

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier 10.1109/ACCESS.2017.DOI

Making Agents' Abilities Explicit

YEDI ZHANG^{1,2,3}, FU SONG¹, TAOLUE CHEN⁴

¹School of Information Science and Technology, ShanghaiTech University, Pudong, Shanghai, China

²Shanghai Institute of Microsystem and Information Technology, Chinese Academy of Sciences, Shanghai, China

³University of Chinese Academy of Sciences, Beijing, China

⁴Department of Computer Science and Information Systems, Birkbeck, University of London, London, UK

Corresponding author: FU SONG (e-mail: songfu@shanghaitech.edu.cn).

Y. Zhang and F. Song were partially supported by the National Natural Science Foundation of China (NSFC) grants (Nos.: 61532019 and 61761136011), and T. Chen was partially supported by UK EPSRC grant (No. EP/P00430X/1), NSFC grant (No. 61662035), Guangdong Science and Technology Department grant (No. 2018B010107004) and Birkbeck BEI School Project (ARTEFACT).

ABSTRACT Alternating-time temporal logics (ATL/ATL^{*}) represent a family of modal and temporal logics for reasoning about strategic abilities of agents in multiagent systems. These logics are usually interpreted over concurrent game structures (CGSs), and their interpretations may vary depending on the abilities of agents, such as perfect versus imperfect information, perfect versus imperfect recall. These different abilities lead to a variety of variants that have been studied extensively in literature. However, all of these variants are defined at the semantic level, which may restrict modeling flexibility, or even give counter-intuitive interpretations. For example, an agent may have different abilities when achieving two different goals on the same CGS. To mitigate these issues, in this paper we propose to extend CGSs with agents' abilities, resulting in Abilities Augmented CGSs where concrete abilities can be defined at the syntactic level. We study ATL/ATL^{*} over this new model. We give formal definitions of the new semantics and present model-checking algorithms for ATL/ATL^{*}. We also identify the computational complexity of ATL/ATL^{*} model checking problem, i.e., Δ_3^P -2EXPTIME-complete. The model-checking algorithms are implemented in a prototype tool. Experimental results show the practical feasibility and effectiveness of our approach.

INDEX TERMS Model-checking, Multi-agent systems, Alternating-time temporal logics, Agents' abilities

I. INTRODUCTION

MULTIAGENT systems (MASs) comprising multiple autonomous agents have become a widely adopted paradigm of intelligent systems. Game-based models and associated logics, as the foundation of MASs, have received tremendous attentions in recent years. The seminal work [1] proposed *concurrent game structures* (CGSs) as the model of MASs and alternating-time temporal logics (typically ATL and ATL^{*}) as specification languages for expressing temporal goals. In a nutshell, a CGS consists of multiple players which are used to represent autonomous agents, components and the environment. The model describes how the MAS evolves according to the collective behavior of agents. ATL/ATL^{*}, an extension of the Computational Tree Logic (CTL/CTL^{*}), features coalition modalities $\langle\langle A \rangle\rangle$, each of which is parameterized with a set of agents A . The formula $\langle\langle A \rangle\rangle\varphi$ expresses the property that the coalition A has a collective strategy to achieve a certain goal specified by φ .

A series of extensions of ATL-like logics have been studied which take different agents' abilities into account. These abilities typically include whether agents can identify the current

state of the system completely or only partially (perfect vs. imperfect information), and whether agents can memorize the whole history of observations or simply part of them (perfect vs. imperfect recall). Different abilities usually induce distinct semantics, which are indeed necessary because of the versatility of problem domains. These semantic variants and their model-checking problems comprise subjects of active research for almost two decades, to cite a few [2]–[7].

While agents' abilities play a prominent role [8], the semantics of ATL-like logics only refers to them *implicitly*. In other words, the logic per se does not specify what ability an agent has; instead one could infer the ability an agent requires by examining the specification expressed in the logic. This approach, being elegant and valuable to understand the relationship between different abilities, suffers from a few shortcomings: (1) From the modeling perspective, it is common in practice that agents in a MAS vary in their abilities (for instance, agents modeling sensors may not identify the complete state of the system so can only use strategies with imperfect information). When constructing a model, these abilities ought to be encoded explicitly. Such modeling flex-

ibility is not supported by the existing formalisms. (2) From the semantic perspective, ATL-like logics may exhibit some counter-intuitive semantics. Using the core modality $\langle\langle A \rangle\rangle$ of ATL, the formula $\langle\langle A \rangle\rangle\varphi$, is interpreted as that the coalition A has a collective strategy to achieve the goal φ “no matter what the other agents do” rather than “no matter which strategies the other agents choose”. The delicate difference suggests that the (multi-player) game nature in the evolution of MASs is not fully captured by ATL. For instance, in the imperfect information/recall setting, only agents that are quantified in $\langle\langle A \rangle\rangle$ are assumed to use imperfect information/recall strategies, while the other agents *not* in $\langle\langle A \rangle\rangle$ may still use perfect information and perfect recall strategies. Even worse, if the coalition modalities are nested, the same agent may have different abilities to fulfill the objectives specified in different subformulae, resulting in inconsistency in the strategies it uses. This phenomenon has also been mentioned, e.g., in [9], which proposed a strategic logic making explicit references to strategies of all agents (including those not in $\langle\langle A \rangle\rangle$), though all agents should have same abilities therein.

To summarize, it occurs to us that the current approach in which temporal formulae are with implicit agents’ abilities at the semantic level impedes necessary modeling flexibilities and often yields unpleasant (even weird) semantics. Instead, we argue that coupling agents’ abilities at the syntactic level of system models would deliver a potentially better approach to overcome the aforementioned limitations. Bearing the rationale in mind, we propose a new MAS model: *Abilities Augmented Concurrent Game Structures* (ACGSs), which encompass agents’ abilities explicitly.

We investigate ATL and ATL^* over ACGSs. We give formal definitions of the new semantics and show that in general the new semantics of ATL/ATL^* over ACGSs is incomparable with others even if the underlying CGSs models are the same. We also study the model-checking problem of ATL/ATL^* over ACGSs. We show that this problem is generally undecidable. However, we manage to show that the model-checking problem for ATL^* (resp. ATL) on ACGSs is 2EXPTIME -complete (resp. Δ_3^P -complete) when the imperfect information and perfect recall strategies are disallowed. We implement our algorithms in a prototype tool MCMAS-ACGS¹ and conduct experiments on some standard applications from the literature. The results confirm the feasibility and effectiveness of our approach.

Organization. The rest of the paper is organized as follows. Section II and Section III recap CGSs and ATL/ATL^* . Section IV introduces ACGSs on which the semantics of ATL/ATL^* are revised. Section V discusses the effects of strategy types. Section VI gives the undecidable results of the ATL/ATL^* model-checking problem on ACGSs. Section VII and Section VIII respectively study the model-checking problem of ATL and ATL^* on ACGSs by disallowing imperfect information and perfect recall strategies. Section IX reports experimental results. Section X discusses

related work. Section XI concludes with a summary and future work.

II. CONCURRENT GAME STRUCTURES

We fix a finite set AP of atomic propositions. Given an infinite word $\rho = s_0s_1\cdots$, we denote by ρ_j the symbol s_j , by $\rho_{[0..j]}$ the prefix $s_0s_1\cdots s_j$, and by $\rho_{[j..\infty]}$ the suffix $s_js_{j+1}\cdots$. Similarly, for a finite word $\rho = s_0s_1\cdots s_m$, we denote by ρ_j the symbol s_j for $0 \leq j \leq m$, and by $\text{1st}(\rho)$ the symbol s_m .

A concurrent game structure (CGS) \mathcal{G} is a tuple

$$\mathcal{G} \triangleq (S, S_0, Ag, (Ac_i)_{i \in Ag}, (\sim_i)_{i \in Ag}, (P_i)_{i \in Ag}, \Delta, \lambda),$$

where

- S is a finite set of *states*;
- $S_0 \subseteq S$ is a set of designated *initial states*;
- $Ag = \{1, \dots, n\}$ is a finite set of *agents*;
- Ac_i for $i \in Ag$ is a finite set of *local actions* of agent i ;
- $\sim_i \subseteq S \times S$ for $i \in Ag$ is an *epistemic accessibility relation* (i.e., an equivalence relation), which is used to characterize observable abilities of agent i , namely, agent i cannot distinguish equivalent states;
- $P_i : S \rightarrow 2^{Ac_i}$ is a *protocol function*, which specifies the set of available local actions of agent i at the each state. We assume that $P_i(s) = P_i(s')$ for every $s \sim_i s'$, as agent i should have the same available local actions at two indistinguishable states;
- $\Delta : S \times Ac \rightarrow S$ is a *transition function* in which $Ac = \prod_{i \in Ag} Ac_i$ is a set of joint actions;
- $\lambda : S \rightarrow 2^{AP}$ is a *labeling function* which assigns each state a set of atomic propositions.

Given a joint action $\vec{a} = \langle a_1, \dots, a_n \rangle \in Ac$, we use $\vec{a}(i)$ to denote the local action of agent i in \vec{a} . For each state $s \in S$, a joint action \vec{a} uniquely determines the successor state $\Delta(s, \vec{a})$ of s . A *path* is an infinite sequence $\rho = s_0s_1\cdots$ of states such that for every $j \geq 0$, $s_{j+1} = \Delta(s_j, \vec{a}_j)$ for some joint action $\vec{a}_j \in \prod_{i \in Ag} P_i(s_j)$. A path ρ yields a *trace* $\tau(\rho) = \alpha_0\alpha_1\cdots$ over the alphabet 2^{AP} , where for every $j \geq 0$, $\alpha_j = \lambda(\rho_j)$. Two finite sequences $\rho = s_0 \dots s_m \in S^+$ and $\rho' = s'_0 \dots s'_m \in S^+$ are *indistinguishable* for agent i , denoted by $\rho \sim_i \rho'$, if for every $j : 0 \leq j \leq m$, $s_j \sim_i s'_j$.

A. STRATEGIES

A strategy of an agent $i \in Ag$ specifies what the agent i plans to do at each state. Typical agents’ abilities are captured by the following types of strategies [2]. For $i \in Ag$,

- **Ir-strategy** $\theta_i : S \rightarrow Ac$ such that for every $s \in S$, $\theta_i(s) \in P_i(s)$, i.e., the local action chosen by agent i depends only on the current state of the system.
- **IR-strategy** $\theta_i : S^+ \rightarrow Ac$ such that for every finite sequence $\rho \in S^+$, $\theta_i(\rho) \in P_i(\text{1st}(\rho))$, i.e., the local action chosen by agent i depends on the whole history of the game so far, instead of only the last state.
- **ir-strategy** $\theta_i : S \rightarrow Ac$, the same as Ir-strategies but with the additional constraint

$$s \sim_i s' \Rightarrow \theta_i(s) = \theta_i(s'),$$

¹Available at <https://github.com/MCMAS-ACGS>.

i.e., agent i has to choose the same local action at the states that are indistinguishable from each other by the agent i ;

- **iR-strategy** $\theta_i : S^+ \rightarrow Ac$, the same as IR-strategies but with the additional constraint

$$\rho \sim_i \rho' \Rightarrow \theta_i(\rho) = \theta_i(\rho'),$$

i.e., agent i has to choose the same local action on the finite paths that are indistinguishable from each other by the agent i .

Intuitively, **i** (resp. **I**) signals that agents can only observe partial information characterized via epistemic accessibility relations \sim_i (resp. complete information with all epistemic accessibility relations being the identity relation). **r** (resp. **R**) signals that agents can make decisions based on the current observation (resp. the whole history of observations). For instance, **Ir** stands for perfect information imperfect recall strategies, while **iR** stands for imperfect information perfect recall strategies. We will, by slightly abusing notation, extend both **Ir**-strategies and **iR**-strategies to the domain S^+ such that for all $\rho \in S^+$, $\theta_i(\rho) = \theta_i(\text{1st}(\rho))$. We denote by T_{str} the set of strategy types $\{\text{Ir}, \text{IR}, \text{ir}, \text{iR}\}$. For each strategy type $\sigma \in T_{\text{str}}$, we denote by Θ_i^σ the set of σ -strategies for agent $i \in Ag$ and by Θ_A^σ the set $\bigcup_{i \in A} \Theta_i^\sigma$, for a coalition $A \subseteq Ag$.

B. OUTCOMES

Given a set of agents $A \subseteq Ag$, a *collective σ -strategy* for the coalition A is a function $v_A^\sigma : A \rightarrow \Theta_A^\sigma$ such that for each agent $i \in A$, $v_A^\sigma(i) \in \Theta_i^\sigma$ is a σ -strategy of agent i . For $i \in A$ and $\rho \in S^+$, we denote the local action $v_A^\sigma(i)(\rho)$ of agent i by $v_A^\sigma(i, \rho)$, and the complementary set $Ag \setminus A$ by \bar{A} .

Given a state s , a collective σ -strategy v_A^σ and a collective σ' -strategy $v_{\bar{A}}^{\sigma'}$, let $\text{play}(s, v_A^\sigma, v_{\bar{A}}^{\sigma'})$ denote the path such that $\rho_0 = s$ and for every $j \geq 0$, $\rho_{j+1} = \Delta(\rho_j, \vec{d}_j)$ for some $\vec{d}_j \in Ac$ such that for every $i \in Ag$:

$$\vec{d}_j(i) = \begin{cases} v_A^\sigma(i, \rho_{[0..j]}), & \text{if } i \in A; \\ v_{\bar{A}}^{\sigma'}(i, \rho_{[0..j]}), & \text{if } i \in \bar{A}. \end{cases}$$

Intuitively, $\text{play}(s, v_A^\sigma, v_{\bar{A}}^{\sigma'})$ is the unique play when the CGS starts from the state s and all the agents enforce strategies specified by v_A^σ and $v_{\bar{A}}^{\sigma'}$.

For every state $s \in S$ and collective σ -strategy v_A^σ of the coalition A , the *outcome* of the CGS \mathcal{G} is defined as follows:

$$O_{\mathcal{G}}^\sigma(s, v_A^\sigma) \triangleq \{\text{play}(s, v_A^\sigma, v_{\bar{A}}^{\text{IR}}) \mid \forall i \in \bar{A}, v_{\bar{A}}^{\text{IR}}(i) \in \Theta_i^{\text{IR}}\}.$$

Intuitively, $O_{\mathcal{G}}^\sigma(s, v_A^\sigma)$ is the set of all the possible plays that may occur when each agent $i \in A$ enforces its σ -strategy $v_A^\sigma(i)$ from the state s no matter which IR-strategies the other agents choose. The subscript \mathcal{G} is dropped from $O_{\mathcal{G}}^\sigma$ when it is clear from the context.

