Probabilistic Alternating-Time μ-Calculus*

Fu Song[1], Yedi Zhang[1], Taolue Chen[2,5], Yu Tang[3], and Zhiwu Xu[4†]

[1]School of Information Science and Technology, ShanghaiTech University, Shanghai, China
[2]Department of Computer Science and Information Systems, Birkbeck, University of London, UK
[3]School of Computer Science and Software Engineering, East China Normal University, Shanghai, China
[4]College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, China
[5]State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

Abstract

Reasoning about strategic abilities is key to an AI system consisting of multiple agents with random behaviors. We propose a probabilistic extension of Alternating μ-Calculus (AMC), named PAMC, for reasoning about strategic abilities of agents in stochastic multi-agent systems. PAMC subsumes existing logics AMC and PμTL. The usefulness of PAMC is exemplified by applications in genetic regulatory networks. We show that, for PAMC, the model checking problem is in UP∩co-UP, and the satisfiability problem is EXPTIME-complete, both of which are the same as those for AMC. Moreover, PAMC admits the small model property. We implement the satisfiability checking procedure in a tool PAMCSolver.

Introduction

Temporal logics play a key role in specification and verification of ICT systems. There are two fundamental decision problems of temporal logics: satisfiability checking and model checking. Given a temporal logic formula, the former asks whether a system satisfying the formula exists, while the latter asks whether a further given system satisfies the formula. Temporal logics for reasoning about strategic abilities in Multi-Agent Systems (MAS) have been proposed, typically in the form of *alternating-time* temporal logics (e.g., ATL, ATL* and AMC) (Alur, Henzinger, and Kupferman 2002), and strategic logics (Chatterjee, Henzinger, and Piterman 2010; Mogavero et al. 2014; 2017). Both model checking and satisfiability checking for these logics have been extensively investigated (Schewe and Finkbeiner 2006; Walther et al. 2006; Schewe 2008).

In practice, agents or the environment may exhibit random behaviors because of unpredictable physical conditions. Hence, it is vital to reason about strategic abilities of agents in a *stochastic* setting. From the modeling perspective, this gives rise to stochastic MAS, which consist of a set of agents operating concurrently in a stochastic environment. Probabilistic variants of ATL and ATL*, for instance PATL, PATL* (Chen and Lu 2007; Chen et al. 2013) and SGL (Baier et al. 2007), have been proposed, for which the model checking problem was also studied. In contrast, the satisfiability problem for these logics turns out to be much more difficult. Indeed, the satisfiability problem for PCTL, a (proper) fragment of these logics, is a long-standing open problem in the formal verification community (Chakraborty and Katoen 2016).

This paper aims to propose a temporal logic with *decidable* satisfiability checking for reasoning about stochastic MAS. Apart from being theoretical appealing, admitting decidability for satisfiability is useful in practice. Examples include debugging specifications (as writing formal specifications is often error-prone (Rozier and Vardi 2010)), social procedure or mechanism design (Pauly 2011), and assertion-based design (i.e., verifying consistency of all requirements by formalizing designer's intention by assertions, and checking the full set of assertions to be satisfied before the model is ready for examination (Foster, Krolnik, and Lacey 2004)). To this end, we propose a probabilistic variant of AMC, *probabilistic alternating-time μ-calculus* (PAMC), which is acquired by equipping the coalition modalities with (*qualitative*) probability quantifiers $\langle\langle A\rangle\rangle^{\bowtie k}$ in AMC. For instance, $\nu Z(a \wedge \langle\langle\{1,2\}\rangle\rangle^{\geq 0.9} Z)$ states that the agent group $\{1,2\}$ has a coalition strategy at each step such that the opponent coalition can escape the a-region in one step with probability less than 0.1. PAMC subsumes both AMC and PμTL (Liu et al. 2015), but is incomparable with PATL and PATL* mentioned before. Remark that we do not extend AMC in a *quantitative* way such as (Mio 2012), in order to retain decidability.

