Alternating-time Temporal Logic with Stochastic Abilities

Gabriel Ballot
SEIDO Lab, EDF R&D and Télécom Paris, Institut Polytechnique de Paris
Palaiseau, France
gabriel.ballot@telecom-paris.fr

Vadim Malvone
LTCI, Télécom Paris, Institut Polytechnique de Paris
Palaiseau, France
vadim.malvone@telecom-paris.fr

Jean Leneutre
LTCI, Télécom Paris, Institut Polytechnique de Paris
Palaiseau, France
jean.leneutre@telecom-paris.fr

Jingxuan Ma
SEIDO Lab, EDF R&D
Palaiseau, France
jingxuan.ma@edf.fr

Mourad Leslous
SEIDO Lab, EDF R&D
Palaiseau, France
mourad.leslous@edf.fr

ABSTRACT

Multi-agent systems strategic verification is a branch of formal methods to model, reason about, and verify strategic behavior in complex environments. The notion of agent *capacity* was introduced alongside the strategic logic CapATL to model multi-agent systems in which each player may exhibit diverse abilities or profiles. These capacities can represent various aspects, such as an agent's experience level, personality traits, type, or version. In real-world applications, domain knowledge or prior statistical analyses may provide a probability distribution over the possible profiles of each agent. This leads to the concept of *stochastic abilities*, where capacities are assigned probabilistically, yet remain private to other agents. In this context, we introduce a novel probabilistic strategic logic, called ATL-SA, that allows the expression of properties concerning the likelihood that agents or coalitions can achieve specific temporal objectives under uncertainty about their capacities. We study the upper and lower complexity bounds of ATL-SA model checking and demonstrate its practical applicability through a use case in cybersecurity, showcasing its potential for analysing systems with probabilistic agent profiles.

KEYWORDS

Strategic Reasoning; Model Checking; Cybersecurity

ACM Reference Format:
Gabriel Ballot, Vadim Malvone, Jean Leneutre, Jingxuan Ma, and Mourad Leslous. 2025. Alternating-time Temporal Logic with Stochastic Abilities. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

As modern systems grow increasingly complex, their specifications become more intricate, making it challenging for humans to ensure that implementations perform as intended. This difficulty often leads to errors, motivating the development of formal verification techniques that rigorously prove system correctness. Among these techniques, model checking [14] stands out as a powerful approach, ensuring that all possible system behaviors meet specified requirements. Model checking involves three essential components: a modeling formalism to abstract the system, a logical formalism to define properties, and a model-checking algorithm to verify if the system model satisfies these properties.

With the growing interconnection of systems, the need to analyse open systems has become more pronounced, particularly in the context of *Multi-Agent Systems* (MAS), that describe the interactions between autonomous entities with distinct objectives. In this domain, *Concurrent Game Structures* (CGS) model agents' actions, the system's states, and the propositions true in each state. From any given state, agents choose actions, and the system transitions to a new state based on the joint actions of the agents. Alur *et al.* introduced *Alternating-time Temporal Logic* (ATL) [3] to express the strategic capabilities of agent coalitions in achieving temporal objectives within a CGS. For example, one might specify that if a `read` command is sent to a memory controller, a register cannot be written until the read operation completes, which can be expressed using the ATL property $readCmd \rightarrow \langle controller\rangle(\neg write) \cup read$. Since ATL model checking is computationally feasible in polynomial time, it has gained widespread use in verifying MAS.

Over time, ATL has been extended to address a variety of objectives, including epistemic properties [18–20, 25, 26, 28, 29], quantitative aspects [2, 23], probabilistic outcomes [7, 10, 13, 18, 21, 23, 25], real-time constraints [12, 15], strategy specifications [22, 30, 31], or action specifications [1, 5, 8, 9, 17]. These extensions have significantly broadened the scope of ATL in practical applications.

Our work resides at the intersection of probabilistic strategic logics and action specifications. We introduce the notion of *stochastic agents* in CGS with ATL objectives, formalized in a new framework called *Alternating-time Temporal Logic with Stochastic Abilities* (ATL-SA). Stochastic agents possess a profile, referred to as a *capacity*, which limits the actions available to them during the system's execution, similar to the approach in [5]. The distribution of these profiles is given as part of the game structure. For example, an agent might be either right- or left-handed, and the agent's handedness could restrict the set of actions it can perform. However, as opposed to [5], we consider a probability distribution over agent profiles. For instance, we can model the fact that 90% of agents are right-handed and 10% are left-handed, and compute optimal strategies for agents aiming to achieve temporal objectives with highest

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), A. El Fallah Seghrouchni, Y. Vorobeychik, S. Das, A. Nowe (eds.), May 19 – 23, 2025, Detroit, Michigan, USA.

probability in this uncertain context. The concept of agents' capacities is relevant in various contexts, including distributed computing, where systems have different resources; heterogeneous robot fleets, where robots exhibit varying capabilities; social structures, where agents possess distinct personality traits (e.g., altruistic, adventurous, or selfish); and cybersecurity, where different attacker profiles require tailored defense strategies. In the cybersecurity domain, defending against all potential attackers may be infeasible, but optimizing the probability of defending against a stochastic attacker becomes a crucial and challenging problem. The stochastic aspect introduced in this work is distinct from the conventional notion of stochastic concurrent game structures, where each joint action results in probabilistic outcomes. It also differs from stochastic strategies, where agents select distributions over actions based on history. The notion of stochastic agents can be combined with these elements, opening up interesting avenues for future research. We refer to Section 6 for further explanation of how stochastic agents differ from previous notions of stochasticity in MAS.

Contributions. This paper makes the following contributions: *(i)* we introduce ATL-SA to reason about strategic and temporal objectives in the context of stochastic agents, considering both communicating and uniform semantics, *(ii)* for both semantics, we prove the NEXPTIME-completeness of ATL-SA model checking with single strategy, the PTIME-completeness for the bounded-capacity case, and the $\mathsf{P}^{\mathsf{NEXPTIME}}$-membership in the general case, *(iii)* we demonstrate the applicability of ATL-SA with a cybersecurity illustration, where we compute optimal defense strategies against stochastic attackers.

Outline. Section 2 introduces the game structure and Section 3 formalizes ATL-SA's syntax and semantics. Section 4 studies the model-checking problem. Section 5 illustrates the practical use of ATL-SA through a cybersecurity case study. Finally, Sections 6 and 7 compare our work to existing research and conclude the paper.

2 GAME STRUCTURE

This section introduces some foundational concepts and outlines the structure of the game and its related definitions.

Let X and Y be two sets. We denote by $\mathcal{P}(X)$ the power set of X, by $f : X \to Y$ a total function from X to Y, and by $g : X \rightharpoonup Y$ a partial function from X to Y. For two functions $f : X_1 \to Y$ and $g : X_2 \to Y$, where $X_1 \cap X_2 = \emptyset$, the function $f \oplus g$ from $X_1 \cup X_2$ to Y assigns $f(x)$ (resp. $g(x)$) to an input x in X_1 (resp. X_2). The set of positive natural numbers is denoted by $\mathbb{N}$. A *signature* is a tuple $\langle \mathrm{Ag}, \Pi \rangle$, where $\mathrm{Ag} = \{1, \ldots, n\}$ is a set of $n \in \mathbb{N}$ *agents*, and Π is a finite set of *atomic propositions*.

A countable *probability space* $(\Omega, \mathbb{P})$ is a non-empty countable set Ω of *outcomes*, defining the set $\mathcal{P}(\Omega)$ of *events*, and a *probability function* $\mathbb{P} : \mathcal{P}(\Omega) \to [0,1]$ that assigns a probability to each event.[1] It must satisfy $\mathbb{P}(\Omega) = 1$ and, for any $W \subseteq \Omega$, $\mathbb{P}(W) = \sum_{\omega \in W} \mathbb{P}(\{\omega\})$. Alternatively, a countable probability space can be defined via a *probability distribution* over Ω, i.e., a function $\delta : \Omega \to [0,1]$ such that $\sum_{\omega \in \Omega} \delta(\omega) = 1$. The corresponding probability space is $(\Omega, \mathbb{P})$, where for $W \subseteq \Omega$, $\mathbb{P}(W) = \sum_{\omega \in W} \delta(\omega)$.

[1]In the general case, a probability space is given by $(\Omega, \mathcal{F}, \mathbb{P})$, where $\mathcal{F} \subseteq \mathcal{P}(\Omega)$ is a σ-algebra. However, since Ω is countable, we implicitly take $\mathcal{F} = \mathcal{P}(\Omega)$.

Action	Description
w	wait
i	intimidate
g	give up
ng	do not give up
cl	compete with left hand
cr	compete with right hand
sr	use strong right hand
sl	use strong left hand
wr	use weak right hand
wl	use weak left hand

(a) Actions.

Agent	Cap	Actions	Proba
1	c_1^r	w, i, cl, cr, sr, wl	0.8
	c_1^l	w, i, cl, cr, sl, wr	0.2
2	c_2^r	w, ng, sr, wl	0.75
	c_2^l	w, ng, sl, wr	0.15
	c_2^{rc}	w, g, sr, wl	0.08
	c_2^{lc}	w, g, sl, wr	0.02

(b) Capacities.

Table 1: Arm wrestling parameters.

A discrete *random variable* over a discrete probability space $(\Omega, \mathbb{P})$ with values in a countable set E is a function $V : \Omega \to E$.

We model the environment using a *Concurrent Game Structure with Stochastic Abilities* (CGS-SA), an extension of a deterministic CGS that includes a probability distribution for each agent over subsets of actions.

