

Verification of Stochastic Multi-Agent Systems with Forgetful Strategies

Francesco Belardinelli

Imperial College

London, United Kingdom

francesco.belardinelli@imperial.ac.uk

MunIQUE MittelmANN

University of Naples Federico II

Naples, Italy

munIQUE.mittelmANN@unina.it

Wojciech Jamroga

SnT, University of Luxembourg, Luxembourg

ICS, Polish Academy of Sciences, Poland

wojciech.jamroga@uni.lu

Aniello Murano

University of Naples Federico II

Naples, Italy

aniello.murano@unina.it

ABSTRACT

Intelligent autonomous agents need to reason about different kinds of uncertainty in a Multi-Agent System (MAS): first, due to the occurrence of randomization and, second, their inability to completely observe the state of the system. In this paper, we investigate the verification of system specifications in probabilistic variants of the logics ATL and ATL* under imperfect information (II). The resulting setting combines these two sources of uncertainty and captures the situation in which agents have qualitative uncertainty about the local state as well as quantitative uncertainty about the occurrence of future events. Since the model-checking problem is undecidable when considered in the context of strategies with perfect recall, we focus on memoryless (positional) strategies. As the main result, we show that, in stochastic MAS under II, model-checking Probabilistic ATL is in EXPTIME when agents play probabilistic strategies. Filling the gap in recent work, we also show that model-checking Probabilistic ATL* is PSPACE-complete when the proponent coalition is restricted to deterministic strategies.

KEYWORDS

Stochastic Multi-Agent Systems, Probabilistic Model Checking, Logics for Strategic Reasoning

ACM Reference Format:

Francesco Belardinelli, Wojciech Jamroga, MunIQUE MittelmANN, and Aniello Murano. 2024. Verification of Stochastic Multi-Agent Systems with Forgetful Strategies. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024)*, Auckland, New Zealand, May 6 – 10, 2024, IFAAMAS, 10 pages.

1 INTRODUCTION

Formal methods for strategic reasoning play a fundamental role in Multi-Agent System (MAS) design and verification [3, 16, 56, 67, 71, 77, 79]. This success story originated from the breakthrough idea of using temporal logics for the specification of behaviors of reactive systems [31, 36, 70]. Temporal logics are traditionally

interpreted over Kripke structures, modeling closed systems, and quantifying the computations of the systems universally and existentially. The need to reason about MAS led to the development of formalisms that enable the specification of strategic behaviors of agents [3, 60, 66, 67]. One of the main developments along this line has been Alternating-time Temporal Logic (ATL) [3], which is a logical formalism for the specification and verification of open systems involving multiple autonomous agents and allows expressing statements about what coalitions of agents can achieve by strategic cooperation.

The autonomous agents that compose a MAS often need to reason about different kinds of uncertainty. One of the sources of uncertainty is their inability to completely observe the current local state (e.g., employees in a company have access to different client information). On the other hand, MAS also face the occurrence of randomization, for instance, due to natural events or the behavior of other agents. While this aspect cannot be known with certainty, it can be measured based on experiments or past observations. For instance, while we cannot know whether a web system will be available when it needs to be used, past observations enable us to measure the probability of its availability. Clearly, both the imperfect information about the local state and the likelihood of stochastic events need to be taken into account by strategic agents.

Probabilistic model-checking is a technique for the formal verification of probabilistic systems that can be modeled by stochastic state-transition models [32]. It can be used to establish the correctness of such systems against probabilistic specifications, which may describe, e.g., the probability of a failure of a system, or the ability of a coalition to protect it from attackers. Alongside model-checking techniques, logic-based formalisms have been widely and successfully applied for the verification of stochastic MAS, including economic mechanisms [65], negotiation games [8], team formation protocols [28], and dispersion games [43], to name a few.

In this paper, we are interested in the model-checking problem in stochastic MAS with partial observability. In particular, we consider the Probabilistic Alternating-time Temporal Logics PATL and PATL* [29, 43] under imperfect information (II). Since model checking PATL* under II for memoryful agents (a.k.a. agents with perfect-recall) is known to be undecidable even for the fragment with a single-player [43], we focus on a classic type of agents [37] called imperfect-recall, i.e., agents who use memoryless strategies, also called Markovian strategies or policies.



This work is licensed under a Creative Commons Attribution International 4.0 License.

	Det. Strat.	Prob. Strat.
PATL _{or}	Δ_2^P -c. [9]	in EXPTIME (new)
PATL _{sr}		
PATL _{or} *	PSPACE-c (new)	?
PATL _{sr} *		

Table 1: Summary of model checking complexity results for PATL and PATL*. The subscripts “o” and “s” denote objective and subjective interpretations, while “r” stands for memoryless strategies. Open problems are indicated with “?”.

Contribution. We study the model-checking problem of the probabilistic logics PATL and PATL* under imperfect information and memoryless strategies. The main advantage of this setting is that it captures MAS with two sources of uncertainty, namely randomization and partial observation, and enables reasoning about the strategic abilities of memoryless agents. We consider two semantic variations for the logics, the objective and subjective interpretations, as well as the cases whereby agents are allowed to play probabilistic v. deterministic strategies, and show how to introduce epistemic operators alongside them. Table 1 summarizes the results presented in this paper and the remaining open questions. Our main result is the solution of the model-checking problem for PATL when agents play probabilistic strategies, which we show can be done in EXPTIME. Filling the gap in recent work [9], we also show that model-checking PATL* is PSPACE-complete when the proponent coalition is restricted to deterministic strategies.

Related Work. One of the most important pioneering work on logics for MAS is the Alternating-time Temporal Logic ATL* and its fragment ATL [3]. ATL* has been extended in various directions, considering for instance strategy contexts [59], or adding imperfect information and epistemic operators [47]. Strategy Logic [27, 67] extends ATL* to represent strategies with first-order variables.

Imperfect information have been extensively considered in the literature on formal verification on MAS (see, for instance, [10, 13, 14, 22, 33, 46, 54, 55, 73]). While imperfect information is a key feature of MAS, where perfect observability is either unrealistic or computationally unattainable, it entails higher complexity, even undecidability when considered in the context of agents with perfect recall [33]. The case of agents with memoryless strategies is interesting to retrieve a decidable model-checking problem [25].

The verification of systems against specifications given in probabilistic logics has been widely studied. In particular, Wan et al. study the model-checking problem for Probabilistic Epistemic Computational Tree Logic with semantics based on probabilistic interpreted systems [81]. In the context of MAS, Kwiatkowska et al. detail how verification techniques for concurrent stochastic games can be developed and implemented using the PRISM model checker [57, 58]. Huang and Luo study an ATL-like logic for stochastic MAS in which agents play deterministic strategies and have probabilistic knowledge about the system [44]. Fu et al. show that the model-checking problem for an epistemic logic with temporal operators is undecidable when considering strategies that depend on agents’ observation history [40]. Chen and Lu propose model-checking algorithms for Probabilistic ATL in MAS with perfect information [29]. This

setting was also considered alongside Probabilistic Alternating-Time μ -Calculus [78] and Probabilistic Strategy Logic (PSL) [4]. All these results cannot be adapted to solve the model-checking problem we are interested in here because they are restricted to the setting of perfect information. Additionally, it is known that Probabilistic Alternating μ -calculus and PATL are incomparable [21, 78]. In the case of PSL, the model checking problem is already 3-EXPTIME-complete, while we show it is in EXPTIME for PATL under the assumption of imperfect information.

ATL-based probabilistic logics were also considered for the verification of unbounded parameterized MAS [61], for resource-bounded MAS [69], alongside behavioral natural strategies [12], and under assumptions over opponents’ strategies [20].

One of the closest related works is [45], which considers the logic PATL* under incomplete information and synchronous perfect recall. The complexity results show that the model-checking problem is in general undecidable even for the single-agent fragment of the logic. PATL with imperfect information was recently considered with the restriction of deterministic memoryless strategies for the proponent coalition [9]. In the present work, we also consider the general case in which agents in the proponent coalition may play probabilistic strategies. The work in [45] considers only the subjective semantics, while we consider both subjective and objective semantics of ability, and extend the logic with epistemic operators. In particular, the technique used in [9] is based on calling an oracle that guesses the successful memoryless strategy. This method cannot be applied when considering probabilistic strategies for the proponent coalition because there are *infinitely many* such strategies, and hence the oracle Turing machine would either have to run in unbounded time, or allow for infinite branching.

Another related problem is the verification of probabilistic observability properties studied in [68]. The main difference with this work is that they consider perfect recall for strategies and observations, whereas we focus on memoryless strategies and memoryless knowledge. Also related is the research of algorithmic solutions for computing winning strategies for two-player stochastic games with imperfect information [24, 34, 35, 41]. Chatterjee and Doyen study the problem of deciding the existence of almost-sure and positive winning strategies in such games with partial-observation [26]. Finally, Gurov et al. investigate strategy synthesis for knowledge-based strategies against a non-deterministic environment [42].