We investigate the model checking and the satisfiability checking problems of PAMC. The former can be done by an adaptation of the algorithm for AMC. As a result, the model checking problem of PAMC is shown in UP∩co-UP. In a sharp contrast, for satisfiability checking the presence of coalition modalities and probabilistic quantifiers pose new challenges. In particular, the decision procedures for (non-probabilistic) alternating-time logics (Schewe and Finkbeiner 2006; Walther et al. 2006; Schewe 2008) and PμTL (Chakraborty and Katoen 2016) cannot be directly adapted. To solve this problem, we propose a novel reduction from the satisfiability problem to solving (turn-based two-player) parity games, resulting in an EXPTIME decision procedure and a *small model property* in the sense that every satisfiable formula ϕ has a model of size exponential in $|\phi|$ (cf. Theorem 3). Remarkably, the com-

*This work is partially supported by NSFC (No. 61532019, 61761136011, and 61502308), EPSRC (EP/P00430X/1) and ARC (DP160101652, DP180100691).

†Corresponding author. Email: xuzhiwu@szu.edu.cn.

plexities of the model checking and satisfiability problems of PAMC lie in the same complexity classes of AMC. We implement a satisfiability solver for PAMC, which is, to our best knowledge, the first satisfiability checker for (probabilistic) alternating-time temporal logics.

Due to space restriction, proofs and experimental results are given in the full version, which, together with our tool PAMC and benchmarks, is available at `http://faculty.sist.shanghaitech.edu.cn/faculty/songfu/Projects/PAMCSolver`.

Probabilistic Concurrent Game Structures

For a natural number $k \in \mathbb{N}$, let $[k]$ denote the set $\{1, ..., k\}$. Given a probability distribution $\mathsf{Pr} : X \to [0, 1]$, let $\mathsf{supp}(\mathsf{Pr})$ denote the set $\{x \in X \mid \mathsf{Pr}(x) > 0\}$. We denote by $\mathcal{D}(X)$ the set of probability distributions on X.

Probabilistic concurrent game structures are a model of concurrent stochastic multi-agent systems, in which the transition probability function gives a probability distribution over the successor states after considering a joint action from the agents. In this work, we consider stochastic multi-agent systems with perfect information only.

Definition 1. *Fix a finite set* AP *of* atomic propositions *(i.e., observations). A* probabilistic concurrent game structure *(PCGS) is a tuple* $\mathcal{M} = (\mathsf{Ag}, \mathsf{Act}, Q, \Gamma, \delta, \lambda, q^0)$*, where:* $\mathsf{Ag} = [n]$ *is a finite set of* agents. $\mathsf{Act} = \bigcup_{i\in\mathsf{Ag}} \mathsf{Act}_i$ *and* Act_i *is a non-empty set of actions of agent* $i \in \mathsf{Ag}$*. A* decision $\mathsf{d} = \langle a_1, ..., a_n\rangle$ *is a joint action of all the agents in which* a_i*, denoted by* $\mathsf{d}(i)$*, is the action chosen by agent i. We denote by* D *the set of decisions.* Q *is a finite set of* states *and* $q^0 \in Q$ *is the initial state.* $\Gamma = (\Gamma_i)_{i\in\mathsf{Ag}}$ *with* $\Gamma_i : Q \to 2^{\mathsf{Act}_i}$ *assigning to agent i at a state q a finite set* $\Gamma_i(q) \subseteq \mathsf{Act}_i$ *of actions available at the state q. We denote by* $\mathsf{D}(q)$ *the set* $\prod_{i\in\mathsf{Ag}} \Gamma_i(q)$ *of available decisions at q.* $\delta : Q \times \mathsf{D} \to \mathcal{D}(Q)$ *is a (partial)* probability transition function *which associates each pair* $(q, \mathsf{d}) \in Q\times\mathsf{D}$ *such that* $\mathsf{d} \in \mathsf{D}(q)$ *with a probability distribution over Q.* $\lambda : Q \to 2^{\mathsf{AP}}$ *is the* labeling function *assigning to each state* $q \in Q$ *a set of atomic propositions.*

Given a state $q \in Q$ and a decision $\mathsf{d} \in \mathsf{D}(q)$, $\mathcal{M}$ moves to a next state $q' \in Q$ with probability $\delta(q, \mathsf{d})(q')$, which is also denoted as $\delta(q, \mathsf{d}, q')$. We indicate by $\mathsf{supp}(q, \mathsf{d})$ the set $\mathsf{supp}(\delta(q, \mathsf{d}))$. The *outdegree* of $\mathcal{M}$ is the maximum of $|\mathsf{supp}(q, \mathsf{d})|$ for $q \in Q$ and $\mathsf{d} \in \mathsf{D}(q)$.