Definition 2.1 (Concurrent Game Structure with Stochastic Abilities). A CGS-SA $\mathcal{G} = \langle \mathrm{Ag}, \mathrm{St}, \Pi, \pi, \mathrm{Ac}, d, o, \Delta_1, \ldots, \Delta_{|\mathrm{Ag}|} \rangle$ over a signature $\langle \mathrm{Ag}, \Pi \rangle$ is a structure with: a finite set of states St, a labeling function $\pi : \mathrm{St} \to \mathcal{P}(\Pi)$, a finite set of actions Ac, a protocol function $d : \mathrm{Ag} \times \mathrm{St} \to \mathcal{P}(\mathrm{Ac})$ where $d(a, s)$ is the set of actions available for the agent $a \in \mathrm{Ag}$ in the state $s \in \mathrm{St}$, a transition function $o : \mathrm{St} \times \mathrm{Ac}^n \rightharpoonup \mathrm{St}$ defined for all $(s, \alpha_1, \ldots, \alpha_n)$ verifying $\alpha_a \in d(a, s)$ for all $a \in \mathrm{Ag}$, and, for each $a \in \mathrm{Ag}$, a probability distribution $\Delta_a : \mathcal{P}(\mathrm{Ac}) \to [0,1]$ over action subsets. We assume that the probability distribution of each agent is independent. A CGS-SA must verify the *progression condition*, that is, for all agents $a \in \mathrm{Ag}$, for all $A \subseteq \mathrm{Ac}$ such that $\Delta_a(A) > 0$, for all states $s \in \mathrm{St}$, we have $A \cap d(a, s) \neq \emptyset$.

Throughout this paper, when referring to a CGS-SA $\mathcal{G}$, we assume that $\mathcal{G} = \langle \mathrm{Ag}, \mathrm{St}, \Pi, \pi, \mathrm{Ac}, d, o, \Delta_1, \ldots, \Delta_{|\mathrm{Ag}|} \rangle$, with a signature $\langle \mathrm{Ag}, \Pi \rangle$ where $|\mathrm{Ag}| = n$. During the game execution, each agent $a \in \mathrm{Ag}$ is assigned a subset of actions $c \subseteq \mathrm{Ac}$ according to a probability $\Delta_a(c)$. The agent must use actions from this subset for the remainder of the game (or until a new subset is assigned). A subset $c \subseteq \mathrm{Ac}$ is called a (positive-probability) *capacity* of agent a if $\Delta_a(c) > 0$. We denote the set of such capacities by $\Delta_a^{>0} = \{c \subseteq \mathrm{Ac} \mid \Delta_a(c) > 0\}$. The *progression condition* from Definition 2.1 ensures that any agent $a \in \mathrm{Ag}$, assigned a capacity $c \in \Delta_a^{>0}$, can perform at least one action in every state. A CGS is a CGS-SA where, for all $a \in \mathrm{Ag}$, $\Delta_a(\mathrm{Ac}) = 1$, allowing us to omit distributions and simply define a CGS by $\langle \mathrm{Ag}, \mathrm{St}, \Pi, \pi, \mathrm{Ac}, d, o \rangle$.

Example 2.2. Consider an arm wrestling match between two agents $\mathrm{Ag} = \{1, 2\}$. Agent 1 decides whether to compete with the right or left arm. Additionally, agent 1 can attempt to intimidate agent 2, hoping for a forfeit. The outcome depends on the agents' handedness. The available actions are $\mathrm{Ac} = \{w, i, g, ng, cl, cr, sr, sl, wr, wl\}$, with interpretations shown in Table 1a. Agent 1 may be right-handed ($c_1^r = \{w, i, cl, cr, sr, wl\}$, probability 0.8) or left-handed ($c_1^l = \{w, i, cl, cr, sl, wr\}$, probability 0.2). Agent 2 has the capacities: right-handed ($c_2^r = \{w, ng, sr, wl\}$, probability 0.75), left-handed ($c_2^l = \{w, ng, sl, wr\}$, probability 0.15), right-handed coward ($c_2^{rc} = \{w, g, sr, wl\}$, probability 0.08), or left-handed coward

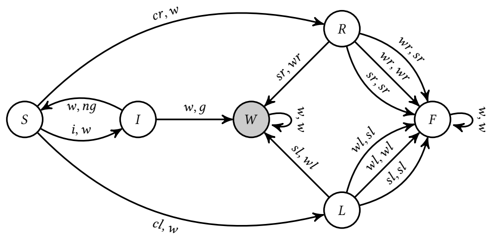

Figure 1: CGS-SA for the arm wrestling game.

($c_2^{lc} = \{w, g, sl, wr\}$, probability 0.02). These probabilities are shown in Table 1b. The game is modeled as the CGS-SA shown in Figure 1. State S is the initial state, I means intimidation, R (resp. L) is for right-hand (resp. left-hand) competition, W (resp. F) is accessed when agent 1 wins (resp. fails).

A CGS-SA is interpreted through *paths*, which specify possible executions of the game, keeping track of both the visited states and actions taken. Unlike ATL, which records only the states, our definition includes actions to reflect the agents' capacities.

Definition 2.3 (Path). A path in a CGS-SA $\mathcal{G}$ is a possibly infinite word of the form $s_1\vec{\alpha}_1 s_2\vec{\alpha}_2\ldots$, where $\vec{\alpha}_i = (\alpha_i^1, \ldots, \alpha_i^n) \in \mathsf{Ac}^n$ represents the joint action of the agents at step i. The path must satisfy the condition that for all i, $s_{i+1} = o(s_i, \alpha_i^1, \ldots, \alpha_i^n)$. If a path is finite, it ends with a state. The set of all paths is denoted by $\mathsf{Pt}_{\mathcal{G}}$, and the set of finite paths is denoted by $\mathsf{Pt}_{\mathcal{G}}^{<\omega}$.

Given a path $\rho = s_1\vec{\alpha}_1 s_2\vec{\alpha}_2\ldots$ in $\mathcal{G}$, for any state index i in ρ, we define the prefix $\rho_{\leq i} = s_1\vec{\alpha}_1\ldots s_i$, the suffix $\rho_{\geq i} = s_i\vec{\alpha}_i\ldots$, the i^{th} state $\rho[i] = s_i$ and, the last state (if finite) $\mathsf{last}(\rho)$. Given a joint action $\vec{\alpha} = (\alpha_1, \ldots, \alpha_n) \in \mathsf{Ac}^n$, we denote by $\vec{\alpha}[a] = \alpha_a$ the action of agent a. A *transition* is a path with exactly two states, and the set of transitions in $\mathcal{G}$ is denoted $\mathsf{Pt}_{\mathcal{G}}^{\langle 2\rangle}$.

We call a *capacity assignment* for a set of agents $Y \subseteq \mathsf{Ag}$ (referred to as *coalition*) a function $\kappa : Y \to \mathcal{P}(\mathsf{Ac})$ such that, for all agents $a \in \mathsf{Ag}$, $\kappa(a) \in \Delta_a^{>0}$. A capacity assignment κ is *complete* if its domain is Ag and the set of complete capacity assignments in $\mathcal{G}$ is denoted by $\Gamma_{\mathcal{G}}$. The distributions $\Delta_1, \ldots, \Delta_n$ induce the discrete probability space $(\Gamma_{\mathcal{G}}, \mathbb{P}_{\mathcal{G}})$ such that, for $\kappa \in \Gamma_{\mathcal{G}}$, we have $\mathbb{P}_{\mathcal{G}}[\{\kappa\}] = \prod_{a\in\mathsf{Ag}} \Delta_a(\kappa(a))$. As the events for different agents are assumed independent, $\mathbb{P}_{\mathcal{G}}[\{\kappa\}]$ measures the probability of the event that all agents in $a \in \mathsf{Ag}$ have respectively the capacity $\kappa(a)$.

The agents' action choice in $\mathcal{G}$ at each time point is formalized as a *strategy*. It is a function $s : \Gamma_{\mathcal{G}} \times \mathsf{Pt}_{\mathcal{G}}^{<\omega} \to \mathsf{Ac}$ that maps each complete capacity assignment and finite path (history) to an action. A strategy assignment is a partial function $\sigma : \mathsf{Ag} \rightharpoonup (\Gamma_{\mathcal{G}} \times \mathsf{Pt}_{\mathcal{G}}^{<\omega} \to \mathsf{Ac})$ that assigns strategies to agents such that, for all $a \in \mathsf{dom}(\sigma)$, $\kappa \in \Gamma_{\mathcal{G}}$, and $\rho \in \mathsf{Pt}_{\mathcal{G}}^{<\omega}$, it verifies $\sigma(a)(\kappa, \rho) \in \kappa(a) \cap d(a, \mathsf{last}(\rho))$. We say that a strategy assignment is *complete* if and only if its domain is Ag. A strategy assignment σ with domain $Y \subseteq \mathsf{Ag}$ is *uniform* (resp. *distributed*) if, for all agents $a \in Y$, all finite paths $\rho \in \mathsf{Pt}_{\mathcal{G}}^{<\omega}$, and all complete capacity assignments κ and κ', we have $\kappa(a) = \kappa'(a)$ (resp. for all $b \in Y$, $\kappa(b) = \kappa'(b)$) implies $\sigma(a)(\kappa, \rho) = \sigma(a)(\kappa', \rho)$. We abbreviate by u- and d-strategy assignment the uniform and distributed strategy assignments, respectively.

An *outcome* of a complete strategy assignment starting from a given state in a CGS-SA is a discrete random variable, which yields the unique path that adheres to the specified strategy assignment.

Definition 2.4 (Outcome). Let $\mathcal{G}$ be a CGS-SA, $s \in \mathsf{St}$, and σ be a complete strategy assignment. The *outcome* $\mathsf{Out}_{\mathcal{G}}(s, \sigma) : \Gamma_{\mathcal{G}} \to \mathsf{Pt}_{\mathcal{G}}$ is a discrete random variable defined such that for a complete capacity assignment κ, $\mathsf{Out}_{\mathcal{G}}(s, \sigma)(\kappa)$ is the unique infinite path $s_1\vec{\alpha}_1 s_2\vec{\alpha}_2\ldots$ with $s_1 = s$. For all $i \in \mathbb{N}$ and $a \in \mathsf{Ag}$, it holds that $\vec{\alpha}_i[a] = \sigma(a)(\kappa, s_1\vec{\alpha}_1\ldots\vec{\alpha}_{i-1}s_i)$.