III. ALTERNATING-TIME TEMPORAL LOGICS

The alternating-time temporal logics: ATL and ATL* are respectively extensions of the branching-time logics CTL and CTL* by replacing the existential path quantifier **E** with coalition modalities $\langle\langle A \rangle\rangle$ [1], each of which is parameterized by

a coalition $A \subseteq Ag$. Intuitively, the formula $\langle\langle A \rangle\rangle\phi$ expresses that the coalition A has a collective strategy to achieve the goal ϕ no matter which strategies the agents in \bar{A} choose. Formally, ATL* is defined by the following grammar:

State formulae $\varphi ::= q \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\phi,$

Path formulae $\phi ::= \varphi \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U}\phi,$

where $q \in AP$ and $A \subseteq Ag$.

The derived operators are defined as usual:

$$\phi_1 \vee \phi_2 \triangleq \neg(\neg\phi_1 \wedge \neg\phi_2) \quad \mathbf{F}\phi \triangleq \text{true } \mathbf{U}\phi$$

$$\phi_1 \rightarrow \phi_2 \triangleq \phi_2 \vee \neg\phi_1 \quad \mathbf{G}\phi \triangleq \neg\mathbf{F}\neg\phi$$

$$[[A]]\phi \triangleq \neg\langle\langle A \rangle\rangle\neg\phi \quad \phi_1 \mathbf{R}\phi_2 \triangleq \mathbf{G}\phi_2 \vee \phi_2 \mathbf{U}(\phi_1 \wedge \phi_2)$$

In this work, ATL* formulae refer to ATL* state formulae. A path formula of ATL* with the state formulae being restricted to atomic propositions is called an LTL formula. Formally, LTL is defined by the following grammar:

$$\phi ::= q \mid \neg\phi \mid \phi \wedge \phi \mid \mathbf{X}\phi \mid \phi \mathbf{U}\phi.$$

The semantics of ATL* is traditionally defined over CGSs. When strategy abilities are considered, it is often parameterized with a strategy type $\sigma \in T_{\text{str}}$, denoted by ATL_σ^* [8]. Formally, let \mathcal{G} be a CGS and s be a state of \mathcal{G} , the semantics of ATL_σ^* (i.e., the satisfaction relation \models_σ) is defined inductively as follows:

Semantics of state formulae

- $\mathcal{G}, s \models_\sigma q$ iff $q \in \lambda(s)$;
- $\mathcal{G}, s \models_\sigma \neg\varphi$ iff $\mathcal{G}, s \not\models_\sigma \varphi$;
- $\mathcal{G}, s \models_\sigma \varphi_1 \wedge \varphi_2$ iff $\mathcal{G}, s \models_\sigma \varphi_1$ and $\mathcal{G}, s \models_\sigma \varphi_2$;
- $\mathcal{G}, s \models_\sigma \langle\langle A \rangle\rangle\phi$ iff there exists a collective σ -strategy v_A^σ for the coalition A such that for each path $\rho \in O^\sigma(s, v_A^\sigma)$: $\mathcal{G}, \rho \models_\sigma \phi$;

Semantics of path formulae

- $\mathcal{G}, \rho \models_\sigma \varphi$ iff $\mathcal{G}, \rho_0 \models_\sigma \varphi$;
- $\mathcal{G}, \rho \models_\sigma \neg\phi$ iff $\mathcal{G}, \rho \not\models_\sigma \phi$;
- $\mathcal{G}, \rho \models_\sigma \phi_1 \wedge \phi_2$ iff $\mathcal{G}, \rho \models_\sigma \phi_1$ and $\mathcal{G}, \rho \models_\sigma \phi_2$;
- $\mathcal{G}, \rho \models_\sigma \mathbf{X}\phi$ iff $\mathcal{G}, \rho_{[1..\infty]} \models_\sigma \phi$;
- $\mathcal{G}, \rho \models_\sigma \phi_1 \mathbf{U}\phi_2$ iff there exists an integer $k \geq 0$ such that $\mathcal{G}, \rho_{[k..\infty]} \models_\sigma \phi_2$ and for all $j : 0 \leq j < k$: $\mathcal{G}, \rho_{[j..\infty]} \models_\sigma \phi_1$.

Given an ATL* formula φ , a CGS \mathcal{G} and a strategy type $\sigma \in T_{\text{str}}$, the *model-checking problem* is to determine whether $\mathcal{G}, s \models_\sigma \varphi$ or not, for each initial state s of the CGS \mathcal{G} .

A. VANILLA ATL

(Vanilla) ATL is a sublogic of ATL* where each occurrence of the coalition modality $\langle\langle A \rangle\rangle$ is immediately followed by a temporal operator. Formally, ATL is defined by the following grammar:

$$\varphi ::= q \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle\mathbf{X}\varphi \mid \langle\langle A \rangle\rangle(\varphi \mathbf{R}\varphi) \mid \langle\langle A \rangle\rangle(\varphi \mathbf{U}\varphi),$$

where $q \in AP$ and $A \subseteq Ag$.

Remark that the release operator **R** cannot be defined using the until (**U**) and global (**G**) operators in ATL [10], so is included for completeness.

An ATL/ATL* formula of the form $\langle\langle A \rangle\rangle\phi$ is *simple* if ϕ is an LTL formula. An ATL/ATL* formula φ is *positive* if (1) each subformula $\langle\langle A \rangle\rangle\phi$ in φ is a simple formula, (2) there is no occurrence of $[[A]]\phi$ in φ , and (3) negations \neg only appear in front of atomic propositions. For example, $\langle\langle A \rangle\rangle X q$ is simple and positive, but $\neg\langle\langle A \rangle\rangle X q$ is neither simple nor positive.

B. SOME SEMANTICS ISSUES

We observe that the semantics of ATL/ATL* refers to the agents' abilities in an implicit manner. For the formula $\langle\langle A \rangle\rangle\varphi$, the specified σ -strategies only apply to agents in the coalition A while the agents in the coalition \bar{A} (i.e., outside of A) could still choose beyond σ -strategies (e.g., IR-strategies). In other words, the coalition A has a collective σ -strategy to achieve φ no matter what the other agents do. When σ is IR as in the original work by [1], this interpretation of $\langle\langle A \rangle\rangle\varphi$ is plausible, as “no matter what the other agents do” is effectively the same as “no matter which strategies the other agents choose”. However, when σ is set to be more restricted than IR, agents not in the coalition A are still allowed to use IR-strategies.

As mentioned in the introduction, this results in a few shortcomings. From a modeling perspective, agents' abilities should be arguably decided by the practical scenario. Namely, they should be fixed when the model is built, and all agents stick to their respective abilities independent of logic formulae. More concretely, from the semantic perspective, the existing semantics only take into account the abilities of agents that are quantified in $\langle\langle A \rangle\rangle$, but does *not* take into account the abilities of agents who are not in the coalition A , and neglects the (multi-player) game nature in the evolution of MASs. As a result, it may exhibit some counter-intuitive semantics. For instance, consider two formulae $\langle\langle A \rangle\rangle\phi$ and $\langle\langle A' \rangle\rangle\phi'$, the agent $i \in A \setminus A'$ may have different abilities to achieve ϕ and ϕ' .

Let us consider an autonomous road vehicle scenario to see why this is not ideal. There are several autonomous cars which can only observe partial information and have bounded memory. A CGS model \mathcal{G} consists of a set A of agents modeling autonomous cars, and an additional environment agent e . We can reasonably assume that all the car agents use IR-strategies, while e uses IR-strategies. The property $\langle\langle A' \rangle\rangle\phi$ expresses that autonomous cars $A' \subset A$ can cooperatively achieve the goal ϕ no matter which strategies the other cars and the environment choose. Verifying that \mathcal{G} satisfies $\langle\langle A' \rangle\rangle\phi$ under the existing semantics would allow car agents $A \setminus A'$ to use IR-strategies. If \mathcal{G} satisfies $\langle\langle A' \rangle\rangle\phi$, then the result is conclusive, i.e., $\langle\langle A' \rangle\rangle\phi$ holds for the system. However, if \mathcal{G} invalidates $\langle\langle A' \rangle\rangle\phi$, we *cannot* deduce that $\langle\langle A' \rangle\rangle\phi$ fails because we overestimate the abilities of agents in $A \setminus A'$ when evaluating $\langle\langle A' \rangle\rangle\phi$. In other words, for the formula $\langle\langle A' \rangle\rangle\phi$ under \models_σ where $\sigma \neq \text{IR}$, it seems to be inappropriate to render the agents in $A \setminus A'$ extra powers of IR to potentially defeat agents from A' when the abilities of the agents in $A \setminus A'$ are actually much weaker and agents

in A' are certainly aware of this. The over-approximation of strategic abilities in such cases are unnecessary and may not be sufficient.

IV. ABILITIES AUGMENTED ACGSs

In this section, we introduce *abilities augmented concurrent game structures* (ACGSs in short), which explicitly equip each agent with a strategy type from T_{str} . As such, agents have fixed strategic abilities for a given ACGS. Formally, an ACGS is a pair

$$\mathcal{M} = (\mathcal{G}, \pi),$$

where \mathcal{G} is a CGS and $\pi : Ag \rightarrow T_{\text{str}}$ is a function that assigns a strategy type $\pi(i)$ to each agent $i \in Ag$. The strategy type $\pi(i)$ explicitly characterizes the abilities of agent i in the CGS model. Recalling that epistemic accessibility relations are used to characterize observable abilities of agents, agents with IR-strategies or IR-strategies are able to distinguish two distinct states, hence we assume that, for each agent $i \in Ag$ with $\pi(i) \in \{\text{IR}, \text{Ir}\}$, the epistemic accessibility relation \sim_i is an identity relation, denoted by id_i .

Paths and traces of ACGSs are defined in the same way as in CGSs, but strategies and outcomes have to be revised as follows.

A. STRATEGIES AND OUTCOMES OF ACGSs

Let A be a set of agents. A *collective strategy* of the coalition A in the ACGS \mathcal{M} is a function $\xi_A : A \rightarrow \bigcup_{i \in A} \Theta_i^{\pi(i)}$ that assigns each agent $i \in A$ a $\pi(i)$ -strategy $\xi_A(i) \in \Theta_i^{\pi(i)}$.

Given a state $s \in S$ and a set of agents $A \subseteq Ag$, for every collective strategy ξ_A of the coalition A , the *outcome* $O_{\mathcal{M}}(s, \xi_A)$ of \mathcal{M} is the set of all possible paths that may occur when each agent $i \in A$ enforces its $\pi(i)$ -strategy $\xi_A(i)$ from the state s , and any other agent $i \in \bar{A}$ can only choose $\pi(i)$ -strategies (rather than general IR-strategies). Formally, $O_{\mathcal{M}}(s, \xi_A)$ is defined as

$$O_{\mathcal{M}}(s, \xi_A) \triangleq \{\text{play}(s, \xi_A, \xi_{\bar{A}}) \mid \forall i \in \bar{A}, \xi_{\bar{A}}(i) \in \Theta_i^{\pi(i)}\}.$$

We will omit the subscript \mathcal{M} from $O_{\mathcal{M}}(s, \xi_A)$ when it is clear from context.

B. SEMANTICS OF ATL AND ATL* ON ACGSs

The difference of outcomes between ACGSs and CGSs induces distinct semantics of ATL/ATL* on ACGSs than CGSs. Let \mathcal{M} be an ACGS and s be a state in \mathcal{M} , the semantics of ATL/ATL* on \mathcal{M} (i.e., the satisfaction relation \models) is defined similar to the one on a CGS, except that the semantics of the state formulae of the form $\langle\langle A \rangle\rangle\phi$ is defined as follows:

$$\begin{aligned} \mathcal{M}, s \models \langle\langle A \rangle\rangle\phi & \text{ if there exists a collective strategy} \\ \xi_A : A & \rightarrow \bigcup_{i \in A} \Theta_i^{\pi(i)} \text{ for the coalition } A \text{ such that} \\ \mathcal{M}, \rho \models \phi, & \text{ for all paths } \rho \in O(s, \xi_A). \end{aligned}$$

Remark that this semantics takes into account whether the agents from \bar{A} have perfect or imperfect information/recall.

Given an ACGS \mathcal{M} and an ATL/ATL* formula φ , the *model-checking problem* is to determine whether $\mathcal{M}, s \models \varphi$

holds, for every initial state s of \mathcal{M} . Given a state formula φ , let $\llbracket \varphi \rrbracket_{\mathcal{M}}$ denote the set of the states of \mathcal{M} that satisfy φ .

We remark that, as per formal semantics, the system model is syntactic and their computations are semantic. In previous work in literature, the abilities of agents are determined when ATL/ATL* formulae are interpreted using computations of the system model. In contrast, in this work, the abilities of agents are defined explicitly in the system model without referring to ATL/ATL* formulae or computations of the system model. Therefore, we assert that the abilities of agents are defined at the semantic level in previous work, whereas at the syntactic level in this work.

V. EFFECTS OF STRATEGY TYPES

Given an ATL/ATL* formula φ , we denote by Ag_{φ} the set of agents that appear in φ . The semantics of ATL/ATL* defined on ACGSs is different from the one defined on CGSs. In general, they are incomparable.

Proposition 1. *There are an ACGS $\mathcal{M} = (\mathcal{G}, \pi)$, an ATL/ATL* formula $\langle\langle A \rangle\rangle\phi$, and a type $\sigma \in T_{\text{str}}$ such that $\pi(i) = \sigma$ for all $i \in A$ and $\mathcal{M}, s \models \langle\langle A \rangle\rangle\phi$ holds, but $\mathcal{G}, s \not\models_{\sigma} \langle\langle A \rangle\rangle\phi$.*

Proof. Let us consider the CGS shown in Figure 1. There are two agents $\{1, 2\}$, four states $\{s_0, s_1, s_2, s_3\}$ (s_0 is the initial state), $\lambda(s_0) = \lambda(s_1) = \lambda(s_2) = \{q\}$ and $\lambda(s_3) = \emptyset$, \sim_1 is the identity relation, $s \sim_2 s'$ for every $s, s' \in \{s_0, s_1, s_2\}$ and $s_3 \sim_2 s_3$.