2 PRELIMINARIES

We start by recalling the basic definitions of stochastic multi-agent models and strategic play [17, 29, 45]. In our presentation, we follow mainly [9]. Fix finite non-empty sets Ag of agents a, a', \dots ; Ac of actions α, α', \dots ; and AP of atomic propositions p, p', \dots . We write \mathbf{o} for a tuple $(o_a)_{a \in \text{Ag}}$ of objects, one for each agent; such tuples are called *profiles*. A *joint action* or *move* c is an element of Ac^{Ag} . Given a profile \mathbf{o} and $C \subseteq \text{Ag}$, we let \mathbf{o}_C be the components for the agents in C . Moreover, we use Ag_{-C} as a shorthand for $\text{Ag} \setminus C$.

Distributions. Let X be a finite non-empty set. A (*probability*) *distribution* over X is a function $d : X \rightarrow [0, 1]$ such that $\sum_{x \in X} d(x) = 1$. $\text{Dist}(X)$ is the set of distributions over X . We write $x \in d$ for $d(x) > 0$. If $d(x) = 1$ for some element $x \in X$, then d is a *point*

(*a.k.a. Dirac*) distribution. If d_i is a distribution over X_i , then, writing $X = \prod_i X_i$, the *product distribution* of the d_i is the distribution $d : X \rightarrow [0, 1]$ defined by $d(x) = \prod_i d_i(x_i)$.

Markov Chains. A *Markov chain* M is a tuple (St, d) where St is a set of states and $d \in \text{Dist}(St \times St)$ is a distribution. The values $d(s, t)$ are called *transition probabilities* of M .

Concurrent Game Structures. A *stochastic concurrent game structure with imperfect information* (or simply *iCGS*) \mathcal{G} is a tuple $(St, L, \delta, \ell, \{\sim_a\}_{a \in \text{Ag}})$ where (i) St is a finite, non-empty set of states; (ii) $L : St \times \text{Ag} \rightarrow 2^{\text{Ac}} \setminus \{\emptyset\}$ is a *legality function* defining the available actions for each agent in each state; we write $L(q)$ for the set of tuples $(L(q, a))_{a \in \text{Ag}}$; (iii) for each state $q \in St$ and each move $c \in L(q)$, the *stochastic transition function* δ gives the (conditional) probability $\delta(q, c)$ of a transition from state q for all $q' \in St$ if each player $a \in \text{Ag}$ plays the action c_a ; we also write this probability as $\delta(q, c)(q')$ to emphasize that $\delta(q, c)$ is a probability distribution on St ; (iv) $\ell : St \rightarrow 2^{\text{AP}}$ is a *labelling function*; (v) $\sim_a \subseteq St \times St$ is an equivalence relation called the *observation relation* of agent a .

A pointed iCGS is a pair (\mathcal{G}, q) where $q \in St$ is a special state designed as initial. Throughout this paper, we assume that iCGSs are *uniform*, that is, if two states are indistinguishable for an agent a , then a has the same available actions in both states. Formally, if $q \sim_a q'$ then $L(q, a) = L(q', a)$, for any $q, q' \in St$ and $a \in \text{Ag}$. For each state $q \in St$ and joint action $c \in L(q)$, we also assume that there is a state $q' \in St$ such that $\delta(q, c)(q')$ is non-zero, that is, every state has a successor state from a legal move. Finally, we say that \mathcal{G} is *deterministic* (instead of stochastic) if every $\delta(q, c)$ is a point distribution.

Plays. A *play* (or path) in a iCGS \mathcal{G} is an infinite sequence $\pi = q_0 q_1 \dots$ of states such that there exists a sequence $c_0 c_1 \dots$ of joint-actions such that for every $i \geq 0$, $c_i \in L(q_i)$ and $q_{i+1} \in \delta(q_i, c_i)$ (i.e., $\delta(q_i, c_i)(q_{i+1}) > 0$). We write π_i for state q_i , $\pi_{\geq i}$ for the suffix of π starting at position i . Finite paths are called *histories*, and the set of all histories is denoted Hist . Write $\text{last}(h)$ for the last state of a history h .

Strategies. A (general) *probabilistic strategy* is a function $\sigma : \text{Hist} \rightarrow \text{Dist}(\text{Ac})$ that maps each history to a distribution of actions. We let Str be the set of all strategies. A *memoryless uniform probabilistic strategy* for an agent a is a function $\sigma_a : St \rightarrow \text{Dist}(\text{Ac})$ in which for all positions q, q' such that $q \sim_a q'$, we have $\sigma_a(q) = \sigma_a(q')$. We let Str_a^r be the set of memoryless uniform strategies for agent a . A deterministic (or *pure*) strategy σ_a for agent a is a strategy in which $\sigma_a(q)$ is a point distribution for any q . A *strategy profile* is a tuple σ of strategies, one for each agent. We write σ_a for the strategy of a in profile σ . For a strategy σ_a for agent a , we assume that $\sigma_a(h)(c) = 0$ if $c \notin L(\text{last}(h), a)$.

3 PROBABILISTIC ATL AND ATL*

Now we present the syntax and semantics of the Probabilistic Alternating-time Temporal Logics PATL* and PATL [9, 29, 45], interpreted under the assumption of imperfect information. Again, we follow [9] in our presentation. Note that [9] adopts the *objective* semantics of strategic ability, where the coalition is supposed to have a strategy that works from the initial state of the game. In contrast, [45] uses the *subjective* semantics of strategic ability, where

the agents need a strategy that wins from all the observationally equivalent states.¹ In this paper, we consider both accounts, as they are equally relevant in the literature. In particular, we integrate the *objective* and *subjective* semantics of probabilistic ability into a single framework.

Definition 1 (PATL*). State formulas Φ and path formulas ψ are defined by the following grammar, where $p \in \text{AP}$, $C \subseteq \text{Ag}$, d is a rational constant in $[0, 1]$, and $\triangleright \in \{\leq, <, >, \geq\}$:

$$\begin{aligned} \Phi &::= p \mid \neg\Phi \mid \Phi \vee \Phi \mid \langle\langle C \rangle\rangle^{\triangleright d} \psi \\ \psi &::= \Phi \mid \neg\psi \mid \psi \vee \psi \mid \mathbf{X}\psi \mid \psi \mathbf{U}\psi \mid \psi \mathbf{R}\psi \end{aligned}$$

Formulas in PATL* are all and only the state formulas Φ .

The intuitive reading of the operators is as follows: $\langle\langle C \rangle\rangle^{\triangleright d} \psi$ means that there exists a strategy for the coalition C of agents to collaboratively enforce ψ with a probability in relation \triangleright with constant d ; “next” \mathbf{X} , “release” \mathbf{R} , and “until” \mathbf{U} are the standard temporal operators. We define the usual derived temporal operators as follows: $\mathbf{F}\psi := \top \mathbf{U}\psi$ and $\mathbf{G}\psi := \perp \mathbf{R}\psi$. Finally, we use $[[C]]^{\triangleright d} \psi := \neg \langle\langle C \rangle\rangle^{\triangleright d} \neg\psi$ to express that no strategy of C can prevent ψ with a probability in relation \triangleright with constant d .

An important syntactic restriction of PATL*, namely PATL, is obtained by restricting path formulas as follows:

$$\psi ::= \mathbf{X}\Phi \mid \Phi \mathbf{U}\Phi \mid \Phi \mathbf{R}\Phi$$

which is tantamount to the following grammar for state formulas:

$$\Phi ::= p \mid \neg\Phi \mid \Phi \vee \Phi \mid \langle\langle C \rangle\rangle^{\triangleright d} \mathbf{X}\Phi \mid \langle\langle C \rangle\rangle^{\triangleright d} (\Phi \mathbf{U}\Phi) \mid \langle\langle C \rangle\rangle^{\triangleright d} (\Phi \mathbf{R}\Phi)$$

where again $p \in \text{AP}$, $C \subseteq \text{Ag}$, and $\triangleright \in \{\leq, <, >, \geq\}$.

Formulas of PATL and PATL* are interpreted over iCGSs.

Probability Space on Outcomes. An *outcome* of a strategy profile σ and a state q is a play π that starts with q and is extended by letting each agent follow the strategies in σ , i.e., $\pi_0 = q$, and for every $k \geq 0$ there exists $c_k \in \sigma(\pi_k)$ such that $\pi_{k+1} \in \delta(\pi_k, c_k)$. The set of outcomes of a strategy profile σ and state q is denoted as $\text{out}(\sigma, q)$. A given iCGS \mathcal{G} , strategy profile σ , and state q induce an infinite-state Markov chain $M_{\sigma, q}$ whose states are the finite prefixes of plays in $\text{Out}(\sigma, q)$. Such finite prefixes of plays are actually *histories*. Transition probabilities in $M_{\sigma, q}$ are defined as $p(h, hq') = \sum_{c \in \text{Ac}^{\text{Ag}}} \sigma(h)(c) \cdot \delta(\text{last}(h), c)(q')$. The Markov chain $M_{\sigma, q}$ induces a canonical probability space on its set of infinite paths [51], which can be identified with the set of plays in $\text{Out}(\sigma, q)$.²

Given a coalitional strategy $\sigma_C \in \prod_{a \in C} \text{Str}_a^r$, we define its *objective outcomes* from state $q \in St$ as the set $\text{out}_{o, C}(\sigma_C, q) = \{\text{Out}((\sigma_C, \sigma_{\text{Ag}-C}), q) \mid \sigma_{\text{Ag}-C} \in \text{Str}^{|\text{Ag}-C|}\}$ of probability measures that the players in C enforce when they follow the strategy σ_C . Note that the opponents can use any strategies within $\sigma_{\text{Ag}-C}$, even if C employ only uniform memoryless strategies.