Paths and tracks. A *path* $\pi \in Q^\omega$ (a.k.a. play) is an infinite sequence of states $q_0q_1...$ such that for all $j \geq 1$, $q_j \in \mathsf{supp}(q_{j-1}, \mathsf{d}_{j-1})$ for some $\mathsf{d}_{j-1} \in \mathsf{D}(q_{j-1})$. A *track* is a finite prefix of a path. Given a track $\pi = q_0q_1...q_m$ (resp. a path $\pi = q_0q_1...$) and $0 \leq i \leq m$ for track, let π_i denote the state q_i and $\pi_{\leq i}$ denote $q_0...q_i$.

Strategies. A *(mixed) strategy* of agent $i \in \mathsf{Ag}$ is a function $\theta_i : Q^+ \to \mathcal{D}(\mathsf{Act}_i)$ that assigns to each track $q_0 \ldots q_m$, representing the past history of the game, a probability distribution on its available actions $\Gamma_i(q_m) \subseteq \mathsf{Act}_i$. Therefore, the choice of the next action can be *history-dependent* and *mixed*. A strategy θ_i is *pure* if for all $\pi \in Q^+$, $\theta_i(\pi)$ is a *Dirac distribution*, i.e., $|\mathsf{supp}(\theta_i(\pi))| = 1$. A strategy θ_i is *memoryless* (i.e., positional or imperfect recall) if for all $\pi, \pi' \in Q^+$ and $q \in Q$, $\theta_i(\pi q) = \theta_i(\pi' q)$, namely, the agent takes an action only depending on the last state. We denote by Θ_i the set of all strategies for agent i and let $\Theta = \bigcup_{i\in\mathsf{Ag}} \Theta_i$.

A *coalition* is a set of agents $A \subseteq \mathsf{Ag}$. A *coalition strategy* for A is a function $\upsilon_A : A \to \Theta$ assigning to each agent $i \in A$ a strategy $\upsilon_A(i) \in \Theta_i$. Let Υ_A denote the set of all coalition strategies for the coalition A. We denote by $\overline{A}$ the set $\mathsf{Ag} \setminus A$.

Outcomes. Given a track $\pi \in Q^+$ such that the last state of π is q and a decision $\mathsf{d} \in \mathsf{D}(q)$, we denote by $\mathsf{Pr}^{\upsilon_A,\upsilon_{\overline{A}}}(\pi, \mathsf{d})$ the probability $\prod_{i\in A} \upsilon_A(i)(\pi)(\mathsf{d}(i))\cdot\prod_{i\in\overline{A}} \upsilon_{\overline{A}}(i)(\pi)(\mathsf{d}(i))$, that is, the probability of the decision d chosen by the agents Ag with respect to π, υ_A and $\upsilon_{\overline{A}}$. We say that a path π is *compatible* with respect to υ_A and $\upsilon_{\overline{A}}$, if for all $j \geq 0$, there is a decision $\mathsf{d}_j \in \mathsf{D}(\pi_j)$ such that $\delta(\pi_j, \mathsf{d}_j, \pi_{j+1}) > 0$ and $\mathsf{Pr}^{\upsilon_A,\upsilon_{\overline{A}}}(\pi_{\leq j}, \mathsf{d}_j) > 0$. We denote by $O_q^{\upsilon_A,\upsilon_{\overline{A}}}$ the set of paths starting from q that are compatible with respect to υ_A and $\upsilon_{\overline{A}}$, which is the set of paths that can be followed by the game when agents enforce strategies υ_A and $\upsilon_{\overline{A}}$.