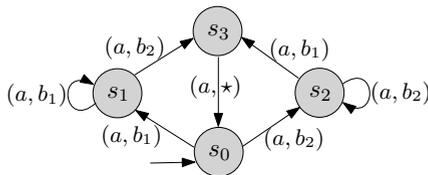


FIGURE 1. An illustrating example, where $\star \in \{b_1, b_2\}$.

Consider the function π such that $\pi(1) = \text{IR}$ and $\pi(2) = \text{ir}$, then $\mathcal{M}, s_0 \models \langle\langle \{1\} \rangle\rangle Gq$, but $\mathcal{G}, s_0 \not\models_{\text{IR}} \langle\langle \{1\} \rangle\rangle Gq$. \square

Proposition 1 reveals that for positive ATL/ATL* formulae φ such that $\pi(i) = \sigma$ for each $i \in Ag_{\varphi}$, even if the agents of Ag_{φ} have the same strategy types in the ACGS (\mathcal{G}, π) and the CGS \mathcal{G} , verifying \mathcal{G} against φ under σ may examine more behavior than verifying (\mathcal{G}, π) against φ . Therefore, if the behavior of a MAS is exactly modeled as an ACGS \mathcal{M} rather than a CGS \mathcal{G} with strategy type σ , verifying \mathcal{G} against φ under σ may lead to incorrect result.

By restricting all the strategy types to IR, straightforwardly we have:

Proposition 2. *Let $\mathcal{M} = (\mathcal{G}, \pi)$ be an ACGS where for each $i \in Ag$, $\pi(i) = \text{IR}$. For each state s of \mathcal{M} and ATL* formula φ , $\mathcal{G}, s \models_{\text{IR}} \varphi$ iff $\mathcal{M}, s \models \varphi$.*

Proof. By applying structural induction on φ , it suffices to show that the result holds for formulae of the form $\langle\langle A \rangle\rangle\phi$.

By the induction hypothesis, for every path ρ , the following holds: $\mathcal{G}, \rho \models_{\text{IR}} \varphi$ iff $\mathcal{M}, \rho \models \varphi$.

For each pair $(\xi_A, \nu_A^{\text{IR}})$ of collective strategies such that $\xi_A = \nu_A^{\text{IR}}$, we have: $\mathcal{O}_{\mathcal{M}}(s, \xi_A) = \mathcal{O}_{\mathcal{G}}^{\text{IR}}(s, \nu_A^{\text{IR}})$. Each agent $i \in A$ has same sets of possible IR-strategies in \mathcal{G} and \mathcal{M} , hence $\mathcal{G}, s \models_{\text{IR}} \langle\langle A \rangle\rangle\phi$ iff $\mathcal{M}, s \models \langle\langle A \rangle\rangle\phi$. \square

In light of Proposition 1 and Proposition 2, in this section we shall investigate the effects of strategy types by considering ACGSs with various different setups of strategy types.

Given a set $A \subseteq Ag$ and two functions $\pi_1, \pi_2 : Ag \rightarrow T_{\text{str}}$, π_1 is *coarser* than π_2 with respect to the coalition A , denoted by $\pi_1 \leq_A \pi_2$, if for every $i \in A$, $\pi_1(i) = \pi_2(i)$ and for every $j \in \bar{A}$, one of the following conditions holds:

- $\pi_1(j) = \text{IR}, \pi_2(j) = \text{IR};$
- $\pi_1(j) = \text{Ir}, \pi_2(j) \in \{\text{IR}, \text{Ir}\};$
- $\pi_1(j) = \text{iR}, \pi_2(j) \in \{\text{IR}, \text{iR}\};$
- $\pi_1(j) = \text{ir}, \pi_2(j) \in \{\text{IR}, \text{Ir}, \text{iR}, \text{ir}\} = T_{\text{str}}.$

Proposition 3. *Let A be a set of agents and s be a state of a CGS \mathcal{G} . For two functions $\pi_1, \pi_2 : Ag \rightarrow T_{\text{str}}$ with $\pi_1 \leq_A \pi_2$, and any collective strategy ξ_A of the coalition A , we have:*

$$\mathcal{O}_{(\mathcal{G}, \pi_1)}(s, \xi_A) \subseteq \mathcal{O}_{(\mathcal{G}, \pi_2)}(s, \xi_A).$$

Proposition 3 reveals the effect of strategy types of \bar{A} on the outcomes. It is easy to observe that if $\pi_2(i) = \sigma$ for all $i \in A$, then for every collective σ -strategy ν_A^{σ} such that $\xi_A = \nu_A^{\sigma}$, we have: $\mathcal{O}_{(\mathcal{G}, \pi_2)}(s, \xi_A) \subseteq \mathcal{O}_{\mathcal{G}}^{\sigma}(s, \nu_A^{\sigma})$. Moreover, if $\pi_2(i) = \text{IR}$ for all $i \in \bar{A}$, then $\mathcal{O}_{(\mathcal{G}, \pi_2)}(s, \xi_A) = \mathcal{O}_{\mathcal{G}}^{\text{IR}}(s, \nu_A^{\text{IR}})$.

By Proposition 3, we have:

Proposition 4. *Given a CGS \mathcal{G} , a state s in \mathcal{G} and a positive ATL/ATL* formula φ , for each pair of two functions $\pi_1, \pi_2 : Ag \rightarrow T_{\text{str}}$ such that $\pi_1 \leq_{Ag_{\varphi}} \pi_2$,*

$$\text{if } (\mathcal{G}, \pi_2), s \models \varphi, \text{ then } (\mathcal{G}, \pi_1), s \models \varphi.$$

Proof. By applying structural induction on φ , it suffices to show that the result holds for formulae of the form $\langle\langle A \rangle\rangle\phi$.

Suppose $(\mathcal{G}, \pi_2), s \models \langle\langle A \rangle\rangle\phi$ (otherwise the proposition immediately holds), then there exists a collective strategy ξ_A for the coalition A such that for each path $\rho \in \mathcal{O}_{(\mathcal{G}, \pi_2)}(s, \xi_A)$, $(\mathcal{G}, \pi_2), \rho \models \phi$ holds. Since $A \subseteq Ag_{\varphi}$ and for every $i \in Ag_{\varphi}$, $\pi_1(i) = \pi_2(i)$ and $\pi_1 \leq_{Ag_{\varphi}} \pi_2$, then $\pi_1 \leq_A \pi_2$. By Proposition 3, we get that $\mathcal{O}_{(\mathcal{G}, \pi_1)}(s, \xi_A) \subseteq \mathcal{O}_{(\mathcal{G}, \pi_2)}(s, \xi_A)$.

By the induction hypothesis, for every state formula ϕ' in ϕ and every state s' , if $(\mathcal{G}, \pi_2), s' \models \phi'$, then $(\mathcal{G}, \pi_1), s' \models \phi'$.

Therefore, for each path $\rho \in \mathcal{O}_{(\mathcal{G}, \pi_1)}(s, \xi_A)$, we get that $(\mathcal{G}, \pi_1), \rho \models \phi$. The result immediately follows. \square

More restrictions on strategy types and ATL/ATL* formulae can make two semantics coincide, as the following proposition shows.

Proposition 5. *Let s be a state of the ACGS $\mathcal{M} = (\mathcal{G}, \pi)$ and $\sigma \in T_{\text{str}}$ be a strategy type. Assume an ATL/ATL* formula φ that satisfies*

- 1) *for every $i \in Ag_{\varphi}$, $\pi(i) = \sigma$,*
- 2) *for every $i \in Ag \setminus Ag_{\varphi}$, $\pi(i) = \text{IR}$, and*

3) for every occurrence of $\langle\langle A' \rangle\rangle\phi$ in φ , $Ag_\varphi = A'$.

Then we have $\mathcal{G}, s \models_\sigma \varphi$ iff $\mathcal{M}, s \models \varphi$.

Proof. The proof directly follows from the fact that $O_{(\mathcal{G}, \pi_2)}(s, \xi_{Ag_\varphi}) = O_{\mathcal{G}}^\sigma(s, \nu_{Ag_\varphi}^\sigma)$ for every state $s \in S$, collective strategy ξ_{Ag_φ} and collective σ -strategy $\nu_{Ag_\varphi}^\sigma$ such that $\xi_{Ag_\varphi} = \nu_{Ag_\varphi}^\sigma$. \square

VI. UNDECIDABLE RESULTS

In this section, we present the following undecidable results.

Theorem 1. *The ATL/ATL* model-checking problem for ACGSs is undecidable.*

Proof. It has been shown [5] that the Halting problem of Turing machines can be reduced to the ATL_{iR} model-checking problem of CGSs against the formula $\varphi = \langle\langle\{1, 2\}\rangle\rangle\mathbf{G} \text{ ok}$, where ok is an atomic proposition.

In other words, one can construct a CGS $\mathcal{G} = (S, \{s_0\}, Ag, (Ac_i)_{i \in Ag}, (\sim_i)_{i \in Ag}, (P_i)_{i \in Ag}, \Delta, \lambda)$ from a Turing machine such that $\mathcal{G}, s_0 \models_{iR} \langle\langle\{1, 2\}\rangle\rangle\mathbf{G} \text{ ok}$ iff the Turing machine does not halt on the empty word.

Let $\mathcal{M} = (\mathcal{G}, \pi)$ be an ACGS such that for every agent $i \in Ag$, $\pi(i) = iR$ if $i \in \{1, 2\}$, otherwise $\pi(i) = IR$. Clearly, $\mathcal{G}, s_0 \models_{iR} \langle\langle\{1, 2\}\rangle\rangle\mathbf{G} \text{ ok}$ iff $\mathcal{M}, s_0 \models \langle\langle\{1, 2\}\rangle\rangle\mathbf{G} \text{ ok}$. The undecidability immediately follows. \square

By Theorem 1, in the rest of this paper, we focus on the model-checking problem of ACGSs by restricting the function π to $Ag \rightarrow T_{str} \setminus \{iR\}$.

VII. ATL MODEL-CHECKING FOR ACGSs

In this section, we show that the ATL model-checking problem for ACGSs is Δ_3^P -complete. We first propose a model-checking algorithm and then prove the Δ_3^P -hardness of the problem.

Our model-checking algorithm iteratively computes the set of states satisfying state formulae from the innermost subformulae. The main challenge is how to compute $\llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$. To this end, we first show how to compute $\llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$ for a simple formula $\langle\langle A \rangle\rangle\phi$, and then present the algorithm for general ATL formulae.

A. MODEL-CHECKING FOR SIMPLE ATL

For a simple ATL formula of the form $\langle\langle A \rangle\rangle\phi$, we can show that whether it is satisfied or not is irrelevant to whether the agents that have perfect information abilities admit perfect recall strategies or not.

Proposition 6. *Given a simple ATL formula $\langle\langle A \rangle\rangle\varphi$, consider an ACGS $\mathcal{M} = (\mathcal{G}, \pi)$ such that $\mathcal{G} = (S, \{s\}, Ag, (Ac_i)_{i \in Ag}, (\sim_i)_{i \in Ag}, (P_i)_{i \in Ag}, \Delta, \lambda)$ and $\pi : Ag \rightarrow T_{str} \setminus \{iR\}$, let π' be a function such that for every $i \in Ag$,*

$$\pi'(i) = \begin{cases} iR, & \text{if } i \in A \wedge \pi(i) = IR; \\ \pi(i), & \text{otherwise.} \end{cases}$$

For every state s in \mathcal{M} , the following holds:

$$(\mathcal{G}, \pi), s \models \langle\langle A \rangle\rangle\varphi \text{ iff } (\mathcal{G}, \pi'), s \models \langle\langle A \rangle\rangle\varphi.$$

Proof. Recalling that for each agent $i \in Ag$ with $\pi(i) \in \{iR, IR\}$, \sim_i is an identity relation, we can then safely regard all the agents in Ag with iR -strategies as iR -strategies with the identity relation. We first construct the tree-unfolding \mathcal{M}_s^* of \mathcal{M} from the state s . Let $\mathcal{M}_s^* = (\mathcal{G}^*, \pi^*)$ such that

$$\mathcal{G}^* = (S^+, S_0^*, Ag, (Ac_i)_{i \in Ag}, (\sim_i^*)_{i \in Ag}, (P_i^*)_{i \in Ag}, \Delta^*, \lambda^*),$$

where

- $S_0^* = \{s\}$;
 - for every $i \in Ag$,
- $$\pi^*(i) = \begin{cases} iR, & \text{if } i \in A \wedge \pi(i) = iR; \\ \pi(i), & \text{otherwise.} \end{cases}$$
- for every $i \in Ag$ and $\rho, \rho' \in S^+$, $\rho \sim_i^* \rho'$, if
 - either $\pi(i) \neq iR$ and $\text{lst}(\rho) \sim_i \text{lst}(\rho')$
 - or $\pi(i) = iR$ and $\rho = \rho'$;
 - $P_i^*(\rho) = P_i(\text{lst}(\rho))$ for every $i \in Ag$ and $\rho \in S^+$;
 - $\Delta^*(\rho, \vec{d}) = \rho \cdot \Delta(\text{lst}(\rho), \vec{d})$ for every $\rho \in S^+$ and $\vec{d} \in Ac$;
 - $\lambda^*(\rho) = \lambda(\text{lst}(\rho))$ for every $\rho \in S^+$.

We observe that the tree-unfolding \mathcal{M}_s^* is a tree-like ACGS, namely, every state can be reached by a unique finite path from the state s . Hence, iR -strategies of the coalition A from the state s in \mathcal{M} correspond exactly to iR -strategies of the coalition A from the state s in the tree unfolding \mathcal{M}_s^* , while the types of other agents are same under π and π^* . We show that $\mathcal{M}, s \models \langle\langle A \rangle\rangle\varphi$ iff $\mathcal{M}_s^*, s \models \langle\langle A \rangle\rangle\varphi$. (We remark that this result does not hold if φ is a general LTL formula.)