To provide context, we compare our definition of outcomes with those used in ATL [3] and its probabilistic extension *Probability ATL* (PATL) [13]. In ATL, the strategy assignment is partial, meaning that the outcome is a set of paths representing all possible responses from the opponents, with one path per response. In PATL, the strategy assignment is also partial, but the outcome is a set of probability distributions over paths, each corresponding to different opponent responses. In contrast, our approach fixes the opponent's strategy by requiring the second argument of Out to be a complete strategy assignment, as described in Definition 2.4. This change simplifies the outcome to a single probability distribution rather than a set of distributions, one for each possible response from the opponent. Consequently, the semantics of ATL-SA (*cf.*, Definition 3.3) involves a universal quantification over the opponent's strategies, rather than over a set of outcomes (as in ATL and PATL). This design choice has several advantages. By avoiding the need to define a function that returns a set of distributions, it clarifies the fact that there is exactly one distribution corresponding to each response of the opponent. Furthermore, this approach makes the formulation of the semantics more intuitive, as it directly quantifies over the strategies of both the strategic coalition and the opponents. Moreover, this allows writing the probability of an outcome satisfying a given formula in a simple form (*cf.*, Definition 3.3).

With the game structure, paths, and outcomes defined, we are now ready to introduce the logic ATL-SA, which will express the properties of CGS-SA in the following section.

3 LOGIC

The properties of CGS-SAs encompass the strategic abilities of agents, temporal conditions, and probabilistic aspects. This section formalizes ATL-SA, an extension of ATL [3], to articulate these properties.

Syntax. ATL-SA extends ATL [3] by introducing a threshold on the strategic operator. This threshold specifies the probability with which a given temporal objective must be satisfied.

Definition 3.1 (Syntax). The following grammar defines an ATL-SA formula ϕ on a signature $\langle \mathsf{Ag}, \Pi\rangle$:

$$\phi ::= \ell \mid \neg\phi \mid \phi \wedge \phi \mid \langle Y\rangle^{\bowtie p}\psi$$
$$\psi ::= \mathsf{X}\,\phi \mid \phi\,\mathsf{U}\,\phi \mid \phi\,\mathsf{R}\,\phi$$

where $\ell \in \Pi$ is an atomic proposition, $Y \subseteq \mathsf{Ag}$ is an agent coalition, $\bowtie \in \{\leq, <, >, \geq\}$ is a comparison operator, and $p \in [0, 1]$ represents a probability threshold.

The Boolean operators $\vee$, $\to$, and $\leftrightarrow$ are defined in the usual manner. The set of agents Y is referred to as a strategic coalition, while $\langle\cdot\rangle$ is called the strategic operator. The expression $\langle Y\rangle^{\bowtie p}\psi$ means "there exists a strategy for Y to ensure ψ with a probability comparable to p with $\bowtie$." A subformula ψ is termed a temporal formula, where X is the "next" operator, U is the "until" operator, and R is the "releases" operator. We also define the "finally" operator F such that $\mathsf{F}\phi := \top \mathsf{U}\, \phi$, and the "globally" operator G such that $\mathsf{G}\,\phi := \bot \,\mathsf{R}\, \phi$. Notably, ATL-SA shares the same syntax as PATL [13] (with the exception of the = operator, which is arguably less interesting in practice and omitted for simplicity; see Remark 1). We can encompass ATL's syntax by restricting ATL-SA with $\bowtie$ set to $\geq$ and $p = 1$. Thus, the p and $\bowtie$ can be omitted when we clearly define an ATL formula.

Example 3.2. Building on Example 2.2, the formula $\phi = \langle 1\rangle^{\geq 0.37}\mathsf{F}(win)$, where win is a label that is true in state W only, means: "there exists a strategy for agent 1 to ensure reaching the state W in the future with a probability greater than or equal to 0.37."

Semantics. ATL-SA semantics extends ATL with the probability of agents' capacities as defined by the CGS-SA. We define the u- and d-semantics for u- and d-strategy assignments, respectively.

Definition 3.3 (Semantics). Let $\mathcal{G}$ be a CGS-SA over a signature $\langle \mathrm{Ag}, \Pi\rangle$, ρ be an infinite path in $\mathcal{G}$, $\ell \in \Pi$ be an atomic proposition, $a \in \mathrm{Ag}$ be an agent, and $Y \subseteq \mathrm{Ag}$ be a coalition of agents. Let (ϕ, ϕ_1, ϕ_2) denote three ATL-SA formulae on $\langle \mathrm{Ag}, \Pi\rangle$, and ψ be a temporal formula on $\langle \mathrm{Ag}, \Pi\rangle$. For $x \in \{\mathrm{d}, \mathrm{u}\}$, we define ATL-SA x-semantics through the following satisfaction relation:

- $(\mathcal{G}, \rho) \models_x \ell$ iff $\ell \in \pi(\rho[1])$,
- $(\mathcal{G}, \rho) \models_x \neg\phi$ iff $(\mathcal{G}, \rho) \not\models_x \phi$,
- $(\mathcal{G}, \rho) \models_x \phi_1 \wedge \phi_2$ iff $(\mathcal{G}, \rho) \models_x \phi_1$ and $(\mathcal{G}, \rho) \models_x \phi_2$,
- $(\mathcal{G}, \rho) \models_x \langle Y\rangle^{\bowtie p}\psi$ iff there exists an x-strategy assignment σ_Y, called a winning strategy assignment for Y, such that for all strategy assignments $\sigma_{\mathrm{Ag}\setminus Y}$ for $\mathrm{Ag} \setminus Y$, we have[2]
$$\mathbb{P}_{\mathcal{G}}\left[(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(\rho[1], \sigma_Y \oplus \sigma_{\mathrm{Ag}\setminus Y})) \models_x \psi\right] \bowtie p$$
- $(\mathcal{G}, \rho) \models_x \mathsf{X}\,\phi$ iff $(\mathcal{G}, \rho_{\geq 2}) \models_x \phi$,
- $(\mathcal{G}, \rho) \models_x \phi_1 \mathsf{U} \phi_2$ iff there exists $i \in \mathbb{N}$ such that $(\mathcal{G}, \rho_{\geq i}) \models_x \phi_2$ and for all $j \in \mathbb{N}$ where $j < i$, we have $(\mathcal{G}, \rho_{\geq j}) \models_x \phi_1$,
- $(\mathcal{G}, \rho) \models_x \phi_1 \,\mathsf{R}\, \phi_2$ iff either *(i)* for all $i \in \mathbb{N}$, $(\mathcal{G}, \rho_{\geq i}) \models_x \phi_2$, or *(ii)* there exists $i \in \mathbb{N}$ such that $(\mathcal{G}, \rho_{\geq i}) \models_x \phi_1 \wedge \phi_2$ and for all $j \in \mathbb{N}$ where $j < i$, $(\mathcal{G}, \rho_{\geq j}) \models_x \phi_2$.

For an ATL-SA formula ϕ, a CGS-SA $\mathcal{G}$, and an infinite path ρ, the satisfaction relation $(\mathcal{G}, \rho) \models_x \phi$ depends only on the first state of ρ. Therefore, we may simply write $(\mathcal{G}, s) \models_x \phi$ as shorthand for $(\mathcal{G}, \rho) \models_x \phi$, where $s = \rho[1]$. Moreover, when ϕ is an ATL formula and $\mathcal{G}$ is a CGS, the relation $\models_x$ coincides with the semantics of ATL for any $x \in \{\mathrm{d}, \mathrm{u}\}$. We explicitly denote this as $\models_{\mathrm{ATL}}$ when the formula is in ATL and $\mathcal{G}$ is a CGS. It is well known that the ATL model-checking problem is PTIME [3].

Example 3.4. Following from the arm wrestling game $\mathcal{G}$ described in Example 2.2, we consider the formula $\phi = \langle 1\rangle^{\geq 0.37}\mathsf{F}(win)$ from Example 3.2. The formula ϕ holds true (in both u- and d-semantics, since there is a single agent in the coalition) if and only if agent 1 has a strategy to win with a probability greater than 0.37. We can intuit that a good strategy for agent 1 is to first intimidate, and if that does not work, to compete using either the right or left hand depending on whether agent 1 has the capacity c_1^r or c_1^l, respectively. We denote this strategy as σ_1 (with domain $\{1\}$). Let σ_2 represent a strategy assignment with domain $\{2\}$. If agent 2 is a coward or has a different handedness than agent 1, the outcome of $\sigma_1 \oplus \sigma_2$ starting from state S necessarily reaches the state win. Thus, we have $(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(S, \sigma_1 \oplus \sigma_2)(\kappa)) \models_x \phi$ under the sufficient condition that $\kappa \notin \{\kappa_{rr}, \kappa_{ll}\}$, where $\kappa_{rr}(1) = c_1^r$, $\kappa_{rr}(2) = c_2^r$, $\kappa_{ll}(1) = c_1^l$, and $\kappa_{ll}(2) = c_2^l$. Therefore, we have $\mathbb{P}_{\mathcal{G}}\left[(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(S, \sigma_1 \oplus \sigma_2)) \models_x \phi\right] \geq 1 - \mathbb{P}_{\mathcal{G}}[\{\kappa_{rr}, \kappa_{ll}\}] = 1 - (0.8 \times 0.75) - (0.2 \times 0.15) = 0.37$. Finally, we conclude that $(\mathcal{G}, I) \models_x \langle 1\rangle^{\geq 0.37}\mathsf{F}(win)$ for both $x \in \{\mathrm{d}, \mathrm{u}\}$.

The distributed semantics allows agents within the strategic coalition to share their capacities, enhancing their strategy space. Notably, since a uniform strategy assignment is a distributed strategy assignment, we establish Proposition 3.5.

Proposition 3.5. *Let $\mathcal{G}$ be a CGS-SA, $Y \subseteq \mathrm{Ag}$, $s \in \mathrm{St}$, $p \in [0,1]$, $\bowtie \in \{\leq, <, >, \geq\}$, and ψ be a temporal formula without a strategic operator. We have $(\mathcal{G}, s) \models_{\mathrm{u}} \langle Y\rangle^{\bowtie p}\psi$ implies $(\mathcal{G}, s) \models_{\mathrm{d}} \langle Y\rangle^{\bowtie p}\psi$.*

Having established a formal semantics for ATL-SA, we now turn our attention to model checking, which involves determining whether a given model satisfies the specified formula.