Probability Space. A σ*-algebra* over a set Ω is a set $\mathcal{E} \subseteq 2^\Omega$ such that $\mathcal{E}$ contains $\emptyset$ and is closed under countable union and complement. An element of $\mathcal{E}$ is called an *event*. A *probability space* is a triple $(\Omega, \mathcal{E}, \mathsf{Pr})$, where Ω is a *sample space*, $\mathcal{E}$ is a σ-algebra over Ω, $\mathsf{Pr} : \mathcal{E} \to [0, 1]$ is a *probability measure* such that $\mathsf{Pr}(\Omega) = 1$ and the countable additivity property is satisfied, i.e., $\mathsf{Pr}(U\cup V) = \mathsf{Pr}(U)+\mathsf{Pr}(V)$ whenever $U \cap V = \emptyset$.

Given a coalition $A \subseteq \mathsf{Ag}$, a state q, *two* coalition strategies υ_A and $\upsilon_{\overline{A}}$, one can construct a probability space over the set of paths $O_q^{\upsilon_A,\upsilon_{\overline{A}}}$ with the probability measure $\mathsf{Pr}_q^{\upsilon_A,\upsilon_{\overline{A}}}$ defined in the standard way (Vardi 1985), where an event is a measurable set of paths.

CGS, MDP and MC. A *concurrent game structure* (CGS) is a PCGS such that all the probability distributions involved in the PCGS (i.e., probability transition function and strategies) are Dirac distributions. A Markov decision process (MDP) is a PCGS such that $|\mathsf{Ag}| = 1$. A *Markov chain* (MC) is an MDP such that $|\mathsf{Act}| = 1$.

Probabilistic Alternating-Time μ-Calculus

Probabilistic alternating-time μ-calculus is a simple and succinct, but natural, *probabilistic* extension of AMC (Alur, Henzinger, and Kupferman 2002). In this logic, coalition modalities $\langle\langle A\rangle\rangle\phi$ from AMC are replaced with *probabilistic* coalition modalities $\langle\langle A\rangle\rangle^{\bowtie k}\phi$, which probabilistically quantify over the strategic choices of a group A of agents.

Definition 2. *Let* $\mathcal{Z}$ *be a finite set of propositional variables. The syntax of probabilistic alternating-time μ-calculus (PAMC for short) is defined as follows:*

$$\phi ::= p \mid \neg\phi \mid Z \mid \phi\wedge\phi \mid \phi\vee\phi \mid \mu Z.\phi \mid \nu Z.\phi \mid \langle\langle A\rangle\rangle^{\bowtie k}\phi \mid [[A]]^{\bowtie k}\phi$$

where $p \in \mathsf{AP}$*,* $Z \in \mathcal{Z}$*,* $\bowtie \in \{\geq, >, <, \leq\}$*,* $k \in [0, 1]$ *is a rational constant,* $A \subseteq \mathsf{Ag}$*, and for each* $\mu Z.\phi$ *and* $\nu Z.\phi$*, each occurrence of* Z *in* ϕ *is under the scope of an even number of negations in* ϕ*.*

Let ϕ be a PAMC formula. $Z \in \mathcal{Z}$ is a *free variable* in ϕ if an occurrence of Z is not in the scope of a fixed-point operator

μZ or νZ. ϕ is *closed* if ϕ does not contain any free variables. A closed formula is called a *sentence*. W.l.o.g., we assume hereafter that each variable Z is quantified by either μ or ν at most once in each PAMC sentence. We denote by $\bot \equiv p \wedge \neg p$ and $\top \equiv p \vee \neg p$. For simplifying the presentation, formulae like $\langle\langle\{i_1, ..., i_m\}\rangle\rangle^{\bowtie k}\phi$ may be written as $\langle\langle i_1, ..., i_m\rangle\rangle^{\bowtie k}\phi$.

Given two PAMC sentences $\eta Z.\phi$ for $\eta \in \{\mu, \nu\}$ and φ, let $\phi[\varphi/Z]$ be the sentence obtained from ϕ by replacing every occurrence of Z with φ. The Fisher-Ladner closure $\mathsf{FL}(\phi)$ of a PAMC sentence ϕ contains all the sub-sentences of ϕ with the rule: if $\eta Z.\varphi \in \mathsf{FL}(\phi)$ for $\eta \in \{\mu, \nu\}$, then $\varphi[\eta Z.\varphi/Z] \in \mathsf{FL}(\phi)$. The size of $\mathsf{FL}(\phi)$ is at most double that of ϕ. We denote by $\mathsf{FL}_\exists(\phi) \subseteq \mathsf{FL}(\phi)$ and $\mathsf{FL}_\forall(\phi) \subseteq \mathsf{FL}(\phi)$ the set of sentences of the form $\langle\langle A\rangle\rangle^{\bowtie k}\varphi$ and $[[A]]^{\bowtie k}\varphi$, respectively.