The *subjective outcomes* are then defined as the set

$$\text{out}_{s, C}(\sigma_C, q) = \bigcup_{q' \sim_a q, a \in C} \text{out}_{o, C}(\sigma_C, q') \quad (1)$$

We will use $\mu_{x, C}^{\sigma_C}$ to range over the elements of $\text{out}_{x, C}(\sigma_C, q)$, for $x \in \{s, o\}$.

¹For a more thorough discussion of objective vs. subjective ability, cf. [1, 22].

²This is a classic construction, see for instance [11, 30].

Remark 1. We note in passing that [45] base their semantics of subjective ability for coalitions upon distributed knowledge (i.e., the intersection of the members' outcome sets), whereas in (1) we use mutual knowledge (i.e., the union of the outcome sets), which is more standard in reasoning about subjective ability [22, 76].

Semantics. For x equal to either s or o , state and path formulas in PATL* are interpreted in a iCGS \mathcal{G} and a state q , resp. path π , according to the x -interpretation of strategy operators, as follows (clauses for Boolean connectives are omitted as immediate):

$$\begin{aligned}
\mathcal{G}, q \models_x p & \quad \text{iff } p \in \ell(q) \\
\mathcal{G}, q \models_x \langle\langle C \rangle\rangle^{\triangleright d} \psi & \quad \text{iff } \exists \sigma_C \in \prod_{a \in C} \text{Str}'_a \text{ such that} \\
& \quad \forall \mu_{x,q}^{\sigma_C} \in \text{out}_{x,C}(\sigma_C, q), \\
& \quad \mu_{x,s}^{\sigma_C}(\{\pi \mid \mathcal{G}, \pi \models_x \psi\}) \triangleright d \\
\mathcal{G}, \pi \models_x X\psi & \quad \text{iff } \mathcal{G}, \pi_{\geq 1} \models_x \psi \\
\mathcal{G}, \pi \models_x \psi_1 U \psi_2 & \quad \text{iff } \exists k \geq 0 \text{ s.t. } \mathcal{G}, \pi_{\geq k} \models_x \psi_2 \text{ and} \\
& \quad \forall j \in [0, k) \mathcal{G}, \pi_{\geq j} \models_x \psi_1 \\
\mathcal{G}, \pi \models_x \psi_1 R \psi_2 & \quad \text{iff } \forall k \geq 0, \mathcal{G}, \pi_{\geq k} \models_x \psi_2 \text{ or} \\
& \quad \exists j \in [0, k) \text{ s.t. } \mathcal{G}, \pi_{\geq j} \models_x \psi_1
\end{aligned}$$

Remark 2. Notice that, by using the subjective interpretation of PATL*, we can introduce the individual knowledge operator K_a of epistemic logic as follows: $K_a \Phi ::= \langle\langle a \rangle\rangle^{>0} \perp U \Phi$.

By definition of the satisfaction relation \models_s , we have that

$$\mathcal{G}, q \models_s K_a \Phi \quad \text{iff for all } q' \sim_a q, \mathcal{G}, q' \models_s \Phi \quad (2)$$

On the other hand, for the objective semantics, K_a can be added as a primitive operator with the semantics defined as in Eq. (2).

The Model Checking Problem. The setting introduced in this paper includes two logics: PATL* and PATL, which are interpreted over stochastic iCGS by using either probabilistic or deterministic (memoryless) strategies for the proponent coalition, according to two different semantics: objective or subjective. This gives a total of 3 different dimensions. We use the notation $\text{PATL}_{\text{or}}^*$, $\text{PATL}_{\text{sr}}^*$, PATL_{or} , and PATL_{sr} to refer to the objective and subjective variants of PATL* and PATL, respectively. As a result, we obtain 8 variants of the model checking problem, defined as follows (see Table 1 for an overview).

Definition 2 (Model Checking Problem). Given a stochastic iCGS \mathcal{G} , a formula $\Phi \in L$, for $L \in \{\text{PATL}_{\text{sr}}^*, \text{PATL}_{\text{or}}^*\}$ and $x \in \{o, s\}$, and a state q , the model checking problem is to determine whether $\mathcal{G}, q \models_x \Phi$, when considering either probabilistic or deterministic strategies for the proponent coalition.

The rest of this paper is devoted to analyzing the decidability and complexity of model checking of these problems. We anticipate that some of the dimensions listed above do not have an impact. For instance, complexity results are the same for objective and subjective interpretation.

4 STRATEGIC REASONING UNDER UNCERTAINTY

In this section, we discuss motivating problems of strategic reasoning in stochastic MAS with agents that have partial observability

of the environment. Our examples are based on security games and probabilistic social choice theory.

4.1 Security Games

Security games are game-theoretic models used to study security problems, such as the protection of biodiversity in conservation areas [38] and wildlife protection from cooperative attackers [82]. In the basic setting, a security game [52] is a two-player game between a defender and an attacker. The attacker may choose to attack any target, while the defender tries to prevent attacks by covering targets using resources. This formalism has been extended to multiple attackers [82], and multiple defenders [62]. Many real-world scenarios are not single-shot games, as the attackers often conduct multiple repeated attacks. This is the case, among others, of security games for protecting the environment [50] (e.g., defending from hunters who continuously try to poach various animals).

Let us consider a multi-defender security game inspired by [62]. The set of agents is $\text{Ag} = \{a\} \cup D$, where D is a non-empty set of defenders and a is the attacker. Each defender $i \in D$ is in charge of protecting a set of targets T_i . The set of all targets is $T = \bigcup_{i \in D} T_i$.

An action for a defender i consist in a subset of targets $o_i \subseteq T_i$, where $t \in o_i$ means that i is covering target t . The action $o_i = \emptyset$ represents that i does not cover any target. However, covering all targets may not be feasible due to resource constraints [53]. We assume each defender i has some given number of resources k that is at most $|T_i|$, used to cover targets. The amount of resources available can change in each time step. The attacker's actions consist of attacks to one of the targets in $t \in T$.

For each $t \in T$, the atomic propositions attacked_t and covered_t , denote whether the target t is attacked or covered resp. The proposition $\text{resource}_{i,k}$ indicates that the defender i has k resources to employ (that is, the maximum capacity to cover targets), for $i \in D$ and $0 \leq k \leq |T_i|$. The resource constraints are represented by the legality function L : given a state q and a defender i , an action o_i is legal for i in q (i.e., $o_i \in L(q, i)$), if $|o_i|$ is smaller or equal to greatest k such that $\text{resource}_{i,k} \in \ell(q)$ (if no such k exists, $L(q, i) = \{\emptyset\}$).

Attacks in a covered target always fail (that is, the proposition attacked_t is false in the next state). On the other hand, an attack in an uncovered target t may fail in some cases (with a given probability). This is captured by the stochastic transition function.

Assume an instance of the problem with two defenders, namely Ann and Bob, who are in charge of defending a forest from the attacker Carol. Ann is in charge of defending north- and south-east zones of the forest (targets NE and SE, resp.), while Bob should defend the north- and south- west zones (targets NW and SW, resp.).

Let q be a state in which Carol can attack any target while Ann and Bob have only one resource each (that is, each one can cover at most one zone). Table 2 illustrates the possible combinations of actions from state q , where X denotes the situations in which Carol would attack an uncovered target.

If the attackers know when they attacked an unprotected target, deterministic memoryless strategies for the defenders are not enough to protect their targets, because, when a situation repeats, the attacker could simply attack the targets left unprotected previously. On the other hand, probabilistic strategies add uncertainty about the behavior of the defenders.

The PATL* formula

$$\langle\langle a \rangle\rangle^{\geq \frac{1}{2}} \text{resources}_{a,k} \rightarrow \mathbf{X} \bigvee_{k \geq k' \geq |T_a|} \text{resources}_{a,k'}$$

says that a has at least $\frac{1}{2}$ probability of ensuring that, if she has k resources at a state, this amount will not decrease in the next state.

For each $t \in T$, let the PATL formula $\text{destroyed}_t := \text{attacked}_t \wedge \neg \text{protected}_t$ denote that target t was destroyed if it was attacked while unprotected.