4 MODEL CHECKING

This section presents ATL-SA model checking for both uniform and distributed semantics. For $x \in \{\mathrm{d}, \mathrm{u}\}$, we denote by Atl-Sa-MC$_x$ the ATL-SA model-checking problem with x-semantics. It takes as input a CGS-SA $\mathcal{G}$, an ATL-SA formula ϕ over the same signature, and a state $s \in \mathrm{St}$, and returns whether $(\mathcal{G}, s) \models_x \phi$. Moreover, we denote by $\langle\cdot\rangle$-Atl-Sa-MC$_x$ the restriction of Atl-Sa-MC$_x$ to formulae of the form $\phi = \langle Y\rangle^{\bowtie p}\psi$ where ψ does not have strategic operators. We begin by establishing the NEXPTIME-completeness of $\langle\cdot\rangle$-Atl-Sa-MC$_x$. Then, we deduce the $\mathsf{P}^{\mathsf{NEXPTIME}}$-membership of Atl-Sa-MC$_x$ and show a PTIME-complete restriction.

Proposition 4.1 demonstrates that we only need to consider $\bowtie \in \{\geq, >\}$ in the semantics, as equivalent formulae can be obtained for other relations. Specifically, for $\bowtie \in \{\leq, <, >, \geq\}$, let $\overline{\bowtie} \in \{\leq, <, >, \geq\}$ be the symbol such that $p_1 \bowtie p_2$ iff $p_2 \,\overline{\bowtie}\, p_1$. Moreover, let $\overline{\psi}$ denote the negation of a temporal formula ψ, i.e., $\overline{\mathsf{X}\,\phi} = \mathsf{X}(\neg\phi)$, $\overline{\phi_1 \,\mathsf{U}\, \phi_2} = (\neg\phi_1)\,\mathsf{R}\,(\neg\phi_2)$, and $\overline{\phi_1 \,\mathsf{R}\, \phi_2} = (\neg\phi_1)\,\mathsf{U}\,(\neg\phi_2)$.

Proposition 4.1. *Let $\mathcal{G}$ be a CGS-SA, Y be a coalition, $s \in \mathrm{St}$, $x \in \{\mathrm{d}, \mathrm{u}\}$, $p \in [0,1]$, $\bowtie \in \{\leq, <, >, \geq\}$, and let ψ be a temporal ATL-SA formula. We have:*

$$(\mathcal{G}, s) \models_x \langle Y\rangle^{\bowtie p}\psi \iff (\mathcal{G}, s) \models_x \langle Y\rangle^{\overline{\bowtie}(1-p)}\overline{\psi}$$

Proof. For all infinite path ρ and temporal formula ψ, we have $(\mathcal{G}, \rho) \not\models_x \psi$ iff $(\mathcal{G}, \rho) \models_x \overline{\psi}$. So, for all complete strategy assignment σ and state s, we have $\mathbb{P}_{\mathcal{G}}\left[(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(s, \sigma)) \models_x \psi\right] \bowtie p$ iff $1 - p \bowtie \mathbb{P}_{\mathcal{G}}\left[(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(s, \sigma)) \not\models_x \psi\right]$. □

For the model-checking procedure, we focus on particular sets of complete capacity assignments called *rectangular* complete capacity

[2] Using a common shortcut from probability theory, for a complete strategy assignment σ and a state s in a CGS-SA $\mathcal{G}$, we let $(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(s, \sigma)) \models_x \psi$ stand for $\{\kappa \in \Gamma_{\mathcal{G}} \mid (\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(s, \sigma)(\kappa)) \models_x \psi\}$.

assignment sets. Given a list of capacities in $\mathcal{G}$ for each agent $(C_1, \dots, C_n) \in \mathcal{P}(\Delta_1^{>0}) \times \dots \times \mathcal{P}(\Delta_n^{>0})$, let $R_{\mathcal{G}}(C_1, \dots, C_n)$ denote the rectangular complete capacity assignment set, where each agent a has a capacity in C_a. Let $\mathbf{R}_{\mathcal{G}}$ denote the set of all such rectangular complete capacity assignments. Given a CGS-SA $\mathcal{G}$, we define the CGS $\mathcal{T}_{\mathcal{G}} = \langle \mathrm{Ag}', \mathrm{St}', \Pi', \pi', \mathrm{Ac}', d', o' \rangle$ as follows:

- $\mathrm{Ag}' = \mathrm{Ag} \cup \{n+1\}$, where agent $n+1$ decides the active capacity assignment in each state.
- $\mathrm{St}' = \mathrm{St} \times \mathbf{R}_{\mathcal{G}}$. A state $(s, K) \in \mathrm{St}'$ means that $\mathcal{G}$ is in state s and the feasible capacity assignments are within K.
- $\Pi' = \Pi \cup \Gamma_{\mathcal{G}}$.
- For $q = (s, K) \in \mathrm{St}'$, set $\pi'(q) = \pi(s) \cup K$.
- Ac' includes Ac, $\Gamma_{\mathcal{G}}$, and partial functions $f : \mathcal{P}(\mathrm{Ac}) \rightharpoonup \mathrm{Ac}$.
- For $q = (s, K) \in \mathrm{St}'$, the protocol for agent $a \in \mathrm{Ag}$, $d'(a, q)$, consists of functions f defined on $\{\kappa(a) \mid \kappa \in K\}$ and such that for each $\kappa \in K$, we have $f(\kappa(a)) \in d(a, s) \cap \kappa(a)$. Agent $n+1$ selects a complete capacity assignment: $d'(n+1, q) = K$.
- For $q = (s, K) \in \mathrm{St}'$ and $(f_1, \dots, f_n, \kappa) \in d'(1, q) \times \dots \times d'(n+1, q)$, we have $o'(q, f_1, \dots, f_n, \kappa) = (s', K')$, where $s' = o(s, f_1(\kappa(1)), \dots, f_n(\kappa(n)))$ and $K' = \{\kappa' \in K \mid \forall a \in \mathrm{Ag}, f_a(\kappa'(a)) = f_a(\kappa(a))\}$. Notice that, indeed, $K' \in \mathbf{R}_{\mathcal{G}}$. Verbally expressed, the game progresses with the actions for Ag of the capacity assignment chosen by agent $n+1$ and the new set of possible capacity assignments contains those where all agents would have chosen the same action.

Let ϕ_1 and ϕ_2 be two propositional formulae, $K \subseteq \Gamma_{\mathcal{G}}$, and $\phi_K = \bigvee_{\kappa \in K} \kappa$. We define the function g as follows:

$$g(\psi, K) = \begin{cases} \mathsf{X}(\phi_K \to \phi_1) & \text{if } \psi = \mathsf{X}\,\phi_1, \\ \phi_1 \,\mathsf{U}\, (\phi_K \to \phi_2) & \text{if } \psi = \phi_1 \,\mathsf{U}\, \phi_2, \\ (\phi_K \to \phi_1) \,\mathsf{R}\, \phi_2 & \text{if } \psi = \phi_1 \,\mathsf{R}\, \phi_2. \end{cases}$$

Given a set of complete capacity assignments $K \subseteq \Gamma_{\mathcal{G}}$, Proposition 4.2 shows that ensuring strategic objectives with u-semantics for a coalition Y is equivalent to solving an ATL formula on $\mathcal{T}_{\mathcal{G}}$.

Proposition 4.2. *Let $\mathcal{G}$ be a CGS-SA, $s \in \mathrm{St}$, $Y \subseteq \mathrm{Ag}$, $K \subseteq \Gamma_{\mathcal{G}}$, and ψ be a temporal formula without a strategic operator. The following two propositions are equivalent:*

(i) There exists a uniform strategy assignment σ_Y for Y such that, for every strategy assignment $\sigma_{\mathrm{Ag}\setminus Y}$ for $\mathrm{Ag} \setminus Y$ and for all $\kappa \in K$, we have $(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(s, \sigma_Y \oplus \sigma_{\mathrm{Ag}\setminus Y})(\kappa)) \models_{\mathrm{u}} \psi$.
(ii) $(\mathcal{T}_{\mathcal{G}}, (s, \Gamma_{\mathcal{G}})) \models_{\mathrm{ATL}} \langle Y \rangle g(\psi, K)$.

Proof. Let $m : \mathrm{Pt}_{\mathcal{T}_{\mathcal{G}}} \to \mathrm{Pt}_{\mathcal{G}}$ be a function such that, for any path $\eta = (s_1, K_1)(f_1^1, \dots, f_1^n, \kappa_1)(s_2, K_2)(f_2^1, \dots, f_2^n, \kappa_2) \dots$ in $\mathcal{T}_{\mathcal{G}}$, we have $m(\eta) = s_1 \vec{\alpha}_1 s_2 \vec{\alpha}_2 \dots$, where $\vec{\alpha}_i = (f_i^1(\kappa_i(1)), \dots, f_i^n(\kappa_i(n)))$ for all $i \in \mathbb{N}$. Notice that m is a surjection. Let κ_0 be the unique complete capacity assignment in the CGS $\mathcal{T}_{\mathcal{G}}$. For a capacity-uniform strategy assignment σ for Y in $\mathcal{G}$, we define the strategy assignment $t(\sigma) = \sigma'$ for Y in $\mathcal{T}_{\mathcal{G}}$, such that, for all $a \in \mathrm{Ag}$ and $\eta \in \mathrm{Pt}_{\mathcal{T}_{\mathcal{G}}}^{<\omega}$ with $\mathsf{last}(\eta) = (s, K)$, we let $\sigma'(a)(\kappa_0, \eta) = f$, where for all $\kappa \in K$, we set $f(\kappa(a)) = \sigma(a)(\kappa, m(\eta))$. This is well-defined because $\sigma(a)(\kappa, m(\eta))$ does not depend on the values $\kappa(b)$ with $b \neq a$ by capacity-uniformity. The function t is surjective.