(\Rightarrow) Suppose $\mathcal{M}, s \models \langle\langle A \rangle\rangle\varphi$, then there exists a collective strategy ξ_A such that for every path $\rho \in O_{\mathcal{M}}(s, \xi_A)$: $\mathcal{M}, \rho \models \varphi$. From the collective strategy ξ_A , we define the function ξ_A^* such that for every $i \in A$ and $\rho \in S^+$:

$$\xi_A^*(i)(\rho) = \begin{cases} \xi_A(i)(\text{lst}(\rho)), & \text{if } \pi(i) \neq iR; \\ \xi_A(i)(\rho), & \text{if } \pi(i) = iR. \end{cases}$$

First, we show that ξ_A^* is a collective strategy of the coalition A in \mathcal{M}_s^* . Consider an agent $i \in A$ and two states $\rho, \rho' \in S^+$, if $\rho \sim_i^* \rho'$, then either $(\pi(i) \neq iR$ and $\text{lst}(\rho) \sim_i \text{lst}(\rho'))$ or $(\rho = \rho'$ and $\pi(i) = iR)$.

- If $\pi(i) \neq iR$ and $\text{lst}(\rho) \sim_i \text{lst}(\rho')$, then we get that $\xi_A(i)(\text{lst}(\rho)) = \xi_A(i)(\text{lst}(\rho'))$, hence $\xi_A^*(i)(\rho) = \xi_A^*(i)(\rho')$.
- If $\rho = \rho'$ and $\pi(i) = iR$, then we get that $\xi_A(i)(\rho) = \xi_A(i)(\rho')$, hence $\xi_A^*(i)(\rho) = \xi_A^*(i)(\rho')$.

Therefore, ξ_A^* is a collective strategy of the coalition A in \mathcal{M}_s^* .

Next, we show that for every collective strategy $\xi_{\bar{A}}^*$ of \bar{A} in \mathcal{M}_s^* , $\text{play}(s, \xi_A^*, \xi_{\bar{A}}^*) \models \varphi$ holds.

Suppose $\text{play}(s, \xi_A^*, \xi_{\bar{A}}^*) = \rho_0 \rho_1 \dots$. Let $\xi_{\bar{A}}$ be the function such that for every $i \in \bar{A}$ and $j \geq 0$,

$$\begin{aligned} \xi_{\bar{A}}(i)(\text{lst}(\rho_j)) &= \xi_{\bar{A}}^*(i)(\rho_j), & \text{if } \pi(i) \neq iR; \\ \xi_{\bar{A}}(i)(\rho_0 \dots \rho_j) &= \xi_{\bar{A}}^*(i)(\rho_0 \dots \rho_j), & \text{if } \pi(i) = iR. \end{aligned}$$

Consider $j, k \geq 0$ such that $\text{lst}(\rho_j) \sim_i \text{lst}(\rho_k)$ for some $i \in \bar{A}$, then either $\pi(i) \neq iR$ or $\pi(i) = iR$.

- If $\pi(i) \neq iR$, then $\rho_j \sim_i^* \rho_k$. This implies that $\xi_A^*(i)(\rho_j) = \xi_A^*(i)(\rho_k)$. Hence $\xi_A^*(i)(\text{lst}(\rho_j)) = \xi_A^*(i)(\text{lst}(\rho_k))$.

- If $\pi(i) = \text{IR}$, then the agent i can choose any action at any state of ρ_j .

Therefore, $\xi_{\bar{A}}$ is a collective strategy of \bar{A} in \mathcal{M} and $\text{play}(s, \xi_A, \xi_{\bar{A}}) = \text{1st}(\rho_0)\text{1st}(\rho_1)\cdots$. Following from the fact that $\lambda^*(\rho) = \lambda(\text{1st}(\rho))$ for every $\rho \in S^+$, we get that $\mathcal{M}_{s^+}^*, s \models \langle\langle A \rangle\rangle\varphi$.

(\Leftarrow) Suppose $\mathcal{M}_{s^+}^*, s \models \langle\langle A \rangle\rangle\varphi$, then there exists a collective strategy ξ_A^* such that for every path $\rho \in \mathcal{O}_{\mathcal{M}_{s^+}^*}(s, \xi_A^*)$: $\mathcal{M}_{s^+}^*, \rho \models \varphi$. Without loss of generality, we assume that there is an arbitrary total order \leq on set S^+ , and denote by $\min(U)$ the minimal one of the set of states $U \subseteq S^+$ with respect to the order \leq .

Let ξ_A be the function such that for every $i \in A$ and $s' \in S$:

$$\xi_A(i)(s') = \xi_A^*(i)(\min(\{\rho \in S^+ \mid \text{1st}(\rho) = s'\})).$$

First, we show that ξ_A is a collective strategy of the coalition A in \mathcal{M} . Consider an agent $i \in A$ and two states $s_1, s_2 \in S$, if $s_1 \sim_i s_2$, then either $\pi(i) \neq \text{IR}$ or ($s_1 = s_2$ and $\pi(i) = \text{IR}$).

- If $\pi(i) \neq \text{IR}$, then for each pair of states $\rho_1, \rho_2 \in S^+$ such that $\text{1st}(\rho_1) = s_1$ and $\text{1st}(\rho_2) = s_2$, we have: $\rho_1 \sim_i^* \rho_2$. This implies that $\xi_A^*(i)(\rho_1) = \xi_A^*(i)(\rho_2)$, hence $\xi_A(i)(s_1) = \xi_A(i)(s_2)$.
- If $s_1 = s_2$ and $\pi(i) = \text{IR}$, we choose $\xi_A(i)(s_1) = \xi_A(i)(s_2) = \xi_A^*(i)(\min(\{\rho \in S^+ \mid \text{1st}(\rho) = s_1\}))$.

Therefore, ξ_A is a collective strategy of the coalition A in \mathcal{M} .

Consider a collective strategy $\xi_{\bar{A}}$ of \bar{A} in \mathcal{M} , let $\rho = \text{play}(s, \xi_A, \xi_{\bar{A}})$, then we have:

$$\rho_{[0..0]}\rho_{[0..1]}\rho_{[0..2]}\cdots \in \mathcal{O}_{\mathcal{M}_{s^+}^*}(s, \xi_A^*).$$

Following from the fact that $\lambda^*(\rho) = \lambda(\text{1st}(\rho))$ for every $\rho \in S^+$, we get that $\mathcal{M}, s \models \langle\langle A \rangle\rangle\varphi$.

Since $\{(\text{1st}(\rho), \rho) \mid \rho \in sS^*\}$ is bisimulation between \mathcal{G} and \mathcal{G}^* (cf. Definition 5.1 and Lemma 5.2 [11]), we get that $(\mathcal{G}^*, \pi^*), s \models \langle\langle A \rangle\rangle\varphi$ iff $(\mathcal{G}, \pi^*), s \models \langle\langle A \rangle\rangle\varphi$. Therefore, we get that $(\mathcal{G}, \pi), s \models \langle\langle A \rangle\rangle\varphi$ iff $(\mathcal{G}^*, \pi^*), s \models \langle\langle A \rangle\rangle\varphi$ iff $(\mathcal{G}, \pi^*), s \models \langle\langle A \rangle\rangle\varphi$. \square

We remark that Alur et al. [1] observed that both semantics of ATL under Ir-strategies and IR-strategies are coincide for CGSs. This result was generalized and formally proved for infinite CGSs (i.e., no finiteness with respect to the set of states and actions) (cf. Proposition 1 [8]). Proposition 6 can be seen as a generalization of the result of [1] and could be extended to the infinite ACGSs similar to [8]. Moreover, Proposition 6 could be generalized to allow all agents that are perfect information to be imperfect recall. We prefer not to do so because it would not improve complexity result, meanwhile it may reduce scalability, as we have to check more strategies of agents not in the coalition A (see the model-checking algorithm below).

By Proposition 6, all the agents in the coalition A with IR-strategies can be seen as with Ir-strategies. Moreover, for each agent $i \in Ag$ with Ir/IR-strategies, \sim_i is an identity relation. Therefore, without loss of generality, we can safely

assume that $\pi(i) = \text{ir}$ for all $i \in A$, and $\pi(i) \in \{\text{ir}, \text{IR}\}$ for all $i \in \bar{A}$. Let \bar{A}_{ir} denotes the set $\{i \in \bar{A} \mid \pi(i) = \text{ir}\}$.

For two collective strategies ξ_A and $\xi_{\bar{A}_{\text{ir}}}$, let

$$\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}}) \triangleq (\mathcal{G}', \pi)$$

be the ACGS obtained from (\mathcal{G}, π) by enforcing strategies ξ_A and $\xi_{\bar{A}_{\text{ir}}}$, namely, by removing transitions whose actions of agents in $A \cup \bar{A}_{\text{ir}}$ do not conform to ξ_A and $\xi_{\bar{A}_{\text{ir}}}$. We have that

Lemma 1. $\llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}} \equiv \bigcup_{\xi_A} \bigcap_{\xi_{\bar{A}_{\text{ir}}}} \llbracket \langle\langle \emptyset \rangle\rangle\phi \rrbracket_{\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})}$.

Proof. (\Rightarrow) Suppose $s \in \llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$, then there exists a collective strategy $\xi_A : A \rightarrow \bigcup_{i \in A} \Theta_i^{\pi(i)}$ such that for each path $\rho \in \mathcal{O}_{\mathcal{M}}(s, \xi_A)$: $\mathcal{M}, \rho \models \phi$.

For every collective strategy $\xi_{\bar{A}_{\text{ir}}} : \bar{A}_{\text{ir}} \rightarrow \bigcup_{i \in \bar{A}_{\text{ir}}} \Theta_i^{\text{ir}}$, we denote by $\text{Paths}_s(\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}}))$ the set of paths in $\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})$ that start from s . Then, $\text{Paths}_s(\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})) \subseteq \mathcal{O}_{\mathcal{M}}(s, \xi_A)$. This implies that for every path $\rho \in \text{Paths}_s(\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}}))$: $\mathcal{M}, \rho \models \phi$ holds. Therefore, we get that $s \in \llbracket \langle\langle \emptyset \rangle\rangle\phi \rrbracket_{\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})}$ for every collective strategy $\xi_{\bar{A}_{\text{ir}}} : \bar{A}_{\text{ir}} \rightarrow \bigcup_{i \in \bar{A}_{\text{ir}}} \Theta_i^{\text{ir}}$. The result immediately follows.

(\Leftarrow) Suppose $s \in \bigcup_{\xi_A} \bigcap_{\xi_{\bar{A}_{\text{ir}}}} \llbracket \langle\langle \emptyset \rangle\rangle\phi \rrbracket_{\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})}$, then there exists a collective strategy $\xi_A : A \rightarrow \bigcup_{i \in A} \Theta_i^{\pi(i)}$ such that $s \in \bigcap_{\xi_{\bar{A}_{\text{ir}}}} \llbracket \langle\langle \emptyset \rangle\rangle\phi \rrbracket_{\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})}$. This implies that for every collective strategy $\xi_{\bar{A}_{\text{ir}}} : \bar{A}_{\text{ir}} \rightarrow \bigcup_{i \in \bar{A}_{\text{ir}}} \Theta_i^{\text{ir}}$, and every path $\rho \in \text{Paths}_s(\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}}))$, $\rho \models \phi$ holds.

Since $\mathcal{O}_{\mathcal{M}}(s, \xi_A) = \bigcup_{\xi_{\bar{A}_{\text{ir}}}} \text{Paths}_s(\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}}))$, we get that for every path $\rho \in \mathcal{O}_{\mathcal{M}}(s, \xi_A)$, $\rho \models \phi$ holds. Therefore, $s \in \llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$. \square

Algorithm for Simple ATL. To compute $\llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$, the Turing machine first existentially guesses a collective strategy $\xi_A : A \rightarrow \bigcup_{i \in A} \Theta_i^{\pi(i)}$ by restricting the transition function of \mathcal{M} . Then the Turing machine reaches the universal state, and explores all collective strategies $\xi_{\bar{A}_{\text{ir}}} : \bar{A}_{\text{ir}} \rightarrow \bigcup_{i \in \bar{A}_{\text{ir}}} \Theta_i^{\text{ir}}$ by restricting the transition function of \mathcal{M} , and finally computes $\llbracket \langle\langle \emptyset \rangle\rangle\phi \rrbracket_{\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})}$ which amounts to CTL model-checking and can be done in polynomial time in the size of $\mathcal{M}(\xi_A, \xi_{\bar{A}_{\text{ir}}})$ and $\langle\langle \emptyset \rangle\rangle\phi$ [12]. Clearly, the number of choices is limited by the size of the transition function and each choice can be doing in polynomial time. Therefore, $\llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$ can be computed in polynomial time by an *alternating* Turing machine with two alternations (starting in an existential state). By the characterization of the polynomial hierarchy (PH), we obtain the following result.

Lemma 2. For a state s and a simple ATL formula $\langle\langle A \rangle\rangle\phi$, checking whether $s \in \llbracket \langle\langle A \rangle\rangle\phi \rrbracket_{\mathcal{M}}$ is in Σ_2^P (i.e., NP^{NP}).

Example 1. Consider the ACGS $\mathcal{M}_r = (\mathcal{G}, \pi)$ defined as follows.

- $\mathcal{G} = (S, \{s_0\}, Ag, (Ac_i)_{i \in Ag}, (\sim_i)_{i \in Ag}, (P_i)_{i \in Ag}, \Delta, \lambda)$, where
 - $S = \{s_0, s_1, s_2, s_3, s_4\}$;
 - $Ag = \{1, 2, 3\}$;
 - $Ac_i = \{a, b\}$ for $i \in Ag$;

Intuitively, the agent g_ψ controls the state s and chooses a clause c to verify by selecting the action a_c . Then, the agent g_d controls the state s_c and chooses a literal ℓ^i (e.g., z or $\neg z$) of c to verify by selecting the action a_i . Next, the agent g_z controls the state s_{ℓ^i} and chooses a truth value for the variable z by selecting a_\top or a_\perp . If the literal ℓ^i is true under z , then \mathcal{M} enters the state q_\top , otherwise \mathcal{M} enters the state q_\perp . The relation \sim_{g_z} ensures that the agent g_z chooses the same truth value at the states s_z and $s_{\neg z}$. The ACGS \mathcal{M} can be constructed in polynomial time in the size of $\exists X.\forall Y.\psi$.