The PATL formula

$$\langle\langle a \rangle\rangle^{\geq c} \mathbf{G} \bigwedge_{t \in T_a} (K_a \text{destroyed}_t \vee K_a \neg \text{destroyed}_t)$$

represents that agent a can ensure with probability c that for each of her targets, she knows whether it was destroyed or not.

The PATL* formula $\langle\langle C \rangle\rangle^{\geq c} \wedge_{a \in C} \bigvee_{t \in T_a} \mathbf{G} \neg \text{destroyed}_t$ represents that the coalition C can ensure, with probability greater or equal to c that at least one target of each member of the coalition will never be destroyed. Assuming each agent in C have always at least one resource (which would allow to cover a target), the formula would be true for $0 < c \leq 1$, since each agent could keep protecting the same target.

The PATL formula

$$\langle\langle a \rangle\rangle^{\geq \frac{1}{4}} \mathbf{G} \bigwedge_{b \in D} \bigwedge_{0 \leq k \leq |T_b|} \text{resources}_{b,k} \rightarrow K_a \text{resources}_{b,k}$$

represents that agent a has at least $\frac{1}{4}$ probability to always ensure that, for each defender b and her possible amount of resources k , if it is the case that b has k resources, then a knows it.

4.2 Probabilistic Social Choice

In recent years, randomization has played an increasingly relevant role in social choice theory and mechanism design [7, 18]. One reason is that deterministically picking a winner is often unfair (e.g., when two agents have the same preference or score). Furthermore, probabilistic approaches enable circumventing impossibility results, such as achieving strategyproofness and non-dictatorship [6].

While in classic voting models, all voters submit their vote at once, in many realistic scenarios committees often follow an informal voting process where members are free to revise their votes. In iterative voting mechanisms, the game proceeds in turns, where single or multiple voters change their vote at each turn until no voter has objections and the final outcome is announced [64].

Voters' actions include reporting an alternative from a finite set of alternatives T and voting in "none". The atomic propositions $\text{vote}_{a,t}$, and $\text{pref}_{a,t}$ denote whether the agent a last vote was to the alternative t and whether t is her most preferred alternative. Finally, the proposition choice_t specifies whether t is the alternative

		Carol			
		NE	SE	NW	SW
(Ann, Bob)	(NE, NW)		X		X
	(NE, SW)		X	X	
	(SE, NW)	X			X
	(SE, SW)	X		X	

Table 2: Example of action profile for an instance of a security game. X denotes that an uncovered target was attacked.

chosen. After the agents vote, choice_t is true for the alternative that received more votes. Ties are broken according to a probability distribution over the most preferred alternatives (e.g., for a tie among n alternatives, each one is chosen with probability $\frac{1}{n}$).

The anonymity of votes can be verified with PATL. For instance, the formula

$$\neg \langle\langle b \rangle\rangle^{\geq \frac{2}{3}} \mathbf{G} \bigwedge_{t \in T} (\neg K_b \text{voted}_{a,t} \wedge \neg K_b \neg \text{voted}_{a,t})$$

expresses that it is not the case that agent b has a strategy to ensure, with probability $\frac{2}{3}$, to know whether a voted in any of the alternatives.

In iterative voting, it is relevant to determine whether the choice will eventually be stable, i.e., to converge to an alternative. This condition can be captured with the formula $\langle\langle \text{Ag} \rangle\rangle^{\geq c} \bigvee_{t \in T} \mathbf{F} \mathbf{G} \text{choice}_t$ which expresses that, with probability c , at a certain point, some alternative is chosen at all future states of the path.

A property that is undesired in a social choice mechanism is called *dictatorship*, which happens when the preferences of a single voter (the dictator) determine the alternative that is chosen, whatever are the preferences of the other individuals [19]. In our example, dictatorship-free is captured by the formula

$$\bigwedge_{a \in \text{Ag}} \neg \left(\langle\langle a \rangle\rangle^{\geq 1} \mathbf{G} \bigwedge_{t \in T} (\text{pref}_{a,t} \rightarrow \text{choice}_t) \right)$$

Recently work has shown how to use variants of Strategy Logic for the verification of economic mechanisms for social choice, first in the deterministic setting with imperfect information [63], and later for stochastic mechanisms with PSL [65]. Since the model-checking of PSL is 3-EXPTIME-complete for memoryless strategies, it is interesting to explore the application of other formal verification techniques with lower computational costs. We have shown how to express a number of properties in iterative voting with PATL and PATL* and we will now focus on establishing the complexity of model-checking these logics under memoryless strategies.

5 MODEL CHECKING STOCHASTIC SYSTEMS WITH FORGETFUL AGENTS

As discussed in Section 4, many multi-agent systems are inherently stochastic and characterized by imperfect information. Moreover, it makes sense to ask about the abilities of agents with bounded memory, i.e., who cannot (or choose not to) remember the whole history of past observations. In that case, the memory of the agent can be encapsulated in its local state, and one can use memoryless strategies to model the agent's strategic decisions.

PATL *model checking* allows us to verify statements about the agents' ability (or inability) to enforce temporal goals within a given range of probabilities. The simpler case of deterministic memoryless strategies has been studied in [9], where it was proven Δ_2^P -complete by a straightforward extension of results for non-probabilistic MAS. Here, we concentrate on the more interesting (and much more difficult) case of agents that can randomize, i.e., employ *probabilistic memoryless strategies with imperfect information*.

In the rest of the section, we present our main technical results, establishing complexity bounds for the problem. The bounds are not tight, but reasonably close for reasoning about probabilistic

policies.³ The proofs are nontrivial, and proceed by reductions to fundamental arithmetic problems rarely used in logic-based approaches (*theory of the reals* and *existential theory of the reals*), which is an interesting contribution in itself.

5.1 Background

Conceptually, model checking of PATL_{or} and PATL_{sr} is closely related to synthesis of memoryless policies for POMDPs, which is known to be in PSPACE , as well as NP-hard and *sum-of-square-roots-hard* [80]. We start by observing that the two problems differ significantly, and cannot be easily reduced to one another.

Firstly, policy synthesis for POMDPs addresses non-nested 1.5-player games with arbitrary rewards. It looks for single-agent strategies that maximize the agent’s expected reward, averaged over all execution paths and future time points. Importantly, the reward decreases with each time step by a given temporal discount that is strictly smaller than 1. No less importantly, the proponent is playing against a purely reactive stochastic environment.

Secondly, model checking of PATL_{or} and PATL_{sr} admits nested strategic properties in games with arbitrarily many players. It seeks coalitional strategies that maximize the probability of enforcing a binary reachability/safety goal against all probabilistic behaviors of the opponents. No temporal discounting is considered.

Our proofs in the rest of this section have been inspired by the results of [80].

Importantly, it is not possible to employ the technique that was used in [9] to establish the complexity of model checking for *deterministic* memoryless strategies of the coalition, i.e., calling an oracle that guesses the best memoryless strategy, pruning the iCGS, and solving the resulting finite set of Markov chains. This is because there are *infinitely many* probabilistic memoryless strategies, and hence the oracle Turing machine would either have to run in unbounded time, or allow for infinite branching. In fact, synthesis of optimal probabilistic strategies is a special case of *jointly constrained bilinear optimization*, which is a notoriously hard problem [2]. Fortunately, our case can be reduced to deciding the second level in the *hierarchical theory of the reals* [75], which is an extension of the *existential theory of the reals* problem [23].

5.2 Probabilistic Strategies: Upper Bounds

We begin by showing that model checking PATL_{sr} for probabilistic strategies of the coalition is decidable in EXPTIME . Moreover, in the special case of formulas that include only the grand coalition of agents, the problem is in PSPACE , analogously to memoryless synthesis for POMDPs. In our proofs, we will use reductions to the following decision problems.

Definition 3 (Existential theory of the reals, $\text{Th}\mathbb{R}_{\exists}$). The problem decides the truth of a first-order formula $\Phi \equiv \exists x_1 \dots \exists x_n P(x_1, \dots, x_n)$ where x_i are interpreted over the reals \mathbb{R} , and P is a Boolean function of atomic predicates of the form $f_i(x_1, \dots, x_n) \geq 0$ or $f_i(x_1, \dots, x_n) > 0$, with each f_i being a polynomial with rational coefficients.

THEOREM 5.1 ([23]). $\text{Th}\mathbb{R}_{\exists}$ is in PSPACE .

³We would be surprised to obtain a completeness result: the exact complexity of solving POMDPs with memoryless policies is a longstanding open problem [49, 80].

Definition 4 (First-order theory of the reals, $\text{Th}\mathbb{R}$). Analogously to Definition 3, only with an arbitrary sequence of quantifiers $Q_1 \dots Q_n$ allowed at the beginning of Φ .

