$ by definition of the rATL semantics, we have $V(\pi, \bigcirc t(\varphi_1)) = 1111$ if and only if $V(\pi[1], t(\varphi_1)) = 1111$. By induction hypothesis, this is equivalent to $\mathcal{S}, \pi[1] \models \varphi_1$ and, by definition of the ATL semantics, $\mathcal{S}, \pi \models \varphi$ since $\varphi = \bigcirc \varphi_1$.

- Let $\varphi = \bigcirc \varphi_1$ and, hence, $t(\varphi) = \bigcirc t(\varphi_1)$. Since $V(\pi, \bigcirc t(\varphi_1)) = b_1 b_2 b_3 b_4$ with $b_k = \max_{i \geq 0} V(\pi[i], t(\varphi_1))[k]$ by definition of the rATL semantics, we know that $V(\pi, \bigcirc t(\varphi_1)) = 1111$ if and only if there exists an $i \geq 0$ with $V(\pi[i], t(\varphi_1)) = 1111$. By induction hypothesis, this is equivalent to $\mathcal{S}, \pi[i] \models \varphi_1$ and, by definition of the ATL semantics, $\mathcal{S}, \pi \models \varphi$ since $\varphi = \bigcirc \varphi_1$.

- Let $\varphi = \square \varphi_1$ and, hence, $t(\varphi) = \square t(\varphi_1)$. By definition of the rATL semantics, we know that $V(\pi, \square t(\varphi_1))$ is a truth value $b_1 b_2 b_3 b_4 \in \mathbb{B}_4$ with $b_1 = \min_{i \geq 0} V(\pi[i], t(\varphi_1))[1]$. Thus, $V(\pi, \square t(\varphi_1)) =$

1 can only hold if $V(\pi[i], t(\varphi_1))[1] = 1$ holds for all $i \geq 0$. Since 1111 is the largest truth value, this is equivalent to the statement that $V(\pi[i], t(\varphi_1)) = 1111$ holds for all $i \geq 0$. Moreover, by applying the induction hypothesis, we obtain that $\mathcal{S}, \pi[i] \models \varphi_1$ holds for all $i \geq 0$. Thus, $\mathcal{S}, \pi \models \varphi$ since $\varphi = \Box\varphi_1$.

This concludes the proof. \square

Lemma 2 now follows immediately from Item 1 of Lemma 5, the fact that every ATL formula can be transformed into negation normal form, and the fact that $\varphi^* = t(\varphi)$.

C Definition of rCTL

For the reader's convenience, we here repeat the definition of rCTL as introduced by Nayak et al. [45].

Syntax of rCTL Formulas of rCTL are classified into state and path formulas. rCTL state formulas are formed according to the grammar

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \varphi \rightarrow \varphi \mid \exists\Phi \mid \forall\Phi,$$

where $p \in \text{AP}$ is an atomic proposition and Φ is a path formula. On the other hand, rCTL path formulas are derived from the grammar

$$\Phi ::= \odot\varphi \mid \diamond\varphi \mid \Box\varphi.$$

If not stated otherwise, an rCTL formula is a state formula. Again, we omit the until-operator for the sake of simplicity since our definition of rATL does not contain it.

Kripke structures An rCTL formula is evaluated on a Kripke structure, a mathematical object modeling a system. Following the notation of Nayak et al. [45], a Kripke structure over a set AP of atomic propositions is a tuple $\mathcal{K} = (S, I, R, L)$ where S is a finite set of states, $I \subseteq S$ is the set of initial states, $R \subseteq S \times S$ is the a transition relation, and $L: S \rightarrow 2^{\text{AP}}$ is the labeling function. Without loss of generality, we assume that for all states $s \in S$, there exists a state s' satisfying $(s, s') \in R$.

A path of the Kripke structure \mathcal{K} is an infinite sequence of states $\pi = s_0 s_1 \dots$ such that $(s_i, s_{i+1}) \in R$ for every $i \geq 0$. Moreover, let $paths(s)$ denote the set of all paths starting from state $s \in S$. Finally, for a path π and $i \geq 0$, we use $\pi[i]$ to denote the i -th state of π and $\pi[i..]$ to denote the suffix of π starting at index i .