Suppose case (i) from Proposition 4.2 holds with a winning strategy σ_Y and $\psi = \phi_1 \,\mathsf{U}\, \phi_2$. Let σ'_Y be such that $t(\sigma'_Y) = \sigma_Y$ (remember the surjectivity), and let $\sigma'_{\mathrm{Ag}\setminus Y}$ be a strategy for $\mathrm{Ag} \setminus Y$ in $\mathcal{T}_{\mathcal{G}}$. Consider the path $\eta = \mathrm{Out}_{\mathcal{T}_{\mathcal{G}}}((s, \Gamma_{\mathcal{G}}), \sigma'_Y \oplus \sigma'_{\mathrm{Ag}\setminus Y})(\kappa_0)$. If $(\mathcal{T}_{\mathcal{G}}, \eta) \not\models_{\mathrm{ATL}} \phi_1$, then $(\mathcal{G}, m(\eta)) \not\models \phi_1$, so $(\mathcal{G}, m(\eta)) \models \phi_2$. Thus, $(\mathcal{T}_{\mathcal{G}}, \eta) \models_{\mathrm{ATL}} \phi_K \to \phi_2$, implying $(\mathcal{T}_{\mathcal{G}}, \eta) \models_{\mathrm{ATL}} g(\psi, K)$. Otherwise, let $i \in \mathbb{N}$ be the smallest index such that $(\mathcal{T}_{\mathcal{G}}, \rho_{\geq i}) \not\models_{\mathrm{ATL}} \phi_1$, and let $(s_i, K_i) = \eta[i]$. If $K \cap K_i = \emptyset$, then $(\mathcal{T}_{\mathcal{G}}, \rho_{\geq i}) \models_{\mathrm{ATL}} \neg\phi_\kappa$, thus $(\mathcal{T}_{\mathcal{G}}, \eta) \models_{\mathrm{ATL}} g(\psi, K)$. Otherwise, let $\kappa \in K \cap K_i$, and define $\sigma_{\mathrm{Ag}\setminus Y} = t(\sigma'_{\mathrm{Ag}\setminus Y})$. The path $\rho = \mathrm{Out}_{\mathcal{G}}(s, \sigma_Y \oplus \sigma_{\mathrm{Ag}\setminus Y})(\kappa)$ satisfies $\rho_{\leq i} = m(\eta)_{\leq i}$, $(\mathcal{G}, \rho_{\geq i}) \not\models \phi_1$, and, for all $j < i$, $(\mathcal{G}, \rho_{\geq j}) \models \phi_1$. Hence, $(\mathcal{G}, \rho_{\geq i}) \models \phi_2$ and $(\mathcal{T}_{\mathcal{G}}, \eta) \models_{\mathrm{ATL}} g(\psi, K)$. Finally, we conclude that $(\mathcal{T}_{\mathcal{G}}, (s, \Gamma_{\mathcal{G}})) \models_{\mathrm{ATL}} \langle Y \rangle g(\psi, K)$. The cases for $\psi = \phi_1 \,\mathsf{R}\, \phi_2$ and $\psi = \mathsf{X}\,\phi_1$ are analogous. Conversely, suppose $(\mathcal{T}_{\mathcal{G}}, (s, \Gamma_{\mathcal{G}})) \models_{\mathrm{ATL}} \langle Y \rangle g(\psi, K)$ with a winning strategy σ'_Y. Then, $t(\sigma'_Y)$ satisfies item (i) of Proposition 4.2. □

Based on this equivalence, we deduce the NEXPTIME-membership of $\langle\cdot\rangle$-Atl-Sa-MC$_x$ for $x \in \{\mathrm{d}, \mathrm{u}\}$.

Proposition 4.3. *For $x \in \{\mathrm{d}, \mathrm{u}\}$, the problem $\langle\cdot\rangle$-Atl-Sa-MC$_x$ is in NEXPTIME.*

Proof. Let $\mathcal{G}$, s, and $\phi = \langle Y \rangle^{\bowtie p} \psi$ be a positive instance of $\langle\cdot\rangle$-Atl-Sa-MC$_\mathrm{u}$, *i.e.*, $(\mathcal{G}, s) \models_{\mathrm{u}} \phi$. By Proposition 4.1, we can assume $\bowtie \in \{\geq, >\}$. By the semantics, $(\mathcal{G}, s) \models_{\mathrm{u}} \phi$ iff there is $K \subseteq \Gamma_{\mathcal{G}}$ and a distributed statregy assignment σ_Y for Y such that $\mathbb{P}_{\mathcal{G}}[K] \bowtie p$ and, for all strategy assignment $\sigma_{\mathrm{Ag}\setminus Y}$ for $\mathrm{Ag} \setminus Y$ and all $\kappa \in K$, we have $(\mathcal{G}, \mathrm{Out}_{\mathcal{G}}(s, \sigma_Y \oplus \sigma_{\mathrm{Ag}\setminus Y})(\kappa)) \models_{\mathrm{d}} \psi$. Such a K is the non-deterministic exponential certificate of $(\mathcal{G}, s) \models_{\mathrm{u}} \phi$. By Proposition 4.2, verifying $(\mathcal{G}, s) \models_{\mathrm{u}} \phi$ with the help of K boils down to verifying $(\mathcal{T}_{\mathcal{G}}, (s, \Gamma_{\mathcal{G}})) \models_{\mathrm{ATL}} \langle Y \rangle g(\psi, K)$. The number of states in $\mathcal{T}_{\mathcal{G}}$ is bounded by $|\mathrm{St}| \cdot |\mathbf{R}_{\mathcal{G}}| = |\mathrm{St}| \cdot 2^{\sum_{a \in \mathrm{Ag}} |\Delta_a^{>0}|}$. The size of $\mathcal{T}_{\mathcal{G}}$ is determined by its number of transitions, which is less than $|\Gamma_{\mathcal{G}}| \cdot \prod_{a \in \mathrm{Ag}} |\mathrm{Ac}|^{|\Delta_a^{>0}|}$ for each state. Therefore, the overall size of $\mathcal{T}_{\mathcal{G}}$ is less than $|\mathrm{St}|^2 \cdot |\mathrm{Ac}|^{\sum_{a \in \mathrm{Ag}} |\Delta_a^{>0}|} \cdot 2^{\sum_{a \in \mathrm{Ag}} |\Delta_a^{>0}|} \leq 2^{P(|\mathcal{G}|)}$ for some polynomial P. Moreover, the size of $g(\psi, K)$ is $O(|\psi| + |K|) = O(|\psi| + 2^{|\mathcal{G}|})$. Since ATL model checking is known to be polynomial [3], we check $(\mathcal{G}, s) \models_{\mathrm{u}} \phi$ in EXPTIME with the help of K. Finally, $\langle\cdot\rangle$-Atl-Sa-MC$_\mathrm{u}$ is in NEXPTIME.

The distributed semantics can be simulated by merging agents in the strategic coalition and using the uniform case. Precisely, to verify $\langle Y \rangle^{\bowtie p} \psi$, we merge the agents in Y into a single agent a_Y, whose capacities correspond to the possible tuples of capacities for each agent in Y. The actions of a_Y are the tuples of actions from the agents in Y. Let $\mathcal{G}_Y$ denote the CGS-SA after merging Y. Although the CGS-SA $\mathcal{G}_Y$ may exhibit an exponentially larger number of capacities, the number of complete capacity assignments in both $\mathcal{G}$ and $\mathcal{G}_Y$ remains the same. Additionally, the size of $\mathcal{T}_{\mathcal{G}}$ is equal to that of $\mathcal{T}_{\mathcal{G}_Y}$. Using the same procedure as the uniform case, we conclude that $\langle\cdot\rangle$-Atl-Sa-MC$_\mathrm{d}$ is in NEXPTIME. □

Remark 1. *We could extend the construction to handle $\bowtie \in \{=\}$. For instance, suppose $\phi = \langle Y \rangle^{=p} \phi_1 \,\mathsf{U}\, \phi_2$, and let $s \in \mathrm{St}$. For $K \subseteq \Gamma_{\mathcal{G}}$, define $\phi_K = \bigvee_{\kappa \in K} \kappa$ and $\phi_{\neg K} = \bigvee_{\kappa \in \Gamma_{\mathcal{G}} \setminus K} \kappa$. We have: $(\mathcal{G}, s) \models_{\mathrm{u}} \phi$ if and only if $(\mathcal{T}_{\mathcal{G}}, (s, \Gamma_{\mathcal{G}})) \models_{\mathrm{ATL}} \phi'$ where*

$$\phi' = \langle Y \rangle \left(\phi_1 \wedge (\phi_{\neg K} \to \neg\phi_2) \,\mathsf{U}\, ((\phi_K \to \phi_2) \wedge (\phi_{\neg K} \to \neg\phi_2)) \right)$$

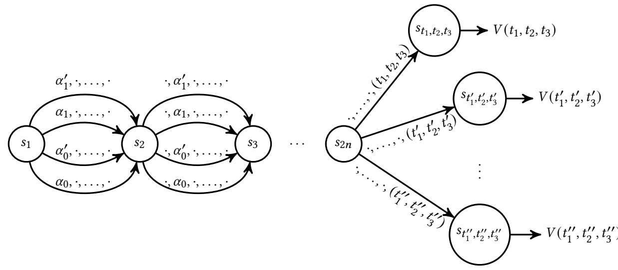

Figure 2: CGS-SA for the NEXPTIME-completeness proof.

for some $K \subseteq \Gamma_{\mathcal{G}}$ such that $\mathbb{P}_{\mathcal{G}}[K] = p$. With similar formula transformations for R *and* X*, we maintain an NEXPTIME model-checking algorithm for the extended logic where $\bowtie$ can be in $\{\leq, <, =, >, \geq\}$.*

In some settings, allowing agents in the strategic coalition Y to share private information (i.e., moving from uniform semantics to distributed semantics) simplifies the model checking [16]. However, in ATL-SA, the cost of finding uniform strategy assignments with respect to distributed ones is negligible when compared to the cost of finding a maximal-probability subset of complete capacity assignment $K \subseteq \Gamma_{\mathcal{G}}$ that cannot prevent Y from achieving their objective. Indeed, Theorem 4.4 demonstrates that $\langle\cdot\rangle$-Atl-Sa-MC$_x$ is NEXPTIME-complete for $x \in \{\mathrm{d}, \mathrm{u}\}$.