THEOREM 5.2 ([74]). *There is an algorithm for $\text{Th}\mathbb{R}$ that requires $(md)^{n \cdot 2^{O(\omega)}}$ operations and $(md)^{O(n)}$ calls to an oracle computing P , where m is the number of atomic predicates in Φ , d is the maximal degree of the polynomials, n is the number of quantifiers, and $\omega - 1$ the number of quantifier alternations in Φ .*

In what follows, we first show that if the opponents can prevent C from winning, they can always achieve it by a memoryless response (Lemma 1). Then, we present a construction that reduces the verification of abilities for reachability goals in iCGSs to policy optimization in multi-agent POMDPs with undiscounted rewards, and we express the latter problem as a formula in the existential theory of the reals (Proposition 1). Further, we show how the verification of safety goals can be reduced to the case of reachability goals (Proposition 2). Finally, we use the standard recursive procedure for model checking formulas with nested strategic operators (Theorem 5.3), and observe that in some special cases the reduction obtains a tighter complexity bound (Theorem 5.4).

Hereafter, let $\bar{\triangleright}$ denote the negated constraint \triangleright , i.e., $\bar{>} = \geq$, $\bar{\leq} = >$, etc.

Lemma 1. *Let (\mathcal{G}, q) be a pointed iCGS, $\langle\langle C \rangle\rangle^{\triangleright d} \varphi$ a formula of PATL_{sr} , and σ_C a memoryless uniform strategy for C . If there exists a general strategy $\sigma_{\text{Ag-C}} \in \text{Str}_{\text{Ag-C}}^{\text{IRP}}$ such that $\text{out}((\sigma_C, \sigma_{\text{Ag-C}}), q) (\{\pi \mid \mathcal{G}, \pi \models \varphi\}) \bar{\triangleright} d$, then there is a memoryless strategy $\sigma'_{\text{Ag-C}} \in \text{Str}_{\text{Ag-C}}^{\text{IRP}}$ (not necessarily uniform!) with $\text{out}((\sigma_C, \sigma'_{\text{Ag-C}}), q) (\{\pi \mid \mathcal{G}, \pi \models \varphi\}) \bar{\triangleright} d$.*

PROOF. We fix σ_C in (\mathcal{G}, q) , remove the epistemic relations, and merge the opponents Ag-C into a single agent. Notice that φ is either a reachability or a safety objective (i.e., of form $\varphi_1 \text{U} \varphi_2$ or $\varphi_1 \text{R} \varphi_2$ respectively). For reachability, we further redirect the transitions to a new “sink” state whenever the objective becomes unattainable (similarly to the construction in the proof of Proposition 1 below). This way, we obtain a Markov Decision Process in which we seek to optimize a reachability reward $T = p_2$, and there always exist deterministic memoryless policies that achieve the minimum and maximum probabilities of reaching T [39].

For safety objectives, we transform it to negation of reachability, by using the equivalence $\varphi_1 \text{R} \varphi_2 \equiv \neg(\neg \varphi_1 \text{U} \neg \varphi_2)$, and proceed analogously. \square

Now we can prove the upper bounds for simple PATL_{sr} formulas.

Proposition 1. *Checking formulas $\varphi = \langle\langle C \rangle\rangle^{\triangleright d} p_1 \text{U} p_2$ is in EXPTIME (with respect to the size of the model).*

PROOF. For the iCGS \mathcal{G} , given as input, first reconstruct it into \mathcal{G}' as follows:

- (i) Add a “sink” state q_{sink} with $\mathcal{G}, q_{\text{sink}} \not\models p_2$ and a self-loop as the only outgoing transition.
- (ii) For all the states q st. $\mathcal{G}, q \models p_2$ or $\mathcal{G}, q \not\models p_1$, remove all outgoing transitions and replace them with an automatic transition to q_{sink} . That is, we stop looking at the rest of the path whenever p_2 has been achieved (and thus $p_1 \text{U} p_2$ already succeeded) or p_1 has been invalidated (and thus $p_1 \text{U} p_2$ already failed).

Note that, on each path in \mathcal{G}' , p_2 can occur at most once. Moreover, the paths that reach p_2 are exactly the paths that satisfy $p_1\text{Up}_2$.

Secondly, formulate a set of constraints Φ as inequalities over the vectors of rewards $r_q^0 \in \mathbb{R}$, $r_q \in \mathbb{R}$ for $q \in St$, and probabilistic decisions $choice_{a,q,\alpha} \in \mathbb{R}$ for $a \in Ag$, $q \in St$, $\alpha \in Ac$. The value r_q^0 captures the immediate level of success at state q , r_q represents the expected probability of success from state q , and $choice_{a,q,\alpha}$ expresses the probability with which agent a takes action α at state q , i.e., represents the probabilistic choices of all the agents. The set of constraints Φ is built as follows:

- (i) For every $q \in St$, if $\mathcal{G}, q \models p_2$ then add constraint ($r_q^0 = 1$) to the set of constraints Φ , else add ($r_q^0 = 0$). That is, the immediate reward at q is 1 if p_2 has just been achieved, and 0 otherwise.⁴
- (ii) For every $q \in St$, add constraint

$$(r_q = r_q^0 + \sum_{\vec{\alpha} \in L_q} r_{\delta(q,\vec{\alpha})} \cdot \prod_{a \in Ag} choice_{a,q,\alpha_a}) \quad (3)$$

expressing that r_q is the sum of the immediate reward at s and the expected reward to be obtained in the future.

- (iii) Add ($choice_{a,q,\alpha} \geq 0$) for each $a \in Ag$, $q \in St$, $\alpha \in L(q,a)$, and ($\sum_{\alpha \in L(q,a)} choice_{a,q,\alpha} = 1$) for each $a \in Ag$, $q \in St$.
- (iv) For every coalition agent $a \in C$, states q, q' with $q \sim_a q'$, and action $\alpha \in L(q,a)$, add ($choice_{a,q,\alpha} = choice_{a,q',\alpha}$), expressing that the probabilistic choices of a at indistinguishable states q and q' must be the same.
- (v) Finally, add ($r_q \approx d$) for every $q \in St$ such that $q_0 \sim_a q$ for some $a \in C$, i.e., the expected probability of success from each state indistinguishable from q_0 is in relation \approx with value d .

By construction, the only value of r_q that satisfies the above constraints captures the expected probability of satisfying $p_1\text{Up}_2$ when the (memoryless probabilistic) choices of agents are given by the vector $choice$. Note that the agents in C are assumed to use memoryless choices by the semantics of $\langle\langle C \rangle\rangle^{\approx d} p_1\text{Up}_2$. Moreover, memoryless choices are sufficient for the opponents in Ag_{-C} by Lemma 1.

Now, checking if $\mathcal{G}, q_0 \models \langle\langle C \rangle\rangle^{\approx d} p_1\text{Up}_2$ is equivalent to deciding the following instance of $\text{Th}\mathbb{R}$:

$$\exists \{r_q \mid q \in St\} \exists \{choice_{a,q,\alpha} \mid a \in C, q \in St, \alpha \in Ac\} \quad (4)$$

$$\forall \{choice_{a,q,\alpha} \mid a \notin C, q \in St, \alpha \in Ac\} \bigwedge \Phi.$$

Note that the number of atomic predicates and the number of quantifiers in Φ are $m = n = O(|Ag| \cdot |St| \cdot |Ac|)$, the number of quantifier groups is $\omega = 2$ (equivalently, the number of quantifier alternations is 1), and the maximal degree of the polynomials is $d = 1$. By Theorem 5.2, the above instance of $\text{Th}\mathbb{R}$ can be decided in $n^n \cdot 2^{O(1)} + n^{O(n)} = 2^{O(n \cdot \log n)} = 2^{O(|Ag| \cdot |St| \cdot |Ac| \cdot \log(|Ag| \cdot |St| \cdot |Ac|))}$ steps. \square

Proposition 2. *Checking formulas $\varphi = \langle\langle C \rangle\rangle^{\approx d} p_1\text{Rp}_2$ is in EXPTIME (with respect to the size of the model).*

PROOF. Recall that $p_1\text{Rp}_2 \equiv \neg(\neg p_1\text{U}\neg p_2)$. Thus, we have $\mathcal{G}, q_0 \models \langle\langle C \rangle\rangle^{\approx d} p_1\text{Rp}_2$ iff $\mathcal{G}, q_0 \models \langle\langle C \rangle\rangle^{\approx (1-d)} (\neg p_1\text{U}\neg p_2)$, which can be verified in EXPTIME by Proposition 1. \square

⁴Note that an equality can be expressed as a pair of inequalities.

THEOREM 5.3. *Model checking PATL_{sr} with probabilistic strategies for the coalition is in EXPTIME.*

PROOF. To check if $\mathcal{G}, q \models \varphi$, we first transform the temporal and Boolean operators in φ to Negation Normal Form by using De Morgan laws and the duality laws for “until” U and “release” R. If the resulting formula contains no nested strategic modalities, then it can be model-checked in EXPTIME by Propositions 1 and 2 (the case of “next” is straightforward). For nested strategic modalities, we proceed recursively (bottom-up), which runs in time $\mathbf{P}^{\text{EXPTIME}} = \text{EXPTIME}$. \square

An interesting special case is when we only consider the abilities of all the agents cooperating on a common goal. Then, the verification problem is in PSPACE.