Let Ag_\exists denote the set $\{g_\ell, g_x \mid x \in X\}$, then we have

$$s \in \llbracket \langle \langle Ag_\exists \rangle \rangle \mathbf{F}q_\top \rrbracket_{\mathcal{M}} \text{ iff } \exists X.\forall Y.\psi \text{ is satisfiable.}$$

Indeed, there exists an assignment $f_1 : X \rightarrow \{0, 1\}$ such that ψ evaluates to 1 under f_1 regardless of the values of the variables in Y iff there is a collective strategy ξ_{Ag_\exists} such that $\mathcal{M}, \rho \models \mathbf{F}q_\top$ for all paths $\rho \in \mathcal{O}(s, \xi_{Ag_\exists})$, where for every $x \in X$, $f_1(x) = 1$ (resp. $f_1(x) = 0$) iff the agent g_x selects the action a_\top (resp. a_\perp) at the states s_x and $s_{\neg x}$.

The proof is completed. \square

We remark that the following result also holds:

$$s \in \llbracket \langle \langle \emptyset \rangle \rangle \mathbf{X} \langle \langle Ag_\exists \rangle \rangle \mathbf{F}q_\top \rrbracket_{\mathcal{M}} \text{ iff } \exists X.\forall Y.\psi \text{ is satisfiable.}$$

which will later be used in the proof of Lemma 5.

Following Lemma 2 and Lemma 3, we have that:

Theorem 2. *The model-checking problem for simple ATL formulae is Σ_2^P -complete.*

B. MODEL-CHECKING FOR GENERAL ATL

Algorithm for General ATL. We now present the model-checking algorithm for general ATL, which computes $\llbracket \varphi \rrbracket_{\mathcal{M}}$ from the innermost subformulae.

Algorithm 1: ATL model-checking algorithm

Input: An ACGS $\mathcal{M} = (\mathcal{G}, \pi)$ and an ATL formula φ
Output: $\llbracket \varphi \rrbracket_{\mathcal{M}}$

```

1 Function MC( $\mathcal{M}, \varphi$ )
2   switch  $\varphi$  :
3     case  $q$  : return  $\{s \in S \mid q \in \lambda(s)\}$ ;
4     case  $\neg\varphi'$  : return  $S \setminus \text{MC}(\mathcal{M}, \varphi')$ ;
5     case  $\varphi_1 \wedge \varphi_2$  : return  $\text{MC}(\mathcal{M}, \varphi_1) \cap \text{MC}(\mathcal{M}, \varphi_2)$ ;
6     case  $\mathbf{K}_i\varphi'$  : return  $\{s \in S \mid [s]^{-i} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
7     case  $\mathbf{E}_A\varphi'$  : return  $\{s \in S \mid [s]^{-E_A} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
8     case  $\mathbf{D}_A\varphi'$  : return  $\{s \in S \mid [s]^{-D_A} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
9     case  $\mathbf{C}_A\varphi'$  : return  $\{s \in S \mid [s]^{-C_A} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
10    case  $\langle \langle A \rangle \rangle \phi$  :
11      foreach sub-state-formula  $\varphi'$  in  $\phi$  do
12        Replace  $\varphi'$  by a fresh atomic proposition  $q_{\varphi'}$ 
13        in  $\varphi$ , and let  $\lambda(q_{\varphi'}) := \text{MC}(\mathcal{M}, \varphi')$ ;
14       $\llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}} := \bigcup_{\xi_A} \bigcap_{\xi_{A^c}} \llbracket \langle \langle \emptyset \rangle \rangle \phi \rrbracket_{\mathcal{M}(\xi_A, \xi_{A^c})}$ ;
15    return  $\llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}}$ ;

```

Algorithm 1 shows the pseudo code, which takes an ACGS $\mathcal{M} = (\mathcal{G}, \pi)$ and an ATL formula φ as inputs, and outputs $\llbracket \varphi \rrbracket_{\mathcal{M}}$ which contains all the states that satisfy φ .

We also incorporate epistemic modalities $\mathbf{K}_i\varphi$, $\mathbf{E}_A\varphi$, $\mathbf{D}_A\varphi$ and $\mathbf{C}_A\varphi$ from [13] into our algorithm with the following semantics:

- $\mathcal{G}, s \models_{\sigma} \mathbf{K}_i\varphi$ iff $\forall s' \in S, s \sim_i s' \implies \mathcal{G}, s' \models_{\sigma} \varphi$;
- $\mathcal{G}, s \models_{\sigma} \mathbf{E}_A\varphi$ iff $\forall s' \in S, s \sim_A^E s' \implies \mathcal{G}, s' \models_{\sigma} \varphi$;
- $\mathcal{G}, s \models_{\sigma} \mathbf{D}_A\varphi$ iff $\forall s' \in S, s \sim_A^D s' \implies \mathcal{G}, s' \models_{\sigma} \varphi$;
- $\mathcal{G}, s \models_{\sigma} \mathbf{C}_A\varphi$ iff $\forall s' \in S, s \sim_A^C s' \implies \mathcal{G}, s' \models_{\sigma} \varphi$,

where φ is a state formula, $\sim_A^E = \bigcup_{i \in A} \sim_i$, $\sim_A^D = \bigcap_{i \in A} \sim_i$, $\sim_A^C = (\sim_A^E)^+$ (i.e., the transitive closure of \sim_A^E). $\mathbf{K}_i\varphi$, $\mathbf{E}_A\varphi$, $\mathbf{D}_A\varphi$ and $\mathbf{C}_A\varphi$ denote that “ i knows”, “every agent in the coalition A knows”, “agents in the coalition A have distributed knowledge”, and “agents in the coalition A have common knowledge” on the fact φ , respectively. The ATL logic extended with these epistemic modalities is called ATLK logic. Given a state $s \in S$ and a binary relation $\simeq \subseteq S \times S$, we denote by $[s]^{\simeq}$ the set $\{s' \in S \mid s \simeq s'\}$.

By Lemma 2, the model-checking problem for ATLK on ACGSs is solvable in Δ_3^P (i.e., P^{NPNP}).

Lemma 4. *The model-checking problem for ATLK on ACGSs is in Δ_3^P .*

Δ_3^P -hardness. We show that the model-checking problem for ATL is Δ_3^P -hard by a reduction from the *sequentially nested satisfiability problem of quantified Boolean formulae* (SNSAT_2), which is known to be Δ_3^P -complete [10], [14].

An instance \mathcal{I}_m of SNSAT_2 is given by m Boolean variables $Z = \{z_1, \dots, z_m\}$ and a list of m equations

$$\begin{aligned} z_1 &\doteq \exists X_1.\forall Y_1.\psi_1(X_1, Y_1) \\ z_2 &\doteq \exists X_2.\forall Y_2.\psi_2(X_2, Y_2, z_1) \\ &\vdots \\ z_m &\doteq \exists X_m.\forall Y_m.\psi_m(X_m, Y_m, z_1, \dots, z_{m-1}) \end{aligned}$$

where for every $i : 1 \leq i \leq m$,

- ψ_i is a 3-CNF Boolean formula over variables $X_i \cup Y_i \cup Z_{<i}$ with $Z_{<i} = \{z_1, \dots, z_{i-1}\}$;
- $X_i = \{x_1^i, \dots, x_{m_i}^i\}$, $Y_i = \{y_1^i, \dots, y_{k_i}^i\}$ are two sets of Boolean variables.

The instance \mathcal{I}_m is *satisfiable* iff there exists an assignment $f_m : Z \rightarrow \{0, 1\}$ such that for every $i : 1 \leq i \leq m$,

$$f_m(z_i) = 1 \text{ iff } \exists X_i.\forall Y_i.\psi_i \text{ is satisfiable under } f_m.$$

Lemma 5. *The model-checking problem for ATL on ACGSs is Δ_3^P -hard.*

Proof. We reduce SNSAT_2 to the ATL model-checking problem. For every equation $z_j \doteq \exists X_j.\forall Y_j.\psi_j$, let

$$\mathcal{M}^j = ((S^j, \{s^j\}, Ag^j, (Ac^j)_{i \in Ag^j}, (\sim^j)_{i \in Ag^j}, (P^j)_{i \in Ag^j}, \Delta^j, \lambda^j), \pi^j)$$

be the ACGS constructed as in the proof of Lemma 3 from the formula $\exists X_j.\forall Y_j.\psi_j$. Let $\overline{\mathcal{M}}^j$ be the ACGS obtained from \mathcal{M}^j , where the initial state s^j is renamed to \overline{s}^j . We recursively construct two families of ACGSs: $(\mathcal{N}^j)_{1 \leq j \leq m}$ and $(\overline{\mathcal{N}}^j)_{1 \leq j \leq m}$.

For $j = 0$, let $\mathcal{N}^1 = \mathcal{M}^1$ and $\overline{\mathcal{N}}^1 = \overline{\mathcal{M}}^1$. For $j > 0$, we define \mathcal{N}^j and $\overline{\mathcal{N}}^j$ as follows.

- 1) \mathcal{N}^j (resp. $\overline{\mathcal{N}}^j$) starts from the initial state s_{z_j} (resp. $s_{\neg z_j}$) which is controlled by the agent g_{z_j} with two available actions a_{\perp} and a_{\top} at the state s_{z_j} (resp. $s_{\neg z_j}$).
- 2) At the state s_{z_j} (resp. $s_{\neg z_j}$),
 - a) if the agent g_{z_j} selects the action a_{\top} , then \mathcal{N}^j (resp. $\overline{\mathcal{N}}^j$) goes to the state s^j (the initial state of \mathcal{M}^j) and then behaves the same as \mathcal{M}^j until some state of the form s_{z_i} or $s_{\neg z_i}$ for some $i < j$ is reached;
 - b) if the agent g_{z_j} selects the action a_{\perp} , then \mathcal{N}^j (resp. $\overline{\mathcal{N}}^j$) goes to the state \overline{s}^j (the initial state of $\overline{\mathcal{M}}^j$) and then behaves the same as $\overline{\mathcal{M}}^j$ until some state of the form s_{z_i} or $s_{\neg z_i}$ for some $i < j$ is reached.
- 3) \mathcal{N}^j and $\overline{\mathcal{N}}^j$ behaves the same as \mathcal{N}^i (resp. $\overline{\mathcal{N}}^i$) after the state s_{z_i} (resp. $s_{\neg z_i}$) for some $i < j$.

Each state of the form s^j is associated with the atomic proposition q , i.e., $\lambda(s^j) = \{q\}$. Let Ag_{\exists} denote the set $\{g_{\ell}, g_x \mid x \in \bigcup_{i=1}^m X_i\} \cup \{g_z \mid z \in Z\}$. The ATL formula is constructed recursively as follows: $\phi_0 \equiv \text{false}$ and for every $j : 1 \leq j < m$,

$$\phi_{j+1} \equiv \langle\langle Ag_{\exists} \rangle\rangle \mathbf{X}(q \leftrightarrow \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_j)))$$

where $a \leftrightarrow b$ denotes $(a \wedge b) \vee (\neg a \wedge \neg b)$.

The result follows from the following claim.

Claim:

- 1) $s^m \in \llbracket q \wedge \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{m-1})) \rrbracket_{\mathcal{N}^m}$ iff the instance \mathcal{I}_m is satisfied by some assignment $f_m : Z \rightarrow \{0, 1\}$ such that $f_m(z_m) = 1$.
- 2) $\overline{s}^m \in \llbracket \neg q \wedge \neg \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{m-1})) \rrbracket_{\overline{\mathcal{N}}^m}$ iff the instance \mathcal{I}_m is satisfied by some assignment $f_m : Z \rightarrow \{0, 1\}$ such that $f_m(z_m) = 0$.

If the instance \mathcal{I}_m is satisfied by an assignment $f_m : Z \rightarrow \{0, 1\}$, then we have that: for every $j : 1 \leq j \leq m$, $f_m(z_i) = 1$ iff $\exists X_i. \forall Y_i. \psi_i$ is satisfiable under f_m . It suffices to prove that

$$s^m \in \llbracket \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{m-1})) \rrbracket_{\mathcal{N}^m} \text{ iff } \mathcal{I}_m \text{ is satisfied by some assignment } f_m : Z \rightarrow \{0, 1\} \text{ such that } f_m(z_m) = 1.$$

We prove this by applying induction on m .

Base case $m = 1$. Following the proof of Lemma 3,

$$s^1 \in \llbracket \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle \mathbf{F}q_{\top} \rrbracket_{\mathcal{M}^1} \text{ iff } \exists X_1. \forall Y_1. \psi_1 \text{ is satisfiable.}$$

Then, the result immediately follows from the fact that ϕ_0 is false. Note that q is always false after the state s^1 .

Inductive step $m > 1$. Recall that $s^m \in \llbracket \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{m-1})) \rrbracket_{\mathcal{N}^m}$ iff a collective strategy $\xi : Ag_{\exists} \rightarrow \bigcup_{i \in Ag_{\exists}} \Theta_i^{(i)}$ exists such that for every path $\rho \in \mathcal{O}_{\mathcal{N}^m}(s_c^m, \xi_A)$ and every success state s_c^m of s^m : $\mathcal{N}^m, \rho \models \neg q \mathbf{U}(q_{\top} \vee \phi_{m-1})$.

For every path $\rho \in \mathcal{O}_{\mathcal{N}^m}(s_c^m, \xi_A)$, we have that:

- ρ visits some state ρ_i of the form s^j or \overline{s}^j for $1 \leq j < m$ iff $\rho_i \in \llbracket \phi_j \rrbracket_{\mathcal{N}^m}$,
- ρ does not visit any state ρ_i of the form s^j or \overline{s}^j for $1 \leq j < m$ iff ρ ends with a loop on the state q_{\top} .

By the induction hypothesis, we have that:

- $s^j \in \llbracket \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{j-1})) \rrbracket_{\mathcal{N}^j}$ iff the instance \mathcal{I}_j is satisfied by an assignment $f_j : Z_{<j+1} \rightarrow \{0, 1\}$ such that $f_j(z_j) = 1$.
- $\overline{s}^j \in \llbracket \neg \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{j-1})) \rrbracket_{\overline{\mathcal{N}}^j}$ iff \mathcal{I}_j is satisfied by an assignment $f_j : Z_{<j+1} \rightarrow \{0, 1\}$ such that $f_j(z_j) = 0$.

Therefore, $s^m \in \llbracket \langle\langle \emptyset \rangle\rangle \mathbf{X} \langle\langle Ag_{\exists} \rangle\rangle (\neg q \mathbf{U}(q_{\top} \vee \phi_{m-1})) \rrbracket_{\mathcal{N}^m}$ iff \mathcal{I}_m is satisfied by an assignment $f_m : Z \rightarrow \{0, 1\}$ such that $f_m(z_m) = 1$ and for every $j : 1 \leq j < m$, $f_m(z_j) = f_j(z_j)$.