Example 1. *Let* $\phi = \nu Z_1.(\langle\langle 1, 2\rangle\rangle^{\geq \frac{1}{2}} Z_1 \wedge \phi')$ *where* $\phi' = \mu Z_2.(p \vee \langle\langle 1, 3\rangle\rangle^{\geq \frac{1}{3}} Z_2)$*, we have* $\mathsf{FL}(\phi) = \{\phi, \langle\langle 1, 2\rangle\rangle^{\geq \frac{1}{2}}\phi \wedge \phi', \langle\langle 1, 2\rangle\rangle^{\geq \frac{1}{2}}\phi, \phi', p \vee \langle\langle 1, 3\rangle\rangle^{\geq \frac{1}{3}}\phi', p, \langle\langle 1, 3\rangle\rangle^{\geq \frac{1}{3}}\phi'\}$.

The semantics of PAMC is defined w.r.t. PCGSs and a valuation $\xi : \mathcal{Z} \to 2^Q$. Let $\xi[S/Z]$ denote the valuation such that $\xi[S/Z](Z) = S$ and $\xi[S/Z](Z') = \xi(Z')$ for $Z \neq Z'$.

Definition 3. *Given a PCGS* $\mathcal{M} = (\mathsf{Ag}, \mathsf{Act}, Q, \Gamma, \delta, \lambda, q^0)$*, the semantics of PAMC is defined via the denotation function* $\llbracket \circ \rrbracket^{\xi}_{\mathcal{M}}$*, which is defined as follows:*

- $\llbracket p \rrbracket^{\xi}_{\mathcal{M}} = \{q \in Q \mid p \in \lambda(q)\}$;
- *Boolean connectors are defined in a standard way;*
- $\llbracket \mu Z.\phi \rrbracket^{\xi}_{\mathcal{M}} = \bigcap\{Q' \subseteq Q \mid \llbracket \phi \rrbracket^{\xi[Q'/Z]}_{\mathcal{M}} \subseteq Q'\}$;
- $\llbracket \nu Z.\phi \rrbracket^{\xi}_{\mathcal{M}} = \bigcup\{Q' \subseteq Q \mid \llbracket \phi \rrbracket^{\xi[Q'/Z]}_{\mathcal{M}} \supseteq Q'\}$;
- $\llbracket \langle\langle A\rangle\rangle^{\bowtie k}\phi \rrbracket^{\xi}_{\mathcal{M}} = \left\{ \begin{array}{l} q \in Q \mid \exists \upsilon_A \in \Upsilon_A,\ \forall \upsilon_{\overline{A}} \in \Upsilon_{\overline{A}} : \\ \mathsf{Pr}_q^{\upsilon_A, \upsilon_{\overline{A}}}(\{\pi \in O_q^{\upsilon_A, \upsilon_{\overline{A}}} \mid \pi_1 \in \llbracket \phi \rrbracket^{\xi}_{\mathcal{M}}\}) \bowtie k \end{array} \right\}$;
- $\llbracket [[A]]^{\bowtie k}\phi \rrbracket^{\xi}_{\mathcal{M}} = \left\{ \begin{array}{l} q \in Q \mid \forall \upsilon_A \in \Upsilon_A,\ \exists \upsilon_{\overline{A}} \in \Upsilon_{\overline{A}} : \\ \mathsf{Pr}_q^{\upsilon_A, \upsilon_{\overline{A}}}(\{\pi \in O_q^{\upsilon_A, \upsilon_{\overline{A}}} \mid \pi_1 \in \llbracket \phi \rrbracket^{\xi}_{\mathcal{M}}\}) \bowtie k \end{array} \right\}$.

We sometimes drop the superscript ξ from $\llbracket \phi \rrbracket^{\xi}_{\mathcal{M}}$ if ϕ is a PAMC sentence. When it is clear from context, the subscript $\mathcal{M}$ is also dropped from $\llbracket \phi \rrbracket^{\xi}_{\mathcal{M}}$ and $\llbracket \phi \rrbracket_{\mathcal{M}}$.