THEOREM 5.4. *Model checking PATL_{sr} with probabilistic strategies for the coalition and formulas that include only the grand coalition (Ag) or the empty coalition (\emptyset) is in PSPACE.*

PROOF. First, notice that $\langle\langle \emptyset \rangle\rangle$ is equivalent to “for all paths”, which reduces our model-checking problem to that of PCTL. For formulas of type $\langle\langle C \rangle\rangle^{\approx d} p_1\text{Up}_2$ and $\langle\langle C \rangle\rangle^{\approx d} p_1\text{Rp}_2$, notice that the universally quantified part in the embedding (4) presented in the proofs of Propositions 1 and 2, is in fact empty. Thus, the constructions define a reduction to the existential theory of the reals, which is in PSPACE. For nested strategic modalities, we proceed recursively, which obtains $\mathbf{P}^{\text{PSPACE}} = \text{PSPACE}$. \square

Thus, in particular, verification of probabilistic memoryless strategies in stochastic single agent iCGS is in PSPACE.

5.3 Probabilistic Strategies: Lower Bounds

THEOREM 5.5. *Model checking PATL_{sr} with probabilistic strategies for the coalition is Δ_2^{P} -hard.*

PROOF. The proof proceeds by a reduction of ATL_{ir} model checking, which is Δ_2^{P} -hard [48].

Consider a iCGS \mathcal{G} , a state q in it, and a formula $\langle\langle C \rangle\rangle\varphi$ of ATL_{ir} . Clearly, \mathcal{G} can be seen as a stochastic iCGS with only Dirac probability distributions for transitions. We begin by recalling that, in ATL_{ir} , it suffices to consider memoryless responses of the opponents. Formally, $\mathcal{G}, q \models_{\text{ATL}_{\text{ir}}} \langle\langle C \rangle\rangle\varphi$ iff there exists a deterministic memoryless strategy with imperfect information σ_C such that, for every deterministic memoryless strategy with perfect information $\sigma_{Ag_{-C}}$, φ holds on the sole path starting from q and consistent with $(\sigma_C, \sigma_{Ag_{-C}})$.⁵

Now, consider the PATL_{sr} evaluation of formula $\langle\langle C \rangle\rangle^=1\varphi$ in \mathcal{G}, q . First, observe that $\mathcal{G}, q \models \langle\langle C \rangle\rangle^=1\varphi$ iff C have a deterministic memoryless strategy to enforce φ with probability 1 against any response. To see this, assume that a probabilistic strategy σ_C enforces φ with probability 1. Then, every deterministic strategy in the support of σ_C also enforces φ with probability 1.

Secondly, by Lemma 1, it suffices to consider only deterministic memoryless strategies of the opponents. Thus, $\mathcal{G}, q \models \langle\langle C \rangle\rangle^=1\varphi$ iff

⁵This follows from the fact that the semantics of ATL with memoryless and perfect recall strategies coincide for agents with perfect information [3].

C have a deterministic memoryless strategy with imperfect information σ_C that enforces φ with probability 1 against every deterministic memoryless strategy with perfect information σ_{Ag_-C} .

Thirdly, the outcome of a deterministic strategy σ_C and counter-strategy σ_{Ag_-C} from state q is always a single path. Thus, enforcing with probability 1 is equivalent to enforcing on that path.

Summing up, $\mathcal{G}, q \models_{\text{ATL}_{\text{ir}}} \langle\langle C \rangle\rangle \varphi$ iff $\mathcal{G}, q \models_{\text{PATL}_{\text{sr}}} \langle\langle C \rangle\rangle^1 \varphi$, which provides a one-to-one polynomial-time reduction from model checking of ATL_{ir} to model checking of PATL_{sr} . \square

5.4 Model Checking Objective Ability

In the previous subsections, we have proved that model checking PATL w.r.t. memoryless probabilistic strategies is between Δ_2^{P} and EXPTIME for the subjective interpretation of ability under imperfect information. Moreover, it is between Δ_2^{P} and PSPACE for stochastic single-agent systems. Now, we show that the same results apply to the objective variant of probabilistic ability.

Proposition 3. *Checking formulas $\varphi = \langle\langle C \rangle\rangle^{\text{p-ad}} p_1 U p_2$ and $\varphi = \langle\langle C \rangle\rangle^{\text{p-ad}} p_1 R p_2$ is in EXPTIME with respect to the size of the model.*

PROOF. Analogously to Proposition 1 and 2. The sole difference is that, in the construction for $\langle\langle C \rangle\rangle^{\text{p-ad}} p_1 U p_2$, only the constraint $(r_{q_0} \text{p-ad } d)$ for the objective initial state q_0 is added to Φ in point (v), instead of all the indistinguishable states. \square

THEOREM 5.6. *Model checking PATL_{or} w.r.t. probabilistic strategies for the coalition is in EXPTIME .*

Moreover, model checking PATL_{or} w.r.t. probabilistic strategies for the coalition and formulas that include only the grand coalition (Ag) or the empty coalition (\emptyset) is in PSPACE .

PROOF. Analogous to the proofs of Theorem 5.3 and 5.4. \square

THEOREM 5.7. *Model checking PATL_{or} with probabilistic strategies for the coalition is Δ_2^{P} -hard.*

PROOF. The proof proceeds by a reduction of ATL_{ir} model checking, which is Δ_2^{P} -hard [48]. We observe that the objective and subjective semantics of ability coincide for the models used in the reduction of SNSAT_2 in [48], and proceed as in Theorem 5.5 \square

We finally comment on the model checking complexity for the logic PATL_{orK} , i.e., PATL_{or} extended with epistemic operators K . Recall that, in contrast to PATL_{sr} , epistemic operators cannot be expressed in PATL_{or} . However, model checking of observational knowledge is in P w.r.t. the size of the model and the length of the model. Thus, the results in Theorems 5.6 and 5.7 carry over to the broader language of PATL_{orK} .

5.5 Beyond PATL

In this section, we make the first step toward establishing the complexity of model-checking for PATL^* with memoryless strategies and imperfect information. In particular, we show that the problem for *memoryless deterministic strategies of the coalition* against probabilistic play of the other agents and a stochastic environment is no more complex than in standard (non-probabilistic) case.

THEOREM 5.8. *Model checking $\text{PATL}_{\text{sr}}^*$ and $\text{PATL}_{\text{or}}^*$ with deterministic strategies for the coalition is PSPACE -complete.*

PROOF. The lower bound follows from the corresponding problem for ATL_{ir}^* , which is also PSPACE -complete.

As for the upper bound, we apply the analogous procedure to model checking of ATL_{ir}^* : for formulas of type $\langle\langle C \rangle\rangle^{\text{p-ad}} \varphi$, we guess a strategy and prune the model accordingly. Then, we check the PCTL^* formula $A^{\text{p-ad}} \varphi$. This procedure gives an algorithm in $\text{NPSpace} = \text{PSPACE}$ (see [5, Theorem 9]). \square

We also speculate that model checking of $\text{PATL}_{\text{sr}}^*$ and $\text{PATL}_{\text{or}}^*$ with probabilistic strategies is between PSPACE and 2EXPTIME . The lower bound follows from an embedding of LTL model checking. For the upper bound, the idea is to extend the construction in Proposition 1 to prefixes of paths that are sufficient to determine what fraction of their infinite extensions satisfy the given LTL objective. It is known that, to determine the *existence* of such extension, it suffices to consider prefixes of length which is polynomial in the size of the model and exponential in the size of the formula [15]. If the same can be proved for bounded model checking of probabilistic LTL objectives, we could combine it with our translation to $\text{Th}\mathbb{R}$, and obtain the inclusion in 2EXPTIME . Moreover, for formulas of bounded length the problem would be in EXPTIME w.r.t. the size of the model. However, the leap from possibilistic to probabilistic bounded model checking for LTL is nontrivial, and it remains to be seen if our proof idea actually works.

6 CONCLUSION

This paper advances the research on the verification of MAS under two combined types of uncertainty: first, the qualitative uncertainty about the local state and second, the quantitative uncertainty about the occurrence of future events. Although the resulting setting is often the case in real-world scenarios, “little progress has been made on developing practical, approximate verification and strategy synthesis algorithms” for stochastic MAS, as noticed by Kwiatkowska et al. [57]. To capture this setting, we have considered the probabilistic logics PATL and PATL^* under imperfect information. We provided novel decidability and model-checking results for memoryless strategies. We have considered two semantic variations for the logics, the objective and subjective interpretations, as well as the cases whereby agents are allowed to play probabilistic versus deterministic strategies. In particular, we have shown that model-checking of PATL when agents play probabilistic memoryless strategies can be done in EXPTIME . We have also shown that the problem is PSPACE -complete for PATL^* when the proponent coalition is restricted to deterministic strategies. For future work, we intend to explore the challenging case of model-checking PATL^* when the proponent coalition plays probabilistic strategies.