We reduce the tiling problem from [11], which is known to be NEXPTIME-complete [11, 27], to $\langle\cdot\rangle$-Atl-Sa-MC$_x$ with a single agent in the coalition (so that uniform and distributed semantics are equivalent). An instance of the NEXPTIME-tiling problem is given by (T, t^*, m), where T is a set of tile types, $t^* \in T$, and $m \in \mathbb{N}$ (written in binary). A tile type $t \in T$ is a tuple of four colors, $t = (\mathsf{left}(t), \mathsf{right}(t), \mathsf{up}(t), \mathsf{down}(t))$. The output of a tiling problem instance (T, t^*, m) is *yes* if the $m \times m$ plane can be tiled with tiles of types in T and with t^* at the plane's origin, and *no* otherwise. Formally, the output is *yes* if there exists a function $\tau : \{0, \ldots, m-1\}^2 \to T$, called a tiling, that satisfies the following conditions:

(i) $\tau(0,0) = t^*$,
(ii) $\mathsf{up}(\tau(x,y)) = \mathsf{down}(\tau(x,y+1))$ for all $x \in \{0,\ldots,m-1\}$ and $y \in \{0,\ldots,m-2\}$, and
(iii) $\mathsf{right}(\tau(x,y)) = \mathsf{left}(\tau(x+1,y))$ for all $x \in \{0,\ldots,m-2\}$ and $y \in \{0,\ldots,m-1\}$.

The main completeness result follows.

Theorem 4.4. *The problem $\langle\cdot\rangle$-Atl-Sa-MC$_x$ is NEXPTIME-complete for both uniform ($x = \mathrm{u}$) and distributed ($x = \mathrm{d}$) semantics.*

Proof. Proposition 4.3 establishes the NEXPTIME-membership. Conversely, we reduce the tiling problem to $\langle\cdot\rangle$-Atl-Sa-MC$_x$ (for both $x \in \{\mathrm{d}, \mathrm{u}\}$). Consider a tiling instance (T, t^*, m) and assume, without loss of generality, that $m = 2^n$. The goal is to construct a CGS-SA $\mathcal{G}$ with a specific state s_1 and an ATL-SA formula $\phi = \langle Y\rangle^{\bowtie p}\psi$ such that $(\mathcal{G}, s_1) \models_x \phi$ iff there exists a tiling for (T, t^*, m). The structure of the CGS-SA is outlined in Figure 2. For the agents in $\mathcal{G}$, we denote them as $X_1, \ldots, X_n, Y_1, \ldots, Y_n, Q, P$ (from 1 to $2n$+2). Let $Z = \{X_1, \ldots, X_n, Y_1, \ldots, Y_n\}$. Agents in Z possess four equiprobable capacities: $c_0 = \{\beta_0, \alpha'_0, \alpha_1, \alpha'_1, \cdot\}$, $c'_0 = \{\beta'_0, \alpha_0, \alpha_1, \alpha'_1, \cdot\}$, $c_1 = \{\beta_1, \alpha_0, \alpha'_0, \alpha'_1, \cdot\}$, and $c'_1 = \{\beta'_1, \alpha_0, \alpha'_0, \alpha_1, \cdot\}$. The action $\cdot$ indicates that the agent performs no action. Note that β'_0 can be used by all capacities except c'_0, and β_0 can only be used by c'_0, with a similar distinction for the other actions. When agent X_i (resp. Y_i) has a capacity in $\{c_0, c'_0\}$, it means that the i^{th} bit of the binary representation of a position x (resp. y) is 0. Conversely, if the capacity is in $\{c_1, c'_1\}$, then the bit is 1. Thus, the capacity assignments of Z encode a position (x, y) on the grid, where we denote x_i and y_i as the i^{th} bit of x and y. Agent Q has l equiprobable capacities represented by $c_t = \{t, \cdot\}$ for each $t \in T$. Agent P, the *picker*, possesses a single capacity $c = \{\alpha_0, \alpha'_0, \alpha_1, \alpha'_1, n\} \cup \{(t_1, t_2, t_3) \in T^3 \mid \mathsf{up}(t_1) = \mathsf{down}(t_3) \wedge \mathsf{right}(t_1) = \mathsf{left}(t_2)\}$, and acts as the strategic agent attempting to prove the existence of a tiling.

The first phase is called the *challenge* phase (from state s_1 to s_{2n} in Figure 2), during which agents in Z commit to a challenge position (x^c, y^c) of their choice by revealing one of the capacities from $\{c_0, c'_0, c_1, c'_1\}$ that each of these agents do not possess. For example, agent X_i will use action α_1 or α'_1 (at least one is available to X_i) to indicate that the bit x_i^c is 1.

In the subsequent *pick* phase, agent P must select three tiles t_1, t_2, and t_3, which P claims are located at positions (x^c, y^c), $(x^c + 1, y^c)$, and $(x^c, y^c + 1)$ of a solution to the tiling problem. By the construction of P's action set, t_1, t_2, and t_3 adhere to the tiling adjacency conditions *(ii)* and *(iii)*.

Finally, there is a *verification* phase $V(t_1, t_2, t_3)$ for each tuple picked by P, where agent P loses if any of the following conditions are met: *(i)* $(x^c, y^c) = (0, 0)$ and $t_1 \neq t^*$ *(ii)* $(x, y) = (x^c, y^c)$ and Q does not have capacity c_{t_1}, *(iii)* $(x, y) = (x^c + 1, y^c)$ and Q does not have capacity c_{t_2}, or *(iv)* $(x, y) = (x^c, y^c + 1)$ and Q does not have capacity c_{t_3}. Notice that, verifying $z = z^c + 1$ can be accomplished by finding $i \in \{1, \ldots, n\}$ such that: $z_i = 1$, $z_i^c = 0$, for all $j < i$, $z_j = z_j^c$, and for all $j > i$, $z_j = 0$ and $z_j^c = 1$. So, during the verification phase $V(t_1, t_2, t_3)$, the opponents (*e.g.*, agent Q) decide which of the four verifications to perform. For verifications involving $x^c + 1$ or $y^c + 1$, Q determines which cutoff i to utilize (totaling $2n + 2$ options). The verification requires checking equality (or difference) between x_i (or y_i), x_i^c (or y_i^c), 1, and 0 for each $i \in \{1, \ldots, n\}$, as well as between the capacities of Q, c_{t_1}, c_{t_2}, c_{t_3}, and c_{t^*} (resulting in $2n + 1$ checks for each of the $2n + 2$ options). Each of these equality or difference checks can be conducted with a fixed-size game structure or a size of $O(l)$ in the case of verifying the tile types. For instance, to test the equality involving x_i, x_i^c, 0, and 1 in a state s, we allow the actions available for P to be $\{\alpha_0, \alpha'_0, \alpha_1, \alpha'_1\}$, the actions for X_i to be $\{\beta_0, \beta'_0, \beta_1, \beta'_1\}$, and other agents are idle. We force agent P to repeat the action used by agent X_i during the challenge phase through transitions of the form $s\vec{\alpha}L$ where $(\vec{\alpha}[X_i], \vec{\alpha}[P]) \in \{(\beta_0, \alpha_0), (\beta'_0, \alpha'_0), (\beta_1, \alpha_1), (\beta'_1, \alpha'_1)\}$ and L is a sink losing state for P. This allows us to easily verify the outcome of our test on x_i, x_i^c, 1, and 0. Finally, the size of the CGS-SA to encode the verification $V(t_1, t_2, t_3)$ is $O(n^2 \cdot l)$, and the overall verification size is $O(n^2 \cdot l^4)$, which is polynomial in the input of the tiling instance. Thus, the entire CGS-SA is also polynomial in size.

Consider the formula $\phi = \langle P\rangle^{\geq 1/l}\mathsf{F}win$, where win is a label in the states such that the verification phase succeeds. Each strategy of P can win against at most one capacity of Q. Therefore, $\mathsf{F}win$ is achieved with a probability of at most $1/l$, implying that ϕ holds if and only if P's strategy wins against a capacity of Q for each capacity assignment of Z. Suppose there exists a tiling τ. Given the challenge (x^c, y^c) from the challenge phase, P can choose $t_1 = \tau(x^c, y^c)$, $t_2 = \tau(x^c+1, y^c)$, and $t_3 = \tau(x^c, y^c+1)$, succeeding in the verification phase whenever Q has capacity $c_{\tau(x,y)}$, where (x, y) is the coordinate encoded by the capacities of Z. This occurs with a probability of $1/l$, so ϕ holds.

Conversely, suppose there is no tiling. If P does not pick $t_1 = t^*$ when the challenge is $(0, 0)$, then P loses against all capacities of Q when the capacities of Z encode $(x, y) = (0, 0)$ and Z gives the challenge $(0, 0)$. (Note that there are several ways to encode a single challenge, so we mean "Z gives the challenge $(0, 0)$ with that particular encoding" here and later.) In this case, ϕ cannot be verified. Otherwise, there must be two challenges (x^c, y^c) and (x'^c, y'^c) such that: *(i)* $x'^c = x^c + 1$ and $y'^c = y^c$ (or $x'^c = x^c$ and $y'^c = y^c + 1$, which case is treated similarly and omitted here), *(ii)* P picks (t_1, t_2, t_3) for the challenge (x^c, y^c), *(iii)* P picks (t'_1, t'_2, t'_3) for the challenge (x'^c, y'^c), and *(iv)* $t'_1 \neq t_2$. Now, assume Z encodes $(x, y) = (x'^c, y'^c)$. If Q encodes neither t'_1 nor t_2, then P loses as usual. When Q encodes t'_1, Z can present the challenge (x^c, y^c), leading to a loss for P since it chooses t_2, which differs from t'_1. Similarly, if Q encodes t_2, Z can present the challenge (x'^c, y'^c), and P loses by choosing t'_1, which differs from t_2. Consequently, P loses against all capacities of Q when Z encodes (x'^c, y'^c), thus ϕ is not satisfied. □

The previous results about $\langle\cdot\rangle$-ATL-SA-MC$_x$ consider formulae without nested strategic operators. Theorem 4.5 states the $\mathsf{P}^{\mathsf{NEXPTIME}}$ upper complexity bound (polynomial with NEXPTIME oracle) and NEXPTIME lower bound in the general case.