AMC (resp. PμTL (Liu et al. 2015)) is an extension of μ-calculus in which the next-modalities are replaced by $\langle\langle A\rangle\rangle\phi$ and $[[A]]\phi$ (resp. $[\mathbf{X}\phi]^{\bowtie k}$) and the semantics is defined over CGSs (resp. MCs). PATL/PATL* (Chen and Lu 2007) are probabilistic variants of ATL/ATL* (Alur, Henzinger, and Kupferman 2002) and the semantics is defined over PCGSs.

Theorem 1. *PAMC subsumes AMC and PμTL. PAMC and PATL/PATL* are incomparable.*

Proof sketch. Each AMC (resp. PμTL) formula can be encoded as a PAMC formula, which can be shown easily by structural induction. For instance, $\langle\langle A\rangle\rangle\phi$ in AMC (resp. $[\mathbf{X}\phi]^{\bowtie k}$ in PμTL) can be encoded as $\langle\langle A\rangle\rangle^{\geq 1}\phi'$ (resp. $\langle\langle\emptyset\rangle\rangle^{\bowtie k}\phi'$), where ϕ' denotes the encoding of ϕ in respective cases. Some PAMC formulae (e.g., $\langle\langle A\rangle\rangle^{>0.5}\phi$) cannot be expressed in AMC or PμTL.

Since PμTL and PCTL are incomparable on Markov chains (Liu et al. 2015), there exists a PCTL formula (and thus a PATL formula) which is inexpressible in PAMC. Moreover, we note that AMC is more expressive than ATL*, by which a formula in PAMC can be constructed which is inexpressible in PATL*. □

Proposition 1 is directly from the PAMC semantics. To simplify the notation, we write $\overline{>}$ for $<$, $\overline{<}$ for $>$, $\overline{\leq}$ for $\geq$ and $\overline{\geq}$ for $\leq$, and write $\widehat{\geq}$ for $<$, $\widehat{>}$ for $\leq$, $\widehat{\leq}$ for $>$ and $\widehat{<}$ for $\geq$.

Proposition 1. *Given a PAMC formula* $\langle\langle A\rangle\rangle^{\bowtie k}\phi$ *or* $[[A]]^{\bowtie k}\phi$ *for* $k \in [0, 1]$, $\bowtie \in \{\geq, >, \leq, <\}$*, a PCGS* $\mathcal{M}$ *and a valuation* ξ*, we have the following deduction rules.*

1. $\llbracket \langle\langle A\rangle\rangle^{\bowtie k}\phi \rrbracket^{\xi} = \llbracket \langle\langle A\rangle\rangle^{\overline{\bowtie} 1-k} \neg\phi \rrbracket^{\xi}$.
2. $\llbracket [[A]]^{\bowtie k}\phi \rrbracket^{\xi} = \llbracket [[A]]^{\overline{\bowtie} 1-k} \neg\phi \rrbracket^{\xi}$.
3. $\llbracket \neg\langle\langle A\rangle\rangle^{\bowtie k}\phi \rrbracket^{\xi} = \llbracket [[A]]^{\widehat{\bowtie} 1-k}\phi \rrbracket^{\xi}$.
4. $\llbracket \neg[[A]]^{\bowtie k}\phi \rrbracket^{\xi} = \llbracket \langle\langle A\rangle\rangle^{\widehat{\bowtie} 1-k}\phi \rrbracket^{\xi}$.
5. $\llbracket \langle\langle \mathsf{Ag}\rangle\rangle^{\bowtie k}\phi \rrbracket^{\xi} = \llbracket [[\emptyset]]^{\bowtie k}\ \phi \rrbracket^{\xi}$.
6. $\llbracket [[\mathsf{Ag}]]^{\bowtie k}\phi \rrbracket^{\xi} = \llbracket \langle\langle\emptyset\rangle\rangle^{\bowtie k}\ \phi \rrbracket^{\xi}$.