ACKNOWLEDGMENTS

This research has been supported by the EU Horizon 2020 Marie Skłodowska-Curie project with grant agreement No 101105549, the PNRR MUR project PE0000013-FAIR, the PRIN 2020 project RIPER, as well as by NCBR Poland and FNR Luxembourg under the PolLux/FNR-CORE project SpaceVote (POLLUX-XI/14/SpaceVote/2023 and C22/IS/17232062/SpaceVote). For the purpose of open access, and in fulfilment of the obligations arising from the grant agreement, the authors have applied CC BY 4.0 license to any Author Accepted Manuscript version arising from this submission.

REFERENCES

- [1] T. Ágotnes, V. Goranko, W. Jamroga, and M. Wooldridge. 2015. Knowledge and Ability. In *Handbook of Epistemic Logic*, H.P. van Ditmarsch, J.Y. Halpern, W. van der Hoek, and B.P. Kooi (Eds.). College Publications, 543–589.
- [2] F.A. Al-Khayyal. 1990. Jointly constrained bilinear programs and related problems: An overview. *Computers & Mathematics with Applications* 19, 11 (1990), 53–62. [https://doi.org/10.1016/0898-1221\(90\)90148-D](https://doi.org/10.1016/0898-1221(90)90148-D)
- [3] R. Alur, T.A. Henzinger, and O. Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (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 *Proc. of IJCAI 2019*. ijcai.org, 32–38.
- [5] Adnan Aziz, Vigyan Singhal, Felice Balarin, Robert K. Brayton, and Alberto L. Sangiovanni-Vincentelli. 1995. It usually works: The temporal logic of stochastic systems. In *Computer Aided Verification*, Pierre Wolper (Ed.), Springer Berlin Heidelberg, Berlin, Heidelberg, 155–165.
- [6] Haris Aziz. 2019. A probabilistic approach to voting, allocation, matching, and coalition formation. *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers* (2019), 45–50.
- [7] Haris Aziz, Felix Brandt, Edith Elkind, and Piotr Skowron. 2019. *Computational Social Choice: The First Ten Years and Beyond*. Springer International Publishing, Cham, 48–65. https://doi.org/10.1007/978-3-319-91908-9_4
- [8] Paolo Ballarini, Michael Fisher, and Michael Wooldridge. 2009. Uncertain Agent Verification through Probabilistic Model-Checking. In *Safety and Security in Multiagent Systems*.
- [9] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, and Aniello Murano. 2023. Strategic Abilities of Forgetful Agents in Stochastic Environments. In *KR*. 726–731.
- [10] 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).
- [11] Raphaël Berthon, Nathanaël Fijalkow, Emmanuel Filiot, Shibashis Guha, Bastien Maubert, Aniello Murano, Laureline Pinault, Sophie Pinchinat, Sasha Rubin, and Olivier Serre. 2020. Alternating Tree Automata with Qualitative Semantics. *ACM Trans. Comput. Logic* 22, 1 (2020), 1–24.
- [12] Raphaël Berthon, Joost-Pieter Katoen, Munyque Mittelmann, and Aniello Murano. 2024. Natural Strategic Ability in Stochastic Multi-Agent Systems. In *Proc. of the AAAI Conference on Artificial Intelligence, AAAI 2024*. AAAI Press.
- [13] Raphaël Berthon, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Y. Vardi. 2021. Strategy Logic with Imperfect Information. *ACM Trans. Comput. Log.* 22, 1 (2021), 1–51.
- [14] Dietmar Berwanger and Laurent Doyen. 2008. On the Power of Imperfect Information. In *Proc. of FSTTCS 2008 (LIPIcs, Vol. 2)*, Ramesh Hariharan, Madhavan Mukund, and V. Vinay (Eds.), 73–82.
- [15] A. Biere, A. Cimatti, E.M. Clarke, and Y. Zhu. 1999. Symbolic Model Checking Without BDDs. In *Proceedings of Tools and Algorithms for Construction and Analysis of Systems (TACAS) (Lecture Notes in Computer Science, Vol. 1579)*. Springer, 193–207.
- [16] Roderick Bloem, Barbara Jobstmann, Nir Piterman, Amir Pnueli, and Yaniv Sa’ar. 2012. Synthesis of reactive (1) designs. *J. Comput. System Sci.* 78, 3 (2012), 911–938.
- [17] Craig Boutilier. 1999. Sequential Optimality and Coordination in Multiagent Systems. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*. 478–485.
- [18] Felix Brandt. 2019. Collective choice lotteries: Dealing with randomization in economic design. *The Future of Economic Design: The Continuing Development of a Field as Envisioned by Its Researchers* (2019), 51–56.
- [19] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. 2016. *Handbook of computational social choice*. Cambridge University Press.
- [20] Nils Bulling and Wojciech Jamroga. 2009. What Agents Can Probably Enforce. *Fundam. Informaticae* 93, 1–3 (2009), 81–96.
- [21] Nils Bulling and Wojciech Jamroga. 2011. Alternating epistemic Mu-calculus. In *Proc. of IJCAI 2011*. IJCAI/AAAI, 109–114.
- [22] N. Bulling and W. Jamroga. 2014. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems* 28, 3 (2014), 474–518.
- [23] John F. Canny. 1988. Some Algebraic and Geometric Computations in PSPACE. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*. ACM, 460–467. <https://doi.org/10.1145/62212.62257>
- [24] Arnaud Carayol, Christof Löding, and Olivier Serre. 2018. Pure strategies in imperfect information stochastic games. *Fundamenta Informaticae* 160, 4 (2018), 361–384.
- [25] Petr Cermák, Alessio Lomuscio, Fabio Mogavero, and Aniello Murano. 2018. Practical verification of multi-agent systems against SLK specifications. *Inf. Comput.* 261 (2018), 588–614.
- [26] Krishnendu Chatterjee and Laurent Doyen. 2014. Partial-observation stochastic games: How to win when belief fails. *ACM Transactions on Computational Logic (TOCL)* 15, 2 (2014), 1–44.
- [27] K. Chatterjee, T. A. Henzinger, and N. Piterman. 2010. Strategy Logic. *Inf. Comput.* 208, 6 (2010), 677–693. <https://doi.org/10.1016/j.ic.2009.07.004>
- [28] Taolue Chen, Marta Kwiatkowska, David Parker, and Aistis Simaitis. 2011. Verifying team formation protocols with probabilistic model checking. In *Proc. of CLIMA 2011 (LNCS 6814)*. Springer, 190–207.
- [29] Taolue Chen and Jian Lu. 2007. Probabilistic alternating-time temporal logic and model checking algorithm. In *Proc. of FSKD*. 35–39.
- [30] E.M. Clarke, O. Grumberg, D. Kroening, D. Peled, and H. Veith. 2018. *Model checking*. MIT press.
- [31] Edmund M Clarke and E Allen Emerson. 1981. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Workshop on logic of programs*. Springer, 52–71.
- [32] Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem. 2018. *Handbook of Model Checking* (1st ed.). Springer Publishing Company, Incorporated.
- [33] C. Dima and F.L. Tiplea. 2011. Model-checking ATL under Imperfect Information and Perfect Recall Semantics is Undecidable. *CoRR* abs/1102.4225 (2011), 17.
- [34] Laurent Doyen. 2022. Stochastic Games with Synchronizing Objectives. In *Proc. of LICS (Haifa, Israel) (LICS ’22)*. Association for Computing Machinery, New York, NY, USA, Article 43, 12 pages. <https://doi.org/10.1145/3531130.3532439>
- [35] Laurent Doyen and Jean-François Raskin. 2011. Games with imperfect information: theory and algorithms. *Lectures in Game Theory for Computer Scientists* 10 (2011).
- [36] E Allen Emerson and Joseph Y Halpern. 1986. “Sometimes” and “not never” revisited: on branching versus linear time temporal logic. *Journal of the ACM (JACM)* 33, 1 (1986), 151–178.
- [37] Ronald Fagin, Joseph Y Halpern, Yoram Moses, and Moshe Vardi. 2004. *Reasoning about knowledge*. MIT press.
- [38] Fei Fang and Thanh H. Nguyen. 2016. Green Security Games: Apply Game Theory to Addressing Green Security Challenges. *SIGecom Exch.* 15, 1 (sep 2016), 78–83. <https://doi.org/10.1145/2994501.2994507>
- [39] Vojtech Forejt, Marta Z. Kwiatkowska, Gethin Norman, and David Parker. 2011. Automated Verification Techniques for Probabilistic Systems. In *Formal Methods for Eternal Networked Software Systems - Proceedings of SFM (Lecture Notes in Computer Science, Vol. 6659)*. Springer, 53–113. https://doi.org/10.1007/978-3-642-21455-4_3
- [40] Chen Fu, Andrea Turrini, Xiaowei Huang, Lei Song, Yuan Feng, and Lijun Zhang. 2018. Model checking probabilistic epistemic logic for probabilistic multiagent systems. In *Proc. of IJCAI 2018*. 4757–4763.
- [41] Vincent Gripon and Olivier Serre. 2009. Qualitative concurrent stochastic games with imperfect information. In *Proc. of ICALP 2009*. Springer, 200–211.
- [42] Dilian Gurov, Valentin Goranko, and Edvin Lundberg. 2022. Knowledge-based strategies for multi-agent teams playing against Nature. *Artificial Intelligence* 309 (2022), 103728.
- [43] Jianye Hao, Songzheng Song, Yang Liu, Jun Sun, Lin Gui, Jin Song Dong, and Ho-fung Leung. 2012. Probabilistic model checking multi-agent behaviors in dispersion games using counter abstraction. In *Proc. of PRIMA 2012 (LNCS 7455)*. Springer, 16–30.
- [44] Xiaowei Huang and Cheng Luo. 2013. A logic of probabilistic knowledge and strategy. In *Proc. of AAMAS 2013*. 845–852.
- [45] Xiaowei Huang, Kaile Su, and Chenyi Zhang. 2012. Probabilistic Alternating-Time Temporal Logic of Incomplete Information and Synchronous Perfect Recall. In *Proc. of AAAI 2012*. 765–771.
- [46] W. Jamroga and T. Ágotnes. 2007. Constructive knowledge: what agents can achieve under imperfect information. *J. Applied Non-Classical Logics* 17, 4 (2007), 423–475.
- [47] W. Jamroga and N. Bulling. 2011. Comparing variants of strategic ability. In *Proc. of IJCAI 2011*. IJCAI/AAAI, 252–257.
- [48] Wojciech Jamroga and Jürgen Dix. 2006. Model Checking Abilities under Incomplete Information Is Indeed Delta2-complete. In *Proc. of EUMAS 2006 (CEUR 223)*. CEUR-WS.org.
- [49] Sebastian Junges, Nils Jansen, Ralf Wimmer, Tim Quatmann, Leonore Winterer, Joost-Pieter Katoen, and Bernd Becker. 2018. Finite-State Controllers of POMDPs using Parameter Synthesis. In *Proceedings of UAI*. AUAI Press, 519–529.
- [50] Debarun Kar, Thanh H. Nguyen, Fei Fang, Matthew Brown, Arunesh Sinha, Milind Tambe, and Albert Xin Jiang. 2018. *Trends and Applications in Stackelberg Security Games*. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-319-44374-4_27
- [51] John G Kemeny, J Laurie Snell, and Anthony W Knapp. 1976. *Stochastic Processes*. In *Denumerable Markov Chains*. Springer, 40–57.
- [52] Christopher Kiekintveld, Manish Jain, Jason Tsai, James Pita, Fernando Ordóñez, and Milind Tambe. 2009. Computing optimal randomized resource allocations for massive security games. In *AAMAS (1)*. IFAAMAS, 689–696.
- [53] Dmytro Korzhuk, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. 2011. Stackelberg vs. Nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research* 41 (2011), 297–327.