THEOREM 4.5. *For $x \in \{\mathrm{d}, \mathrm{u}\}$, the problem ATL-SA-MC$_x$ is in $\mathsf{P}^{\mathsf{NEXPTIME}}$ and is NEXPTIME-hard.*

PROOF. Starting from innermost strategic subformulae ϕ, we label the states s with a new atomic proposition ℓ_ϕ iff $(\mathcal{G}, s) \models_x \phi$ and replace the subformulae ϕ by ℓ_ϕ. □

Sometimes, the number of agents and their capacities are small or fixed relative to the size of the game structure. Thus, it is worthwhile to study the complexity of ATL-SA model checking when the number of complete capacity assignments grows exponentially slower than the number of transitions in a CGS-SA.

PROPOSITION 4.6. *For any function $f : \mathbb{N} \to \mathbb{N}$ in $O(\log n)$ and $x \in \{\mathrm{d}, \mathrm{u}\}$, the problem ATL-SA-MC$_x$, restricted to CGS-SA such that $|\Gamma_\mathcal{G}| \leq f(|\mathrm{Pt}^{\langle 2 \rangle}_\mathcal{G}|)$ is PTIME-complete.*

PROOF. Given that $|\Gamma_\mathcal{G}| \leq f(|\mathrm{Pt}^{\langle 2 \rangle}_\mathcal{G}|)$, the certificates K (from the procedure of Proposition 4.3) can be iterated over in polynomial time and the size of $\mathcal{T}_\mathcal{G}$ is polynomial in $|\mathcal{G}|$. This gives the PTIME membership, and the hardness derives from the PTIME-hardness of *Computation Tree Logic* (CTL) which is included in ATL-SA with one agent and one capacity [24]. □

In the next section, we provide evidence of the applicability of ATL-SA in practice and demonstrate why the properties expressible in ATL-SA can be of interest.

5 CYBERSECURITY ILLUSTRATION

A team of cybersecurity experts aims to identify the best strategy to defend their industrial system. They rely on *Moving Target Defense* (MTD), a defense paradigm that promotes regularly changing the system configuration to increase security. The experts also possess a set of probes that notify them of attacker activities, as well as knowledge of their system's vulnerabilities and potential attack methods.

System Model. The system is characterized by a set of configuration parameters $p_1, \ldots, p_n$, and a set of attack subgoals $g_1, \ldots, g_k$. Each parameter p_i can be configured to one of the values $v_i^1, \ldots, v_i^{k_i}$. Each attack subgoal is either active or inactive. MTDs can be triggered to modify a configuration parameter. If a configuration parameter p_i has a value v_i^j, then the MTD is defined by the set of accessible configuration values $V(p_i, v_i^j) \subseteq \{v_i^1, \ldots, v_i^{k_i}\}$. Moreover, the security experts are aware of a set of attacker exploits $\{e_1, \ldots, e_t\}$, where each exploit e is defined as a tuple $e = (\chi_e, pre_e, post_e)$. Here, χ_e is a partial configuration parameter valuation, and $pre_e, post_e \subseteq \{g_1, \ldots, g_k\}$ represent the attack subgoals necessary before the exploit and obtained after the attack success, respectively.

Modeling the Attacker-Defender Interaction. The interaction between the attacker and the defender can be modeled as a CGS-SA with two agents: the defender $D = 1$ and the attacker $A = 2$. The set of states St takes the form (χ, G, x), where:

- $\chi(p_i) \in \{v_i^1, \ldots, v_i^{k_i}\}$ is a complete configuration parameter valuation,
- $G \subseteq \{g_1, \ldots, g_k\}$ is the set of active attack subgoals,
- $x \in \{D, A\}$ indicates whose turn it is, defender or attacker.

From a state (χ, G, A), there is an outgoing transition controlled by agent D for each possible MTD activation. For example, an MTD activation that reconfigures p_i to value v_i^j leads to a new state $(\chi[p_i \mapsto v_i^j], G, A)$. The defender can also choose the action *nothing*, which indicates no reconfiguration, thus passing the turn to the attacker and transitioning to the state (χ, G, D). From state (χ, G, D), the attacker can execute one of the exploits $e = (\chi_e, pre_e, post_e)$, provided that χ_e is a subfunction of χ and $pre_e \subseteq G$. This results in a new state $(\chi, G \cup post_e, D)$. We assume that the defender has a single profile, *i.e.*, $\Delta_D(\mathrm{Ac}) = 1$. However, there exists a set of attacker profiles with associated probabilities. For example, if 90% of attackers are novices and can only perform exploits e_1, e_2, and e_3, we set $\Delta_A(\{e_1, e_2, e_3\}) = 0.9$.

System Objectives. Each state (χ, G, x) is labeled with $q_{p,v}$ whenever $\chi(p) = v$, and with the set of labels G representing active attack subgoals. The primary objective is to prevent severe system compromise, which is defined by a propositional formula ϕ_g over $\{g_1, \ldots, g_k\}$, *e.g.*, $\phi_g = (g_3 \wedge g_4) \vee g_7$. Additionally, we need to avoid specific undesirable configurations, represented by a propositional formula ϕ_p over $\{q_{p_1, v_1^1}, \ldots, q_{p_n, v_n^{k_n}}\}$, for instance, $\phi_p = \neg(q_{p_3, v_2} \wedge q_{p_4, v_1})$. These bad configurations could, for example, lead to poor service quality for regular users. The defense team's goal can be expressed using the formula: $\phi = \langle D \rangle^{\geq 0.98}\, \mathrm{G}(\neg\phi_g \wedge \neg\phi_p)$, *i.e.*, find a strategy for the defender that guarantees, with at least

98% probability, that the system will always avoid ϕ_g (system compromise) and ϕ_p (bad configurations).

6 RELATED WORK

First, this section relates ATL-SA to *Capacity Alternating-time Temporal Logic* (CapATL) [5]. Then, it demonstrates the novelty of our probabilistic approach in comparison to the two notions of stochasticity commonly found in the literature on strategic verification in MAS: *stochastic CGS* and *stochastic strategies*.

CapATL. The concept of capacity is shared between ATL-SA and CapATL: each agent has a set of capacities corresponding to a subset of actions the agent can perform. However, there are two significant differences. First, CapATL assumes that the actions of other agents are indistinguishable. For example, the paths $s(\alpha, \beta_1)s'$ and $s(\alpha, \beta_2)s'$ are indistinguishable for agent 1 but distinguishable for agent 2. Moreover, CapATL includes a knowledge operator to express whether agents know certain facts about capacities compatible with the history. In contrast, ATL-SA assumes that actions are publicly observable and does not include a knowledge operator. This design choice emphasizes the probabilistic aspect on capacities with as little changes to ATL as possible. Second, in CapATL, agents can choose their capacities. Specifically, a coalition wins if there is a capacity assignment for them that succeeds against all capacity assignments of their opponents. In ATL-SA, however, agents do not choose their capacities; instead, capacities are assigned probabilistically according to predefined distributions. This enables ATL-SA to express more nuanced objectives, such as winning with a certain probability. It is important to note that CapATL is not expressible within ATL-SA (ignoring the indistinguishability aspect). One might attempt to use a probability threshold of 1, but this would imply that all capacities of the coalition (instead of *some*, in CapATL) must win against all capacities of the opponents.

Probabilistic concepts in strategic logics. The literature on stochastic MAS and probabilistic model checking introduces two key probabilistic concepts: *stochastic CGS* and *stochastic strategies*. We compare our notion of stochastic abilities with these two.

Stochastic CGS were introduced by Chen and Lu in 2007 [13], referred to as probabilistic CGS in their paper. At a low level, these extend CGS by incorporating a stochastic transition function into the game structure, where a state and a joint action by all agents correspond to a probability distribution over the next states. At a higher level, given the strategies of all agents, the stochastic CGS becomes a Markov chain, providing a probability distribution over the outcomes. Semantically, the strategic coalition first chooses its strategy, followed by the opponent, and finally random events occur at runtime. Stochastic CGS have been explored to define probabilistic versions of ATL [13], imperfect information ATL and ATL* [18, 25] with memoryless agents [6, 7] or natural strategies [10], *Strategy Logic* (SL) [4], and resource-bounded ATL [23]. However, stochastic CGS differ from stochastic abilities in ATL-SA because, in ATL-SA, an agent's capacities imply a probabilistic commitment to the rest of the outcome, whereas in stochastic CGS, random events occurring at one moment are independent of those occurring later.

A second probabilistic concept in strategic logics is *stochastic strategies*. While a deterministic strategy maps each history (or state, depending on the definition) to the action that the agent performs, a stochastic strategy maps each history (or state) to a distribution over actions. At runtime, the agent selects an action according to that distribution. Opponents are similarly constrained to a stochastic strategy, which transforms a CGS or stochastic CGS into a Markov chain. Conceptually, stochastic strategies expand the strategy space available to agents. To the best of our knowledge, stochastic strategies do not increase an agent's power in deterministic game structures (except for purposes of deliberate failure). Nonetheless, they become relevant when applied to probabilistic CGS. Stochastic strategies have been used in [4, 7, 23] and in combination with natural strategies in [10]. Once again, stochastic strategies do not imply the commitment found in ATL-SA: the probability that a stochastic strategy selects an action at a given history cannot be explicitly restricted by past events in the game.

The main difference we emphasize, is that ATL-SA's stochastic aspect arises only when agents choose their strategies, as agents are probabilistically assigned capacities. This probabilistic event implies a commitment for the remainder of the play, which is absent in stochastic CGS and stochastic strategies. As a result, ATL-SA's probabilistic nature is distinct and could be explored in conjunction with these other concepts in future research. Last but not least, the outcome of the stochastic capacity assignment is private to each agent (or shared within a coalition), aligning our work more closely with imperfect information PATL [18, 25]. However, we adopt perfect recall semantics, which is known to be undecidable in general for PATL under perfect recall conditions.