Semantics of rCTL The rCTL semantics is a mapping V_{CTL} that assigns an element of \mathbb{B}_4 to every pair of state and state formula and to every pair of path

and path formula. Let us begin with state formulas, where $s \in S$ is a state and $p \in AP$ is an atomic proposition:

$$\begin{aligned} V_{rCTL}(s, p) &= \begin{cases} 0000 & \text{if } p \notin L(s); \text{ and} \\ 1111 & \text{if } p \in L(s), \end{cases} \\ V_{rCTL}(s, \neg\Phi) &= \begin{cases} 0000 & \text{if } V_{rCTL}(s, \Phi) = 1111; \text{ and} \\ 1111 & \text{if } V_{rCTL}(s, \Phi) \prec 1111, \end{cases} \\ V_{rCTL}(s, \Phi \vee \Psi) &= \max\{V_{rCTL}(s, \Phi), V_{rCTL}(s, \Psi)\}, \\ V_{rCTL}(s, \Phi \wedge \Psi) &= \min\{V_{rCTL}(s, \Phi), V_{rCTL}(s, \Psi)\}, \\ V_{rCTL}(s, \Phi \rightarrow \Psi) &= V_{rCTL}(s, \Phi) \Rightarrow V_{rCTL}(s, \Psi). \end{aligned}$$

where, for two truth values $a, b \in \mathbb{B}_4$, we have

$$a \Rightarrow b = \begin{cases} 1111 & \text{if } a \preceq b; \text{ and} \\ b & \text{if } a \succ b. \end{cases}$$

For the existential and universal path quantification, the definition of V_{rCTL} is as follows:

$$\begin{aligned} V_{rCTL}(s, \exists\varphi) &= \max_{\pi \in \text{paths}(s)} V_{rCTL}(\pi, \varphi), \\ V_{rCTL}(s, \forall\varphi) &= \min_{\pi \in \text{paths}(s)} V_{rCTL}(\pi, \varphi). \end{aligned}$$

Finally, the semantics for path formulas is as follows:

$$\begin{aligned} V_{rCTL}(\pi, \odot\Phi) &= V_{rCTL}(\pi[1], \Phi), \\ V_{rCTL}(\pi, \diamond\Phi) &= \max_{i \geq 0} V_{rCTL}(\pi[i], \Phi), \\ V(\pi, \square\Phi) &= b_1 b_2 b_3 b_4, \end{aligned}$$

where

$$\begin{aligned} b_1 &= \min_{i \geq 0} V_{rCTL}(\pi[i], \varphi)[1], \\ b_2 &= \max_{j \geq 0} \min_{i \geq j} V_{rCTL}(\pi[i], \varphi)[2], \\ b_3 &= \min_{j \geq 0} \max_{i \geq j} V_{rCTL}(\pi[i], \varphi)[3], \\ b_4 &= \max_{i \geq 0} V_{rCTL}(\pi[i], \varphi)[4]. \end{aligned}$$

We omit the semantics of the Boolean operators as it is same we have given for rATL. Recall that for a truth value $b = b_1 b_2 b_3 b_4 \in \mathbb{B}_4$ and $i \in \{1, \dots, 4\}$, we use $b[i]$ to denote the bit b_i , as for rATL.

D Proof of Lemma 3

Recall the statement of Lemma 3: for every rCTL formula φ , there exists an rATL formula φ^* such that for every Kripke structure \mathcal{K} and every state s of \mathcal{K} we have $V(s, \varphi^*) = V_{rCTL}(s, \varphi)$.

We proceed with the proof of Lemma 3 in three steps, as described in Subsection 3.3.

Step 1 First, we transform the Kripke structure $\mathcal{K} = (S, I, R, L)$ into an “equivalent” concurrent game structure $\mathcal{S}_{\mathcal{K}}$ with a single-agent a . This transformation is required because rCTL formulas are evaluated over Kripke structures, whereas rATL formulas are evaluated over concurrent game structures. Formally, we define $\mathcal{S}_{\mathcal{K}} = (St, Ag, Ac, \delta, \ell)$ with

- $St = S$;
- $Ag = \{a\}$;
- $Ac = S$;
- $\delta(s, s') = \begin{cases} s' & \text{if } (s, s') \in R; \\ s'' & \text{for some } (s, s'') \in R \text{ if } (s, s') \notin R; \end{cases}$
- $\ell(s) = L(s)$ for each $s \in S$.

Note that actions in $\mathcal{S}_{\mathcal{K}}$ are the states of the Kripke structure, and δ is well defined because \mathcal{K} has no dead ends. Moreover, the transition function δ mimics the transition relation R , except that it (potentially) contains additional parallel edges to make the function complete. Overall, it is not hard to verify that \mathcal{K} and $\mathcal{S}_{\mathcal{K}}$ have the same set of paths.