Note that $s_{z_j} \sim_{g_{z_j}} s_{\neg z_j}$ for every $j : 1 \leq j \leq m$, hence the agent g_{z_j} always chooses the same action at the states s_{z_j} and $s_{\neg z_j}$. This ensures that f_m is well-defined. \square

Following Lemma 4 and Lemma 5, we get that:

Theorem 3. *The model-checking problem for ATL (hence ATLK) formula is Δ_3^P -complete.*

VIII. ATL* MODEL-CHECKING FOR ACGSS

In this section, we show that the ATL* model-checking problem for ACGSS is 2EXPTIME-complete. The model-checking algorithm mainly follows Algorithm 1 which iteratively computes the set of states satisfying state formulae from the innermost subformulae. The main challenge is how to compute $\llbracket \langle\langle A \rangle\rangle \phi \rrbracket_{\mathcal{M}}$, as Proposition 6 does not hold if ϕ is a general LTL formula. To solve this problem, we propose a novel reduction to the solving of parity games.

A. MODEL-CHECKING SIMPLE ATL*

Given a simple ATL* formula $\langle\langle A \rangle\rangle \phi$ and an ACGS $\mathcal{M} = (\mathcal{G}, \pi)$, we compute $\llbracket \langle\langle A \rangle\rangle \phi \rrbracket_{\mathcal{M}}$ by a reduction to the problem of computing the winning region of a turn-based two-player parity game. We first introduce some basic concepts which will be used in our reduction.

A *deterministic parity automaton* (DPA) \mathcal{A} is a tuple $(P, \Sigma, \delta, p_0, R)$, where

- P is a finite set of *states*;
- Σ is a finite input *alphabet*;
- $\delta : P \times \Sigma \rightarrow P$ is a *transition function*;
- $p_0 \in P$ is an *initial state*;
- $R : P \rightarrow \{0, \dots, k\}$ is a *rank function*.

A *run* ρ of \mathcal{A} over an ω -word $\alpha_0 \alpha_1 \dots \in \Sigma^\omega$ is an infinite sequence of states $\rho = p_0 p_1 \dots$ such that for every $i \geq 0$, $p_{i+1} = \delta(p_i, \alpha_i)$. Let $\text{inf}(\rho)$ be the set of states visited infinitely often in ρ . An infinite word is *recognized* by \mathcal{A} if \mathcal{A} has a run ρ over this word such that $\min_{p \in \text{inf}(\rho)} R(p)$ is *even*. For every LTL formula ϕ , one can construct a DPA $\mathcal{A}_\phi = (P, 2^{AP}, \delta, p_0, R)$ with $2^{2^{O(|\phi|)}}$ states and rank $k = 2^{O(|\phi|)}$ such that \mathcal{A}_ϕ recognizes all the ω -words satisfying ϕ [15], where each ω -word corresponds to a trace $\tau(\rho)$ of a path ρ in the ACGS.

A *(turned-based, two-player) parity game* \mathcal{P} is a tuple $(V = V_0 \uplus V_1, E, \Xi)$, where

- V_i for $i \in \{0, 1\}$ is a finite set of *vertices* controlled by *Player- i* ;

- $E \subseteq V \times V$ is a finite set of *edges*;
- $\Xi : V \rightarrow \{0, \dots, k\}$ is a *rank function*.

A *play* ρ starting from the vertex v_0 is an infinite sequence of vertices $v_0 v_1 \dots$ such that for every $i \geq 0$, $(v_i, v_{i+1}) \in E$. ρ is accepting if $\min_{v \in \text{inf}(\rho)} \Xi(v)$ is *even*. A *strategy* of Player- i is a function $\theta : V^* V_i \rightarrow V$ such that for every $\rho \in V^*$ and $v \in V_i$, $(v, \theta(\rho \cdot v)) \in E$. Given a strategy θ_0 for Player-0 and a strategy θ_1 for Player-1, let $\mathcal{P}(\theta_0, \theta_1)$ be the play where Player-0 and Player-1 enforce their strategies θ_0 and θ_1 , respectively. θ_0 is a winning strategy for Player-0 if $\mathcal{P}(\theta_0, \theta_1)$ is accepting for all strategies θ_1 of Player-1. The winning region of Player-0, denoted by WR_0 , is the set of vertices from which Player-0 has a winning strategy.

We will use the following notations in our reduction.

- $\text{dom}(g)$ denotes the domain of the function g .
- $A_\sigma := \{i \in A \mid \pi(i) = \sigma\}$ and $\bar{A}_\sigma := \{i \in \bar{A} \mid \pi(i) = \sigma\}$.
- F_{IR} is the set of (total) functions $f : A_{\text{IR}} \times S \rightarrow \bigcup_{i \in A_{\text{IR}}} A_i$ such that for all $(i, s) \in A_{\text{IR}} \times S$, $f(i, s) \in P_i(s)$ and $s \sim_i s'$ entails that $f(i, s) = f(i, s')$.
- Given a state s , let F_{IR}^s be the set of functions $f : A_{\text{IR}} \rightarrow \bigcup_{i \in A_{\text{IR}}} A_i$ such that $f(i) \in P_i(s)$ for every $i \in A_{\text{IR}}$, and $F_{\text{IR}} := \bigcup_{s \in S} F_{\text{IR}}^s$.
- G_{IR} is the set of *partial* functions $g : \bar{A}_{\text{IR}} \times S \rightarrow \bigcup_{i \in \bar{A}_{\text{IR}}} A_i$, such that for each $(i, s) \in \bar{A}_{\text{IR}} \times S$, if $g(i, s) \in \text{dom}(g)$, then for all $s' \in S$ with $s \sim_i s'$: $g(i, s) = g(i, s') \in P_i(s)$.
- $\Pi_{\text{IR}} := \{G \subseteq G_{\text{IR}} \mid \forall g, g' \in G, \text{dom}(g) = \text{dom}(g')\}$.

We construct a parity game \mathcal{P}_ϕ as follows:

$$\mathcal{P}_\phi = (V = V_0 \uplus V_1, E, \Xi),$$

where

- $V_0 = S \cup (S \times P \times F_{\text{IR}} \times \Pi_{\text{IR}})$;
- $V_1 = (S \times F_{\text{IR}}) \cup (S \times P \times F_{\text{IR}} \times F_{\text{IR}} \times \Pi_{\text{IR}})$;
- $\Xi : V \rightarrow \{0, \dots, k\}$ is a rank function such that:
 - $\Xi(s) = \Xi(s, f) = 0$ for every $s \in S$ and $f \in F_{\text{IR}}$,
 - $\Xi(s, p, f_1, G) = \Xi(s, p, f_1, f_2, G) = R(p)$, for every $s \in S, p \in P, f_1 \in F_{\text{IR}}, f_2 \in F_{\text{IR}}$ and $G \in \Pi_{\text{IR}}$.
- E is defined as follows:
 - 1) $(s, (s, f_1)) \in E$ for $(s, f_1) \in S \times F_{\text{IR}}$;
 - 2) $((s, f_1), (s, p_0, f_1, \emptyset)) \in E$ for $(s, f_1) \in S \times F_{\text{IR}}$;
 - 3) $((s, p, f_1, G), (s, p, f_1, f_2, G)) \in E$ for $(s, p, f_1, G) \in V_0$ and $f_2 \in F_{\text{IR}}^s$;
 - 4) $((s, p, f_1, f_2, G), (s', \delta(p, \lambda(s)), f_1, G')) \in E$ for every $(s, p, f_1, f_2, G) \in V_1$ and $s' \in S$, where $G' \in \Pi_{\text{IR}}$ is the largest set such that the following conditions hold: for every $g' \in G'$,
 - (1) either $G = \emptyset$ or there exists $g \in G$ such that

$$\text{dom}(g') = \text{dom}(g) \cup \{(i, s'') \in \bar{A}_{\text{IR}} \times S \mid s \sim_i s''\}$$
 and for every $(i, s'') \in \text{dom}(g)$, $g'(i, s'') = g(i, s'')$;
 - (2) there exists $\vec{a} \in A_c$ such that $s' = \Delta(s, \vec{a})$, and for every $(i, s) \in A_g \times S$,

$$\vec{a}_i = \begin{cases} f_1(i, s), & \text{if } (i, s) \in \text{dom}(f_1); \\ g'(i, s), & \text{if } (i, s) \in \text{dom}(g'); \\ f_2(i), & \text{if } i \in \text{dom}(f_2). \end{cases}$$

In this reduction, intuitively, the function $f_1 \in F_{\text{IR}}$ encodes *ir*-strategies of agents in A_{IR} , and the collection of the functions $f_2 \in F_{\text{IR}}$ in plays of \mathcal{P}_ϕ from the vertex (s, f_1) encodes *IR*-strategies of agents in A_{IR} . These functions together encode a collective strategy of the coalition A . Each function $g \in G_{\text{IR}}$ encodes *ir*-strategies of agents in \bar{A}_{IR} . The imperfect information abilities of agents are ensured by the definitions of the functions $f \in F_{\text{IR}}$ and $g \in G_{\text{IR}}$.

Intuitively, to check whether $s \in \llbracket \langle\langle A \rangle\rangle \phi \rrbracket_{\mathcal{M}}$, \mathcal{P}_ϕ starts with the vertex s . At the first step, Player-0 chooses a function $f_1 \in F_{\text{IR}}$ meaning that the *ir*-strategies of the agents in A_{IR} are chosen. Next, \mathcal{P}_ϕ moves from (s, f_1) to (s, p_0, f_1, \emptyset) which lets the DPA \mathcal{A}_ϕ start with p_0 (note that Player-1 has only one choice at this step). At a vertex (s, p, f_1, G) controlled by Player-0, Player-0 chooses actions for agents in A_{IR} by choosing one function $f_2 \in F_{\text{IR}}^s$. Then Player-1 chooses actions for agents in \bar{A} with respect to the chosen actions of agents in \bar{A}_{IR} tracked by G . These selections of actions together with f_1 and G determine a joint action \vec{a} , based on which \mathcal{P}_ϕ moves to $(s', \delta(p, \lambda(s)), f_1, G')$ such that s' is the successor state of the state s after the joint action \vec{a} , and $\delta(p, \lambda(s))$ is the successor state of the state p in \mathcal{A}_ϕ which allows to mimic the run of \mathcal{A}_ϕ over the trace $\tau(\rho)$ induced by the play ρ of \mathcal{M} . During this step, f_2 is dropped from the vertex of \mathcal{P}_ϕ , as f_2 corresponds to actions of agents in A_{IR} and needs not track. The actions of agents in \bar{A}_{IR} are preserved in G' from G . This ensures imperfect recall abilities of agents in \bar{A}_{IR} . Note that it is important to associate the functions $g \in G_{\text{IR}}$ to the same state s' that is reached via same sequence of states for a given function $f_1 \in F_{\text{IR}}$ (i.e., Item 4) in the definition of E), otherwise the agents from A_{IR} may choose different actions on the same path of \mathcal{M} .

Lemma 6. $WR_0 \cap S = \llbracket \langle\langle A \rangle\rangle \phi \rrbracket_{\mathcal{M}}$.

Proof. According to the definition of \mathcal{P}_ϕ , each $\mathcal{P}_\phi(\theta_0, \theta_1)$ must be of the form

$$\begin{aligned} & (s, f_1)(s, p_0, f_1, \emptyset)(s, p_0, f_1, f_2, \emptyset) \\ & (s_1, p_1, f_1, G_1)(s_1, p_1, f_1, f_2^1, G_1) \\ & (s_2, p_2, f_1, G_2)(s_2, p_2, f_1, f_2^2, G_2) \\ & (s_3, p_3, f_1, G_3)(s_3, p_3, f_1, f_2^3, G_3) \\ & \dots \end{aligned}$$

Let $\rho_{\theta_0, \theta_1}$ be the path $ss_1s_2s_3 \dots$ obtained from $\mathcal{P}_\phi(\theta_0, \theta_1)$. (\Rightarrow) Suppose $s \in WR_0 \cap S$, then Player-0 has a winning strategy $\theta : V^* V_0 \rightarrow V$ such that $\min_{v \in \text{inf}(\mathcal{P}_\phi(\theta_0, \theta_1))} \Xi(v)$ is even for every strategy θ_1 of Player-1.

It is well-known that, for parity games, if Player-0 has a winning strategy, then it also has a *pure memoryless* winning strategy [16]. Let $\theta_0 : V_0 \rightarrow V$ be a pure memoryless winning strategy for Player-0 such that $\min_{v \in \text{inf}(\mathcal{P}_\phi(\theta_0, \theta_1))} \Xi(v)$ is even for every strategy θ_1 of Player-1. Then, the DPA \mathcal{A}_ϕ recognizes the ω -word $\tau(\rho_{\theta_0, \theta_1})$. This implies that $\mathcal{M}, \rho_{\theta_0, \theta_1} \models \phi$.

Let ξ_A be the collective strategy for the coalition A such that

- for every $(i, s') \in A_{\text{IR}} \times S$: $\xi_A(i)(s') = f_1(i, s')$; and
- for every $i \in A_{\text{IR}}$, strategy θ_1 of Player-1 and prefix $s_1 \cdots s_j$ of $\rho_{\theta_0, \theta_1}$
 - $\xi_A(i)(s_1 \cdots s_j) = f_2^j(i)$, if $j \geq 1$;
 - $\xi_A(i)(s) = f_2(i)$, otherwise.

For every collective strategy $\xi_{\bar{A}}: \bar{A} \rightarrow \bigcup_{i \in \bar{A}} \Theta_i^{\pi(i)}$, let θ_1 be a strategy of Player-1 such that

- for every $(s, f_1) \in V_1$: $\theta_1(s, f_1) = (s, p_0, f_1, \emptyset)$;
- for every $(s, p, f_1, f_2, G) \in V_1$:

$$\theta_1(s, p, f_1, f_2, G) = (s', \delta(p, \lambda(s)), f_1, G'),$$

where G' satisfies the conditions in the reduction. Moreover, for every $g \in G'$ and $(i, s'') \in \text{dom}(g)$,

$$g(i, s'') = \xi_{\bar{A}}(i)(s).$$

Then, $\rho_{\theta_0, \theta_1} = \text{play}(s, \xi_A, \xi_{\bar{A}})$. Therefore, $s \in \llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}}$.