7 CONCLUSION

This article introduces a novel probabilistic dimension to ATL, inspired by the concept of capacities in CapATL. In our framework, ATL-SA, each agent is associated with a probability distribution over subsets of actions (their capacities), representing different agent profiles. Agents are restricted to using only actions from their assigned capacity in their strategies, enabling the modeling of MAS with uncertainty regarding agent profiles and the expression of properties related to the probability that a coalition can achieve a temporal objective. We examine scenarios where coalitions may or may not communicate their capacities, and establish the NEXPTIME-completeness for single strategy formulae and both semantics, along with PTIME-completeness under certain capacity restrictions. The general model-checking problem is between NEXPTIME and $\mathsf{P}^{\mathsf{NEXPTIME}}$. Finally, we illustrate the applicability of ATL-SA in a cybersecurity use case.

In the future, we plan to combine stochastic CGS with stochastic abilities. We will also analyse the impact of stochastic strategies in this setting. Finally, we aim to incorporate imperfect information on strategies and CapATL's knowledge operators regarding capacities.

ACKNOWLEDGMENTS

This work was carried out within SEIDO Lab, a joint research laboratory covering research topics in the field of smart grids, *e.g.*, distributed intelligence, service collaboration, cybersecurity, and privacy. It involves researchers from academia (Télécom Paris, Télécom SudParis, CNRS LAAS) and industry (EDF R&D).

REFERENCES

[1] Thomas Ågotnes. 2006. Action and Knowledge in Alternating-Time Temporal Logic. *Synth.* 149, 2 (2006), 375–407. https://doi.org/10.1007/S11229-005-3875-8

[2] Natasha Alechina, Brian Logan, Nguyen Hoang Nga, and Abdur Rakib. 2010. Resource-Bounded Alternating-Time Temporal Logic. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1* (Toronto, Canada) *(AAMAS '10)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 481–488.

[3] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-Time Temporal Logic. *Journal of the ACM* 49, 5 (Sept. 2002), 672–713. https://doi.org/10.1145/585265.585270

[4] Benjamin Aminof, Marta Kwiatkowska, Bastien Maubert, Aniello Murano, and Sasha Rubin. 2019. Probabilistic Strategy Logic. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 32–38. https://doi.org/10.24963/IJCAI.2019/5

[5] Gabriel Ballot, Vadim Malvone, Jean Leneutre, and Youssef Laarouchi. 2024. Strategic Reasoning under Capacity-Constrained Agents. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024, Auckland, New Zealand, May 6-10, 2024*, Mehdi Dastani, Jaime Simão Sichman, Natasha Alechina, and Virginia Dignum (Eds.). ACM, 123–131. https://doi.org/10.5555/3635637.3662859

[6] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, and Aniello Murano. 2023. Strategic Abilities of Forgetful Agents in Stochastic Environments. In *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023, Rhodes, Greece, September 2-8, 2023*, Pierre Marquis, Tran Cao Son, and Gabriele Kern-Isberner (Eds.). 726–731. https://doi.org/10.24963/KR.2023/71

[7] Francesco Belardinelli, Wojtek Jamroga, Munyque Mittelmann, and Aniello Murano. 2024. Verification of Stochastic Multi-Agent Systems with Forgetful Strategies. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2024, Auckland, New Zealand, May 6-10, 2024*, Mehdi Dastani, Jaime Simão Sichman, Natasha Alechina, and Virginia Dignum (Eds.). International Foundation for Autonomous Agents and Multiagent Systems / ACM, 160–169. https://doi.org/10.5555/3635637.3662863

[8] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. 2017. Verification of Multi-Agent Systems with Imperfect Information and Public Actions. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, Kate Larson, Michael Winikoff, Sanmay Das, and Edmund H. Durfee (Eds.). ACM, 1268–1276. http://dl.acm.org/citation.cfm?id=3091301

[9] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. 2020. Verification of Multi-Agent Systems with Public Actions against Strategy Logic. *Artif. Intell.* 285 (2020), 103302. https://doi.org/10.1016/J.ARTINT.2020.103302

[10] Raphaël Berthon, Joost-Pieter Katoen, Munyque Mittelmann, and Aniello Murano. 2024. Natural Strategic Ability in Stochastic Multi-Agent Systems. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (Eds.). AAAI Press, 17308–17316. https://doi.org/10.1609/AAAI.V38I16.29678

[11] P. E. Boas. 1996. The Convenience of Tiling. (1996). https://www.semanticscholar.org/paper/31089fbcd06121ce8f19ef70f27c702db72dcac9

[12] Thomas Brihaye, Véronique Bruyere, and Jean-François Raskin. 2005. On Optimal Timed Strategies. In *International Conference on Formal Modeling and Analysis of Timed Systems (SpringerLink)*, Paul Pettersson and Wang Yi (Eds.). Springer, Springer Berlin Heidelberg, Berlin, Heidelberg, 49–64.

[13] Taolue Chen and Jian Lu. 2007. Probabilistic Alternating-Time Temporal Logic and Model Checking Algorithm, In Fourth International Conference on Fuzzy Systems and Knowledge Discovery (FSKD 2007). *Fourth International Conference on Fuzzy Systems and Knowledge Discovery (FSKD 2007)* 2, 35–39. https://doi.org/10.1109/FSKD.2007.458

[14] Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, Roderick Bloem, et al. 2018. *Handbook of Model Checking*. Springer eBook Collection, Vol. 10. Springer Cham, Cham. https://doi.org/10.1007/978-3-319-10575-8

[15] Alexandre David, Peter G. Jensen, Kim Guldstrand Larsen, Axel Legay, Didier Lime, Mathias Grund Sørensen, and Jakob H. Taankvist. 2014. On Time with Minimal Expected Cost!. In *Automated Technology for Verification and Analysis*, Franck Cassez and Jean-François Raskin (Eds.). Springer, Cham, 129–145.

[16] Cătălin Dima, Constantin Enea, and Dimitar Guelev. 2010. Model-Checking an Alternating-Time Temporal Logic with Knowledge, Imperfect Information, Perfect Recall and Communicating Coalitions. *Electronic Proceedings in Theoretical Computer Science* 25 (June 2010), 103–117. https://doi.org/10.4204/eptcs.25.12

[17] Andreas Herzig, Emiliano Lorini, and Dirk Walther. 2013. Reasoning about Actions Meets Strategic Logics. In *Logic, Rationality, and Interaction*, Davide Grossi, Olivier Roy, and Huaxin Huang (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 162–175.

[18] Xiaowei Huang, Kaile Su, and Chenyi Zhang. 2012. Probabilistic Alternating-Time Temporal Logic of Incomplete Information and Synchronous Perfect Recall, In Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada, Jörg Hoffmann and Bart Selman (Eds.). *Proceedings of the AAAI Conference on Artificial Intelligence* 26, 1, 765–771. https://doi.org/10.1609/aaai.v26i1.8214

[19] W. Jamroga. 2003. Some Remarks on Alternating Temporal Epistemic Logic, In Formal Approaches to Multi-Agent Systems. *Formal Approaches to Multi-Agent Systems*. https://www.semanticscholar.org/paper/37e42747212142ef57bb4b36dcd12920225765b1

[20] Wojciech Jamroga and Wiebe van der Hoek. 2004. Agents That Know How to Play. *Fundamenta Informaticae* 63 (2004), 185–219. 2-3.

[21] Munyque Mittelmann, Bastien Maubert, Aniello Murano, and Laurent Perrussel. 2023. Formal Verification of Bayesian Mechanisms. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, February 7-14, 2023*, Brian Williams, Yiling Chen, and Jennifer Neville (Eds.). AAAI Press, 11621–11629. https://doi.org/10.1609/AAAI.V37I10.26373

[22] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. 2014. Reasoning about Strategies: On the Model-Checking Problem. *ACM Transactions on Computational Logic* 15, 4 (Aug. 2014), 1–47. https://doi.org/10.1145/2631917

[23] Hoang Nga Nguyen and Abdur Rakib. 2019. Probabilistic Resource-Bounded Alternating-Time Temporal Logic. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems* (Montreal QC, Canada) *(AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2141–2143.

[24] Philippe Schnoebelen. 2002. The Complexity of Temporal Logic Model Checking. In *Advances in Modal Logic 4, papers from the fourth conference on "Advances in Modal logic," held in Toulouse, France, 30 September - 2 October 2002*, Philippe Balbiani, Nobu-Yuki Suzuki, Frank Wolter, and Michael Zakharyaschev (Eds.). King's College Publications, 393–436. http://www.aiml.net/volumes/volume4/Schnoebelen.ps

[25] Henning Schnoor. 2010. Strategic Planning for Probabilistic Games with Incomplete Information. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1* (Toronto, Canada) *(AAMAS '10)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1057–1064.

[26] Pierre-Yves Schobbens. 2004. Alternating-Time Logic with Imperfect Recall. *Electronic Notes in Theoretical Computer Science* 85, 2 (2004), 82–93. https://doi.org/10.1016/S1571-0661(05)82604-0 LCMAS 2003, Logic and Communication in Multi-Agent Systems.

[27] F. Schwarzentruber. 2019. The Complexity of Tiling Problems. (2019). https://doi.org/10.48550/arxiv.1907.00102 arXiv:1907.00102

[28] Wiebe van der Hoek and Michael J. Wooldridge. 2002. Tractable Multiagent Planning for Epistemic Goals. In *The First International Joint Conference on Autonomous Agents & Multiagent Systems*. ACM, Bologna, Italy, 1167–1174. https://doi.org/10.1145/545056.545095

[29] Wiebe van der Hoek and Michael J. Wooldridge. 2003. Cooperation, Knowledge, and Time: Alternating-Time Temporal Epistemic Logic and Its Applications. *Stud Logica* 75, 1 (2003), 125–157. https://doi.org/10.1023/A:1026185103185

[30] Dirk Walther, Wiebe van der Hoek, and Michael Wooldridge. 2007. Alternating-Time Temporal Logic with Explicit Strategies. In *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge* (Brussels, Belgium) *(TARK '07)*. Association for Computing Machinery, New York, NY, USA, 269–278. https://doi.org/10.1145/1324249.1324285

[31] Thomas Ågotnes, Valentin Goranko, and Wojciech Jamroga. 2007. Alternating-Time Temporal Logics with Irrevocable Strategies. In *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge* (Brussels, Belgium) *(TARK '07)*. Association for Computing Machinery, New York, NY, USA, 15–24. https://doi.org/10.1145/1324249.1324256