Step 2 Second, given an rCTL formula φ , we construct the rATL formula φ^* . To this end, we use a mapping t' that replaces the path quantifiers \exists and \forall with the corresponding strategy quantifiers $\langle\langle\{a\}\rangle\rangle$ and $\langle\langle\emptyset\rangle\rangle$, respectively. For state formulas, we define

$$\begin{aligned}
 t'(p) &:= p, \\
 t'(\neg\varphi) &:= \neg t'(\varphi), \\
 t'(\varphi_1 \vee \varphi_2) &:= t'(\varphi_1) \vee t'(\varphi_2) \\
 t'(\varphi_1 \wedge \varphi_2) &:= t'(\varphi_1) \wedge t'(\varphi_2) \\
 t'(\varphi_1 \rightarrow \varphi_2) &:= t'(\varphi_1) \rightarrow t'(\varphi_2), \\
 t'(\exists\Phi) &:= \langle\langle\{a\}\rangle\rangle t'(\Phi), \\
 t'(\forall\Phi) &:= \langle\langle\emptyset\rangle\rangle t'(\Phi).
 \end{aligned}$$

For path formulas, we define

$$\begin{aligned}
 t'(\odot\varphi) &:= \odot t'(\varphi), \\
 t'(\diamond\varphi) &:= \diamond t'(\varphi), \\
 t'(\square\varphi) &:= \square t'(\varphi).
 \end{aligned}$$

Given an rCTL formula φ , we then simply set $\varphi^* := t'(\varphi)$.

Step 3 It is left to show that $V(s, \varphi^*) = V_{rCTL}(s, \varphi)$ holds for all states s . To this end, we show the following, slightly stronger statement.

Lemma 6. *For every rCTL state formula φ , every rCTL path formula Φ , every Kripke structure \mathcal{K} , every state s of \mathcal{K} , and every path π of \mathcal{K} , the transformed formulas $t'(\varphi)$ and $t'(\Phi)$ satisfy the following in the concurrent game structure $\mathcal{S}_{\mathcal{K}}$:*

1. $V(s, t'(\varphi)) = V_{rCTL}(s, \varphi)$; and
2. $V(\pi, t'(\Phi)) = V_{rCTL}(\pi, \Phi)$.

Proof of Lemma 6. The base case (i.e., atomic propositions) and the induction step for the Boolean operators ($\neg, \vee, \wedge, \rightarrow$), and the robust temporal operators (\odot, \diamond, \square) follow from applying the semantics of rCTL and rATL as expected. Hence, we skip them here and only investigate the remaining two cases for the existential and universal path quantifiers:

- Let $\varphi = \exists\Phi$ and, therefore, $t'(\varphi) = \langle\langle\{a\}\rangle\rangle t'(\Phi)$.

Recall that the paths of \mathcal{K} and $\mathcal{S}_{\mathcal{K}}$ are identical and $\mathcal{S}_{\mathcal{K}}$ is a concurrent game structure with a single-agent. Thus, for every path $\pi \in \text{paths}(s)$, there exists a strategy f_π such that $\text{out}(s, \{f_\pi\}) = \{\pi\}$. Conversely, every strategy f_a of agent a produces exactly one path $\pi_{f_a} \in \text{paths}(s)$.

Next, we apply the induction hypothesis and obtain $V(\pi, t'(\Phi)) = V_{rCTL}(\pi, \Phi)$ for all paths $\pi \in \text{paths}(s)$. Thus,

$$\max_{f_a \in F_a} V(\pi_{f_a}, t'(\Phi)) = b = \max_{\pi \in \text{paths}(s)} V_{rCTL}(\pi, \Phi),$$

where F_a is the set of all strategies of agent a starting in state s . This observation implies $V(s, \langle\langle\{a\}\rangle\rangle t'(\Phi)) = b$ by definition of the $\langle\langle \rangle\rangle$ -operator (recall that $\text{out}(s, f_a)$ is a singleton set of all strategies $f_a \in F_a$). On the other hand, we have $V_{rCTL}(s, \exists\Phi) = b$ by definition of the \exists -operator. In total, we obtain $V(s, \langle\langle\{a\}\rangle\rangle t'(\Phi)) = V_{rCTL}(s, \exists\Phi)$, proving the claim.

- Let $\varphi = \forall\Phi$ and, therefore, $t'(\varphi) = \langle\langle\emptyset\rangle\rangle t'(\Phi)$.

Again, recall now that the paths of \mathcal{K} and $\mathcal{S}_{\mathcal{K}}$ are identical. Moreover, note that $\text{out}(s, \emptyset) = \text{paths}(s)$ because \mathcal{S} is a concurrent game structure with a single agent.