- [54] Orna Kupferman and Moshe Y Vardi. 1997. Module checking revisited. In *Computer Aided Verification: 9th International Conference, CAV'97 Haifa, Israel, June 22–25, 1997 Proceedings 9*. Springer, 36–47.
- [55] O. Kupferman and M. Y. Vardi. 2000. Synthesis with incomplete informatio. In *Advances in Temporal Logic*. Springer, Berlin, 109–127.
- [56] Orna Kupferman, Moshe Y Vardi, and Pierre Wolper. 2001. Module checking. *Information and Computation* 164, 2 (2001), 322–344.
- [57] Marta Kwiatkowska, Gethin Norman, and David Parker. 2022. Probabilistic model checking and autonomy. *Annual Review of Control, Robotics, and Autonomous Systems* 5 (2022), 385–410.
- [58] Marta Kwiatkowska, Gethin Norman, David Parker, Gabriel Santos, and Rui Yan. 2022. Probabilistic Model Checking for Strategic Equilibria-based Decision Making: Advances and Challenges. In *Proc. of MFCS 2022*. 4–22.
- [59] F. Laroussinie and N. Markey. 2015. Augmenting ATL with strategy contexts. *Inf. Comput.* 245 (2015), 98–123. <https://doi.org/10.1016/j.ic.2014.12.020>
- [60] François Laroussinie and Nicolas Markey. 2015. Augmenting ATL with strategy contexts. *Information and Computation* 245 (2015), 98–123.
- [61] Alessio Lomuscio and Edoardo Pirovano. 2020. Parameterised verification of strategic properties in probabilistic multi-agent systems. In *Proc. of AAMAS 2020*. 762–770.
- [62] Jian Lou, Andrew M Smith, and Yevgeniy Vorobeychik. 2017. Multidefender security games. *IEEE Intelligent Systems* 32, 1 (2017), 50–60.
- [63] B. Maubert, M. Mittelmann, A. Murano, and L. Perrussel. 2021. Strategic Reasoning in Automated Mechanism Design. In *KR-21*.
- [64] Reshef Meir. 2017. Iterative voting. *Trends in computational social choice* 4 (2017), 69–86.
- [65] Munyque Mittelmann, Bastien Maubert, Aniello Murano, and Laurent Perrussel. 2023. Formal Verification of Bayesian Mechanisms. In *Proc. of AAAI 2023*.
- [66] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. 2012. What Makes ATL* Decidable? A Decidable Fragment of Strategy Logic. In *CONCUR*.
- [67] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi. 2014. Reasoning About Strategies: On the Model-Checking Problem. *ACM Trans. Comput. Log.* 15, 4 (2014).
- [68] Chunyan Mu and Jun Pang. 2023. On Observability Analysis in Multiagent Systems. In *ECAI (Frontiers in Artificial Intelligence and Applications, Vol. 372)*. IOS Press, 1755–1762.
- [69] Hoang Nga Nguyen and Abdur Rakib. 2019. A Probabilistic Logic for Resource-Bounded Multi-Agent Systems.. In *Proc. of IJCAI 2019*. 521–527.
- [70] Amir Pnueli. 1977. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*. iee, 46–57.
- [71] A. Pnueli and R. Rosner. 1989. On the Synthesis of a Reactive Module.. In *Symposium on the Principles of Programming Languages (POPL 1989)*. ACM, New York, 179–190.
- [72] Martin L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- [73] J. H. Reif. 1984. The Complexity of Two-Player Games of Incomplete Information. *J. Comput. System Sci.* 29, 2 (1984), 274–301.
- [74] James Renegar. 1992. On the Computational Complexity and Geometry of the First-Order Theory of the Reals, Part I: Introduction. Preliminaries. The Geometry of Semi-Algebraic Sets. The Decision Problem for the Existential Theory of the Reals. *J. Symb. Comput.* 13, 3 (1992), 255–300. [https://doi.org/10.1016/S0747-7171\(10\)80003-3](https://doi.org/10.1016/S0747-7171(10)80003-3)
- [75] Marcus Schaefer and Daniel Stefankovic. 2022. Beyond the Existential Theory of the Reals. *CoRR* abs/2210.00571 (2022). <https://doi.org/10.48550/arXiv.2210.00571>
- [76] Pierre-Yves Schobbens. 2004. Alternating-time logic with imperfect recall. *Electr. Notes Theor. Comput. Sci.* 85, 2 (2004), 82–93. [https://doi.org/10.1016/S1571-0661\(05\)82604-0](https://doi.org/10.1016/S1571-0661(05)82604-0)
- [77] Sanjit A Seshia, Dorsa Sadigh, and S Shankar Sastry. 2022. Toward verified artificial intelligence. *Commun. ACM* 65, 7 (2022), 46–55.
- [78] Fu Song, Yedi Zhang, Taolue Chen, Yu Tang, and Zhiwu Xu. 2019. Probabilistic alternating-time μ -calculus. In *Proc. of AAAI 2019*. 6179–6186.
- [79] Wolfgang Thomas. 1995. On the synthesis of strategies in infinite games. In *Annual Symposium on Theoretical Aspects of Computer Science*. Springer, 1–13.
- [80] Nikos Vlassis, Michael L. Littman, and David Barber. 2012. On the Computational Complexity of Stochastic Controller Optimization in POMDPs. *ACM Trans. Comput. Theory* 4, 4 (2012), 12:1–12:8. <https://doi.org/10.1145/2382559.2382563>
- [81] Wei Wan, Jamal Bentahar, and Abdessamad Ben Hamza. 2013. Model checking epistemic–probabilistic logic using probabilistic interpreted systems. *Knowledge-Based Systems* 50 (2013), 279–295.
- [82] Binru Wang, Yuan Zhang, Zhi-Hua Zhou, and Sheng Zhong. 2019. On repeated stackelberg security game with the cooperative human behavior model for wildlife protection. *Applied Intelligence* 49 (2019), 1002–1015.