(\Leftarrow) Suppose $s \in \llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}}$, then there exists a collective strategy ξ_A for the coalition A such that for every collective strategy $\xi_{\bar{A}}$ of the coalition \bar{A} : $\mathcal{M}, \text{play}(s, \xi_A, \xi_{\bar{A}}) \models \phi$.

Let θ_0 be a strategy of Player-0 such that

- $\theta_0(s) = (s, f_1)$ where $f_1 \in F_{\text{IR}}$ such that for every $(i, s') \in A_{\text{IR}} \times S$: $f_1(i, s') = \xi_A(i)(s')$;
- and $\theta_0(s_j, p_j, f_1, G_{j-1}) = (s_j, p_j, f_1, f_2^j, G_{j-1})$ where $f_2^j \in F_{\text{IR}}^{s_j}$ such that for every $i \in A_{\text{IR}}$: $f_2^j(i) = \xi_A(i)(s_1 \cdots s_j)$.

Similarly, for every strategy θ_1 for Player-1, we can construct a collective strategy $\xi_{\bar{A}}$ of the coalition \bar{A} such that $\text{play}(s, \xi_A, \xi_{\bar{A}}) = \rho_{\theta_0, \theta_1}$. Since $\mathcal{M}, \text{play}(s, \xi_A, \xi_{\bar{A}}) \models \phi$, $\min_{v \in \text{inf}(\mathcal{P}_\phi(\theta_0, \theta_1))} \Xi(v)$ is even. Therefore, θ_0 is a winning strategy of Player-0, namely, $s \in WR_0 \cap S$. \square

The winning region of Player-0 in \mathcal{P}_ϕ can be computed in time polynomial in $|V| \cdot |E| \cdot 2^k$ [17]. In this reduction, each $G \in \Pi_{\text{IR}}$ contributes at most $O(|S|)$ sets of G' . Therefore, $|V| \cdot |E|$ is exponential in $|\mathcal{G}| \cdot 2^{|\phi|}$. Recall that $k = 2^{O(|\phi|)}$. Consequently, we have

Lemma 7. For the simple ATL* formula $\langle \langle A \rangle \rangle \phi$, $\llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}}$ can be computed in 2EXPTIME.

Example 2. Let us continue with the ACGS $\mathcal{M}_r = (\mathcal{G}, \pi)$ and ATL formula φ_r from Example 1. The DPA $\mathcal{A}_{q_1, u_{q_2}} = (P, 2^{\{q_1, q_2\}}, \delta, p_0, R)$ is shown in Figure 3, where $R(p_0) = R(p_2) = 1$ and $R(p_1) = 0$. The parity game \mathcal{P}_φ , starting from

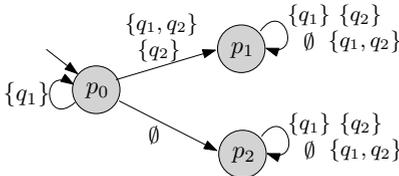


FIGURE 3. The DPA $\mathcal{A}_{q_1, u_{q_2}}$.

s_0 is shown in Figure 4, where the circle-shape vertices are controlled by Player-0, the others are controlled by Player-1, the ranks of the vertices containing state p_1 are 0, and the ranks of the other vertices are 1. We can see that $s_0 \notin WR_0$. Hence, $\mathcal{M}_r, s_0 \not\models \varphi_r$.

It is known that the model-checking problem of simple ATL* on CGSs is 2EXPTIME-complete [1]. We have that:

Theorem 4. The model-checking problem for simple ATL* is 2EXPTIME-complete.

B. MODEL-CHECKING FOR GENERAL ATL*

The model-checking algorithm (shown in Algorithm 2) for general ATL* follows Algorithm 1, where the procedure for simple ATL is replaced by the one for simple ATL*. We also extend ATL* with the epistemic modalities $\mathbf{K}_i\varphi$, $\mathbf{E}_A\varphi$, $\mathbf{D}_A\varphi$ and $\mathbf{C}_A\varphi$, leading to the logic ATLK*.

Algorithm 2: ATL* model-checking algorithm

Input: An ACGS $\mathcal{M} = (\mathcal{G}, \pi)$ and an ATL* formula φ
Output: $\llbracket \varphi \rrbracket_{\mathcal{M}}$

```

1 Function MC( $\mathcal{M}, \varphi$ )
2   switch  $\varphi$  :
3     case  $q$  : return  $\{s \in S \mid q \in \lambda(s)\}$ ;
4     case  $\neg\varphi'$  : return  $S \setminus \text{MC}(\mathcal{M}, \varphi')$ ;
5     case  $\varphi_1 \wedge \varphi_2$  : return  $\text{MC}(\mathcal{M}, \varphi_1) \cap \text{MC}(\mathcal{M}, \varphi_2)$ ;
6     case  $\mathbf{K}_i\varphi'$  : return  $\{s \in S \mid [s]^{-i} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
7     case  $\mathbf{E}_A\varphi'$  : return  $\{s \in S \mid [s]^{-E} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
8     case  $\mathbf{D}_A\varphi'$  : return  $\{s \in S \mid [s]^{-D} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
9     case  $\mathbf{C}_A\varphi'$  : return  $\{s \in S \mid [s]^{-C} \subseteq \text{MC}(\mathcal{M}, \varphi')\}$ ;
10    case  $\langle \langle A \rangle \rangle \phi$  :
11      foreach sub-state-formula  $\varphi'$  in  $\phi$  do
12        Replace  $\varphi'$  by a fresh atomic proposition  $q_{\varphi'}$ 
13        in  $\varphi$ , and let  $\lambda(q_{\varphi'}) := \text{MC}(\mathcal{M}, \varphi')$ ;
14       $\llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}} := WR_0 \cap S$ ;
15    return  $\llbracket \langle \langle A \rangle \rangle \phi \rrbracket_{\mathcal{M}}$ ;

```

By Lemma 7, the model-checking problem for ATLK* on ACGSs can be solved in 2EXPTIME.

Theorem 5. The model-checking problem for ATL* (hence ATLK*) on ACGSs is 2EXPTIME-complete.

IX. IMPLEMENTATION AND EXPERIMENTS

We implement the ATLK/ATLK* model-checking algorithms in MCMAS [18] and carry out several experiments. The tool GOAL [19] is used to transform LTL formulae to DPA and compute winning regions of parity games. All models in our experiments are based on the existing benchmarks in the literature. All experiments were conducted on a desktop with 1.70GHz Intel Core E5-2603 CPU and 32GB of memory.

A. CASTLE GAME

In Castle Game [20], there are several agents modeling workers and an environment agent. Each worker works for the benefit of a castle, and the environment keeps track of the Health Points (HP) of castles. Each castle preserves an HP up to 3, and 0 means it is defeated. Workers are able to attack a castle which they do not work for, or defend the castle which they work for, or do nothing. Any agent cannot defend its castle twice in a row, it must wait 1 step before being able to defend again. The castle gets damaged if the number of attackers is greater than the number of defenders, and the

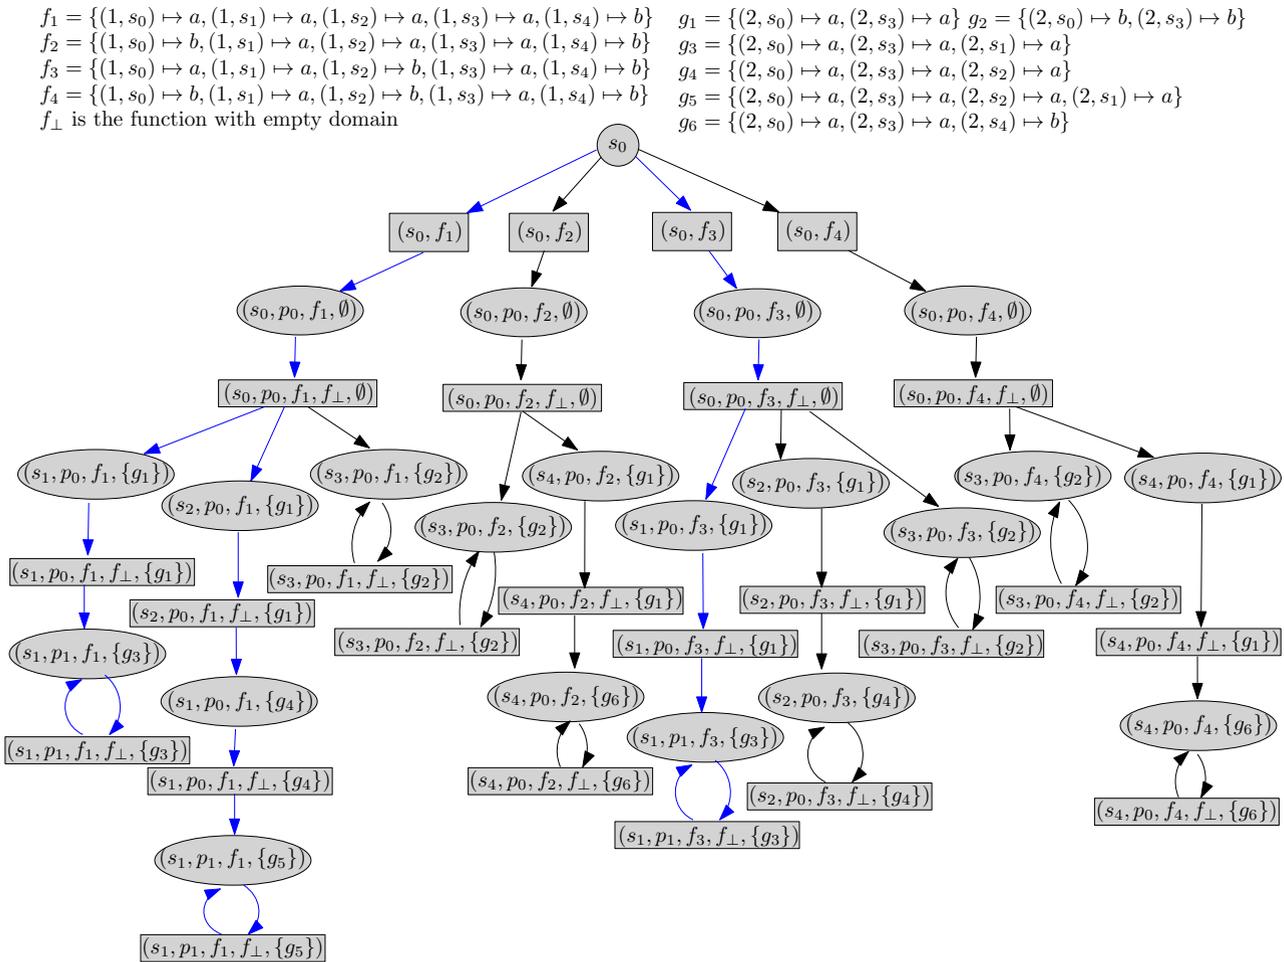


FIGURE 4. The parity game \mathcal{P}_{per} starting from the state s_0 with accepting plays highlighted in blue color.

difference influences its HP. In this model, the number of states is $8000 \times 4n$, the environment agent has 1 local action, and each worker agent has 4 local actions, where n denotes the number of workers.

In this experiment, we consider an ACGS consisting of three worker agents w_1, w_2, w_3 and an environment agent e , where worker w_i works for the castle c_i .

- $\varphi_1 \equiv \langle\langle\{w_1, w_2\}\rangle\rangle \mathbf{F}(\text{castle3Defeated})$: expresses that workers w_1 and w_2 can make castle c_3 defeated, no matter which strategies the worker w_3 uses.
- $\varphi_2 \equiv \langle\langle\{w_1, w_2\}\rangle\rangle \mathbf{F}(\text{allDefeated})$: expresses that workers w_1 and w_2 can make all the castles defeated, no matter which strategies the worker w_3 uses.

The results are shown in Table 1, where $(\sigma_1, \sigma_2, \sigma_3, \sigma_4)$ in each row denotes the strategy types of agents e, w_1, w_2, w_3 , N/A denotes timeout (2.5 hours), Y (resp. N) denotes that the model satisfies (resp. fails) the formula, and columns 2–4 (resp. 5–7) show total time (in seconds) and result of verifying φ_1 (resp. φ_2) using Algorithm 1 and Algorithm 2, respectively.

We observe that: (1) the strategy types of agents do affect

TABLE 1. Results of Castle Game

π	φ_1			φ_2		
	Alg. 1	Alg. 2	SAT	Alg. 1	Alg. 2	SAT
(IR, IR, IR, IR)	N/A	20.295	Y	N/A	18.178	Y
(IR, IR, IR, ir)	N/A	7523.67	Y	N/A	7377.44	Y
(IR, IR, ir, IR)	N/A	31.904	Y	N/A	30.578	N
(IR, ir, IR, IR)	N/A	32.446	Y	N/A	31.259	N
(IR, IR, ir, ir)	N/A	3402.56	Y	N/A	3451.59	N
(IR, ir, IR, ir)	N/A	3294.51	Y	N/A	3366.71	N
(IR, ir, ir, IR)	5.822	24.254	Y	77.514	23.37	N
(IR, ir, ir, ir)	13.791	113.493	Y	45.679	113.647	N

the performance and results. In particular, the time significantly increases when w_3 is ir-typed while w_1 or w_2 is IR-typed; (2) Algorithm 1 is more efficient when both w_1 and w_2 are ir-typed; otherwise and Algorithm 2 is more efficient. This is because the number of possible strategies of w_1 and w_2 is small (using Lemma 2) if both w_1 and w_2 are ir-typed.