Next, we apply the induction hypothesis and obtain $V(\pi, t'(\Phi)) = V_{rCTL}(\pi, \Phi)$ for all paths $\pi \in \text{paths}(s)$, as before. Thus,

$$\min_{\pi \in \text{out}(s, \emptyset)} V(\pi, t'(\Phi)) = b = \min_{\pi \in \text{paths}(s)} V_{rCTL}(\pi, \Phi).$$

This observation implies $V(s, \langle\langle\emptyset\rangle\rangle t'(\Phi)) = b$ by definition of the $\langle\langle \rangle\rangle$ -operator (recall that $\text{out}(s, f_a) = \text{paths}(s)$). On the other hand, we have $V_{rCTL}(s, \forall\Phi) = b$ by definition of the \forall -operator. In total, we obtain $V(s, \langle\langle\emptyset\rangle\rangle t'(\Phi)) = V_{rCTL}(s, \forall\Phi)$, proving the claim.

This concludes the proof. \square

Lemma 3 now follows immediately from Item 1 of Lemma 6 and the fact $\varphi^* = t'(\varphi)$.

E Proof of Lemma 4

Recall that we want to prove that for every rATL* formula φ and every truth value $t \in \mathbb{B}_4$, there is an ATL* formula φ_t such that $V(s, \varphi) \succeq t$ if and only if $\mathcal{S}, s \models \varphi_t$. Furthermore, we will show function mapping φ and t to φ_t is polynomial-time computable.

Proof of Lemma 4. Before we show the inductive construction of φ_t , let us remark that we can restrict ourselves to $t \succ 0000$, as $V(s, \varphi) \succeq 0000$ is true for every state s and every formula φ . Hence, we can pick $\varphi_{0000} = p \vee \neg p$ for some atomic proposition p .

In the following, we assume $t \succ 0000$ and define

- $p_t = p$,
- $(\neg\varphi)_t = \neg\varphi_t$,
- $(\varphi_1 \vee \varphi_2)_t = (\varphi_1)_t \vee (\varphi_2)_t$,
- $(\varphi_1 \wedge \varphi_2)_t = (\varphi_1)_t \wedge (\varphi_2)_t$,
- $(\varphi_1 \rightarrow \varphi_2)_{1111} = \bigwedge_{t \succeq 0000} (\varphi_2)_t \vee \neg(\varphi_1)_t$ and
- $(\varphi_1 \rightarrow \varphi_2)_t = (\varphi_1 \rightarrow \varphi_2)_{1111} \vee (\varphi_2)_t$ for $t \prec 1111$,
- $(\langle\langle A \rangle\rangle\Phi)_t = \langle\langle A \rangle\rangle\Phi_t$, and
- $(\llbracket A \rrbracket\Phi)_t = \llbracket A \rrbracket\Phi_t$.

For Boolean combinations of path formulas, the translation is defined analogously as for state formulas while for temporal operators, the translation is defined as

- $(\odot\Phi)_t = \odot\Phi_t$,
- $(\diamond\Phi)_t = \diamond\Phi_t$, and
- $(\square\Phi)_{1111} = \square\Phi_{1111}$,
- $(\square\Phi)_{0111} = \diamond\square\Phi_{0111}$,
- $(\square\Phi)_{0011} = \square\diamond\Phi_{0011}$, and
- $(\square\Phi)_{0001} = \diamond\Phi_{0001}$.

An induction over the construction of φ shows that φ_t has the desired properties. \square

F Proof of Theorem 5

Proof of Theorem 5. The upper bounds follow directly from Lemma 4 and that fact that ATL* model-checking and satisfiability are in 2EXPTIME [5,49], while the lower bounds follow from the fact that ATL* (and thus CTL*) is a fragment of rATL*: Generalizing the proof of Lemma 2, given an ATL* formula φ , we eliminate all implications, push all negations to the atomic propositions, and then dot every operator, thereby obtaining an rATL* formula φ^* such that $\mathcal{S}, s \models \varphi$ if and only if $V(s, \varphi^*) = 1111$.

Now, as the CTL* satisfiability problem is 2EXPTIME-complete (even for the fragment without until and release [42]), the rATL* satisfiability problem is 2EXPTIME-hard. To obtain 2EXPTIME-hardness of the model-checking problem

we rely on the reduction from the LTL realizability problem to the ATL* model-checking problem [5] and the fact that LTL realizability is 2EXPTIME-hard, even for the fragment without until and release [6]. \square