B. DINING CRYPTOGRAPHERS PROTOCOL

Dining Cryptographers Protocol is one of anonymity protocols aimed at establishing the privacy of principals during an exchange [13]. The dining cryptographers protocol can

TABLE 2. Results of Dining Cryptographers Protocol

#Crypts	#States	Alg. 1			Alg. 2		
		ψ_1	ψ_2	ψ_3	ψ_1	ψ_2	ψ_3
3	160	0.022	0.016	0.013	6.439	5.838	5.852
4	384	0.059	0.049	0.028	6.928	6.744	7.242
5	896	0.133	0.114	0.049	8.839	8.874	8.88
6	2048	0.315	0.328	0.163	12.567	12.724	12.865
7	4608	0.929	1.388	0.382	22.938	23.411	23.654
8	10240	3.463	4.022	0.834	60.642	60.583	63.064
9	22528	9.19	8.913	1.721	266.844	240.003	254.293
10	49152	21.988	21.927	5.094	1712.62	1588.06	1762.88

be modeled as a MAS. In this game, n cryptographers share a meal around a circular table. Either one of them or their employer paid the bill. They want to know whether it was sponsored by their employer without revealing the identity of the payer (if one of them did pay). The protocol works as follows: each cryptographer 1) tosses a coin and shows the outcome to his/her right-hand neighbour, 2) announces whether the two coins agree or not if he/she is not payer, otherwise announces the opposite of what he/she sees. Their employer is the payer if an even number of cryptographers claiming that the two coins are different, otherwise not. For experimental purpose, we allow the cryptographer who paid for the meal announces either the two coins agree or not no matter what he/she saw.

In this experiment, n ranges from 3 to 10, two cryptographers use *ir*-strategies, and one of them should be the payer. The others all use *IR*-strategies. We verify three formulae ψ_1 , ψ_2 and ψ_3 , where ψ_i expresses that if the number of “sayd-ifferent” is odd and the i -th cryptographer is not the payer, then he/she knows that the bill is paid by one of the others, but cannot tell exactly who is the payer. For instance, in the three cryptographers case, $\psi_1 \equiv \langle\langle\emptyset\rangle\rangle\mathbf{G}((\text{odd} \wedge \neg c1\text{paid}) \rightarrow ((\mathbf{K}_{c1}(c2\text{paid} \vee c3\text{paid})) \wedge \neg\mathbf{K}_{c1}c2\text{paid} \wedge \neg\mathbf{K}_{c1}c3\text{paid}))$.

The results are shown in Table 2, where column 1 gives the number of cryptographers, column 2 gives the number of states, columns 3–5 (resp. columns 6–7) show the total time of respectively verifying ψ_1 , ψ_2 and ψ_3 using Algorithm 1 (resp. Algorithm 2). Both ψ_1 and ψ_2 are satisfied by all the models, while ψ_3 not. We observe that Algorithm 1 is more efficient than Algorithm 2, as the coalitions in all the formulae are \emptyset . From this experiment, one may conclude the reasonable scalability of our tool.

C. BOOK STORE SCENARIO

The Book Store Scenario depicts a deal between two agents: a supplier (S) and a purchaser (P) [18]. Initially, S is waiting for an order from P, and P is ready for initiating a trade. Upon receiving an order of some e-good from P, S can make a decision to accept the order or not, and later notifies P. If S accepts, then P can pay the fee. Once paid, S can either reject the payment or accept and deliver the good. If P received the good, then trade is completed. During the trade, P can revoke the order, both S and P can terminate the trade, after which the information of the trade should be symmetric at any time.

TABLE 3. Results of Book Store Scenario

π	φ_1			φ_2		
	Alg. 1	Alg. 2	SAT	Alg. 1	Alg. 2	SAT
(IR, IR)	4.237	11.264	Y	0.08	5.566	Y
(IR, Ir)	4.102	12.185	Y	0.081	5.124	Y
(IR, ir)	4.094	11.459	Y	0.081	5.26	Y
(Ir, IR)	4.095	17.398	Y	0.081	6.096	Y
(Ir, Ir)	4.086	30.649	Y	0.082	7.183	Y
(Ir, ir)	4.112	32.985	Y	0.082	8.009	Y
(ir, IR)	4.162	17.842	N	0.082	5.96	Y
(ir, Ir)	4.144	31.155	N	0.082	7.592	Y
(ir, ir)	4.157	30.73	N	0.082	7.473	Y

In this model, S has 15 local states and 13 actions, and P has 12 local states and 7 actions. In this experiment, we verify the model against the following formulae:

- $\varphi_1 \equiv \langle\langle\emptyset\rangle\rangle\mathbf{G}((S \& P_no_T) \rightarrow (\mathbf{K}_S \langle\langle\{S, P\}\rangle\rangle\mathbf{F} \text{trd_end}))$: expresses that if neither S nor P terminates the trade (i.e., $S \& P_no_T$ is true), then S knows that they can cooperatively complete the trade eventually.
- $\varphi_2 \equiv \langle\langle\{S, P\}\rangle\rangle(S \& P_no_T \mathbf{U} (\text{trd_end} \wedge \neg \text{trd_succ}))$: expresses that the trade can end by P’s requiring for refund.

The results are shown in Table 3. Each row presents the result of one of strategy type combinations of S and P, for instance, (IR, ir) denotes that S has an IR-strategy while P has ir-strategy. Columns 2–4 (resp. columns 5–7) show total time and results of verifying φ_1 (resp. φ_2) using Algorithm 1 (resp. Algorithm 2). The results of φ_1 confirm that strategy types affect the truth of formula. Algorithm 1 performs better than Algorithm 2 both on φ_1 and φ_2 in this experiment.

X. RELATED WORK

The family of alternating-time temporal logics (ATL, ATL* and AMC [1]) for reasoning about games was introduced with motivations partially from MASs. Model-checking algorithms were also given with IR-strategies. [21] extended ATL with knowledge operators and proposed corresponding model-checking algorithms. In their work, epistemic accessibility relations are considered in the interpretation of knowledge operators, but not for the strategies and outcomes. This means that agents still use IR strategies for collation modalities $\langle\langle A \rangle\rangle\psi$. This issue was discussed in [22] which proposed an idea of *ir*-strategies. [2] introduced the notion of imperfect recall into ATL/ATL*, and systematically investigated the complexity of model-checking problems for ATL/ATL* under four different strategic types. Importantly, with *ir*-strategies the model-checking problem becomes undecidable [5]. Authors in [6] introduced knowledge operators into AMC, studied its semantics and proposed a model-checking algorithm for the alternation-free fragment under the imperfect information setting. [8] further conducted a comprehensive comparison of variants of ATL/ATL* with different strategic abilities. The study corroborates that the agents’ strategic abilities play a prominent role in logic semantics.

In the previous work, strategies of agents are revocable, i.e., when it comes to achieve a goal in the (nested) subformulae, previously selected strategies are deleted. [4] introduced a variant of ATL with *irrevocable* strategies under the imperfect recall setting. It was generalized into ATL/ATL* with strategy contexts [23], which allowed agents to drop or inherit previously selected strategies.

Two versions of strategic logics were introduced by [24] and [25], and the model-checking problems were investigated therein under the IR-setting. Strategic logics extend LTL with first-order quantifications over strategies which naturally capture the multi-player game nature in the evolution of MASSs. Knowledge operators were introduced in the strategic logic [25] where a model-checking algorithm with *ir*-strategies was given [9]. Here all agents must take *ir*-strategies (so the potential inconsistency can be ruled out), but no other strategic abilities were considered. To gain decidability under *ir*-setting, specific restrictions on the abilities of the agents were proposed in, for instance, [26]–[32]. Several subsets of the strategic logics [25] such as BSIL [33], TCL [34] were proposed and studied under the IR-setting in order to maintain a low complexity.

Our work is orthogonal to the existing work which defines the strategic abilities at the semantics level, but takes a more syntactic level by strengthening the model.

XI. CONCLUSION AND FUTURE WORK

In this paper, we discussed the problem of existing semantics of ATL/ATL*, and advocated the approach to make agents' abilities explicit in modeling. For this purpose, we introduced an extension of standard CGS model, named ACGS, which defines agents' abilities at the syntactic level of the system model. We explored the effects of strategy types in the semantics, in particular model-checking, of ATL/ATL* over ACGSs, and provided model-checking algorithms with identified complexity. The algorithms are implemented in a tool MCMAS-ACGS, which has been applied to several applications to demonstrate the feasibility and effectiveness. This work represents the first systematic study towards different agents' abilities at the syntactic level, which is in contrast to previous approaches at the semantic level.

Currently we use ATL/ATL* as the specification, but the methodology can be extended to other logics such as Strategy Logic, and other agents' abilities such as strategy contexts. Several questions are left open such as axiomatization and satisfiability problem. We leave them for future work.

REFERENCES

- [1] R. Alur, T. A. Henzinger, and O. Kupferman, "Alternating-time temporal logic," *Journal of the ACM*, vol. 49, no. 5, pp. 672–713, 2002.
- [2] P. Schobbens, "Alternating-time logic with imperfect recall," *Electronic Notes in Theoretical Computer Science*, vol. 85, no. 2, pp. 82–93, 2004.
- [3] W. Jamroga and W. van der Hoek, "Agents that know how to play," *Fundamenta Informaticae*, vol. 63, no. 2-3, pp. 185–219, 2004.
- [4] T. Ágotnes, V. Goranko, and W. Jamroga, "Alternating-time temporal logics with irrevocable strategies," in *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 2007, pp. 15–24.
- [5] C. Dima and F. L. Tiplea, "Model-checking ATL under imperfect information and perfect recall semantics is undecidable," *CoRR*, vol. abs/1102.4225, 2011.
- [6] N. Bulling and W. Jamroga, "Alternating epistemic mu-calculus," in *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 2011, pp. 109–114.
- [7] F. Laroussinie and N. Markey, "Augmenting ATL with strategy contexts," *Information and Computation*, vol. 245, pp. 98–123, 2015.
- [8] N. Bulling and W. Jamroga, "Comparing variants of strategic ability: how uncertainty and memory influence general properties of games," *Autonomous Agents and Multi-Agent Systems*, vol. 28, no. 3, pp. 474–518, 2014.
- [9] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano, "Practical verification of multi-agent systems against SLK specifications," *Information and Computation*, vol. 261, no. Part, pp. 588–614, 2018.
- [10] F. Laroussinie, N. Markey, and G. Oreiby, "On the expressiveness and complexity of ATL," *Logical Methods in Computer Science*, vol. 4, no. 2, 2008.
- [11] M. Cohen, M. Dam, A. Lomuscio, and F. Russo, "Abstraction in model checking multi-agent systems," in *Proceedings of the 8th International Joint Conference on Autonomous Agents and Multiagent Systems*, 2009, pp. 945–952.
- [12] E. M. Clarke, E. A. Emerson, and A. P. Sistla, "Automatic verification of finite state concurrent systems using temporal logic specifications: A practical approach," in *Proceedings of the Conference Record of the 10th Annual ACM Symposium on Principles of Programming Languages*, 1983, pp. 117–126.
- [13] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano, "MCMAS-SLK: A model checker for the verification of strategy logic specifications," in *Proceedings of the 26th International Conference on Computer Aided Verification*, 2014, pp. 525–532.
- [14] F. Laroussinie, N. Markey, and P. Schnoebelen, "Model checking CTL+ and FCTL is hard," in *Proceedings of the 4th International Conference on Foundations of Software Science and Computation Structures*, 2001, pp. 318–331.
- [15] N. Piterman, "From nondeterministic Büchi and Streett automata to deterministic parity automata," in *Proceedings of the 21st IEEE Symposium on Logic in Computer Science*, 2006, pp. 255–264.
- [16] Y. Gurevich and L. Harrington, "Trees, automata, and games," in *Proceedings of the 14th Annual ACM Symposium on Theory of Computing*, 1982, pp. 60–65.
- [17] M. Jurdzinski, "Small progress measures for solving parity games," in *Proceedings of the 17th Symposium on Theoretical Aspects of Computer Science*, 2000, pp. 290–301.
- [18] A. Lomuscio, H. Qu, and F. Raimondi, "MCMAS: an open-source model checker for the verification of multi-agent systems," *International Journal on Software Tools for Technology Transfer*, vol. 19, no. 1, pp. 9–30, 2017.
- [19] Y. Tsay, Y. Chen, M. Tsai, K. Wu, and W. Chan, "GOAL: A graphical tool for manipulating büchi automata and temporal formulae," in *Proceedings of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, 2007, pp. 466–471.
- [20] J. Pilecki, M. A. Bednarczyk, and W. Jamroga, "Model checking properties of multi-agent systems with imperfect information and imperfect recall," in *Proceedings of the 7th International Conference on Intelligent Systems*, 2014, pp. 415–426.
- [21] W. van der Hoek and M. Wooldridge, "Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications," *Studia Logica*, vol. 75, no. 1, pp. 125–157, 2003.
- [22] W. Jamroga, "Some remarks on alternating temporal epistemic logic," *Proceedings of the Workshop on Formal Approaches to Multi-Agent Systems*, pp. 133–140, 2003.
- [23] A. D. C. Lopes, F. Laroussinie, and N. Markey, "ATL with strategy contexts: Expressiveness and model checking," in *Proceedings of the 30th Conference on Foundations of Software Technology and Theoretical Computer Science*, 2010, pp. 120–132.
- [24] K. Chatterjee, T. A. Henzinger, and N. Piterman, "Strategy logic," *Information and Computation*, vol. 208, no. 6, pp. 677–693, 2010.
- [25] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi, "Reasoning about strategies: On the model-checking problem," *ACM Transactions on Computational Logic*, vol. 15, no. 4, pp. 34:1–34:47, 2014.
- [26] R. Berthon, B. Maubert, and A. Murano, "Decidability results for ATL* with imperfect information and perfect recall," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 1250–1258.

- [27] R. Berthon, B. Maubert, A. Murano, S. Rubin, and M. Y. Vardi, "Strategy logic with imperfect information," in Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, 2017, pp. 1–12.
- [28] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin, "Verification of broadcasting multi-agent systems against an epistemic strategy logic," in Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, 2017, pp. 91–97.
- [29] —, "Verification of multi-agent systems with imperfect information and public actions," in Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, 2017, pp. 1268–1276.
- [30] B. Maubert and A. Murano, "Reasoning about knowledge and strategies under hierarchical information," in Proceedings of the Sixteenth International Conference on Principles of Knowledge Representation and Reasoning, 2018, pp. 530–540.
- [31] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin, "Decidable verification of multi-agent systems with bounded private actions," in Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, 2018, pp. 1865–1867.
- [32] W. Jamroga, V. Malvone, and A. Murano, "Natural strategic ability under imperfect information," in Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, 2019, pp. 962–970.
- [33] F. Wang, S. Schewe, and C. Huang, "An extension of ATL with strategy interaction," ACM Transactions on Programming Languages and Systems, vol. 37, no. 3, p. 9, 2015.
- [34] C. Huang, S. Schewe, and F. Wang, "Model-checking iterated games," Acta Inf., vol. 54, no. 7, pp. 625–654, 2017.

• • •