A PAMC sentence is in *negation normal form* (NNF) if $\neg$ only appear in front of atomic propositions. Using Proposition 1 and the rule $\neg\mu Z.\phi \equiv \nu Z.\neg\phi[\neg Z/Z]$, every PAMC sentence can be equivalently transformed into NNF. Hereafter, we assume that all PAMC sentences are in NNF.

In this work, for a given PAMC sentence ϕ, we consider the following two fundamental problems. **Model checking** is to determine whether $q^0 \in \llbracket \phi \rrbracket_{\mathcal{M}}$ for the initial state q^0 of a given PCGS $\mathcal{M}$. **Satisfiability checking** is to determine whether there exists a PCGS $\mathcal{M}$ with the initial state q^0 such that $q^0 \in \llbracket \phi \rrbracket_{\mathcal{M}}$.

In PAMC, only probabilistic next-time coalition modalities $\langle\langle A\rangle\rangle^{\bowtie k}\phi$ and $[[A]]^{\bowtie k}\phi$ are allowed, hence it suffices to consider memoryless strategies for the coalition A. By leveraging the model checking algorithm for AMC (Alur, Henzinger, and Kupferman 2002) that computes $\llbracket \phi \rrbracket$ recursively, we can get a model checking algorithm for PAMC. The key difference lies in the handling of $\langle\langle A\rangle\rangle^{\bowtie k}\phi$ and $[[A]]^{\bowtie k}\phi$, which can be done in polynomial time by linear programming. The model checking problem for AMC lies in $\mathsf{UP}\cap$co-UP, and can be solved in exponential time in the alternation depth of the formula. We have that

Theorem 2. *The model checking problem for PAMC is in* $\mathsf{UP}\cap$*co-*UP *and can be decided in* $\mathbf{O}((|\phi| \cdot |\mathcal{M}|)^{c \cdot d})$ *time for some constant c, where d denotes the alternation depth of* ϕ.

An application: Genetic Regulatory Networks

To demonstrate the usage of PCGSs and PAMC, we present an example from precision medicine based on the probabilistic Boolean network (PBN) model for genetic regulatory networks (Shmulevich and Dougherty 2010). Consider a PBN with n genes, each of which has two local states $\{0, 1\}$ where 0 and 1 indicate that the corresponding gene is not expressed and expressed, respectively. A (global) state of the PBN is a Boolean vector of local states with width n. Each gene i has a finite set of predictor functions $F_i = \{f^i_1, ..., f^i_{k_i}\}$ denoting the inter-gene relationship, where each function f^i_j determines the local state of gene i at the next step when the j-th element of F_i (i.e., f^i_j) is chosen, and c^i_j is the probability that f^i_j is selected (so $\sum_{j=1}^{k_i} c^i_j = 1$). The probability of the PBN evolving from a state $\vec{q} = [q_1, ..., q_n]$ to the state $\vec{q}' = [q'_1, ..., q'_n]$ is

$\sum_{f_1\in F_1:f_1(\vec{q})=q'_1}\cdots\sum_{f_n\in F_n:f_n(\vec{q})=q'_n}\prod_{i=1}^{n} c^i$, where c^i is the probability that f_i is selected by the gene i.

For intervention purposes, control inputs are introduced into PBN to reason about treatment strategies. For instance, in cancer therapy, control inputs (e.g. radiation, chemotherapy) may be employed to move the state probability distribution vector away from one associated with uncontrolled cell proliferation or markedly reduced apoptosis (Shmulevich and Dougherty 2010, Section 5.2). Suppose there are $x_1, ..., x_m$ control inputs ranging over binary values $\{0, 1\}$, which control the probabilities of predictor functions. The control input being 1 indicates that a corresponding predictor function (intuitively a treatment) is applied; 0 otherwise. Under the values $\vec{v} = v_1, ..., v_m$ of control inputs, the probability of the predictor function $f^i_j \in F_i$ becomes $c^i_{j,\vec{v}}$ with $\sum_{j=1}^{k_i} c^i_{j,\vec{v}} = 1$. (The specific definitions of $c^i_{j,\vec{v}}$ are irrelevant here.) The main problem is to synthesize values for the control inputs at each state to ensure that the network behaves as desired.

The control of PBN can be naturally modelled as a PCGS $\mathcal{M} = (\mathsf{Ag}, \mathsf{Act}, Q, \Gamma, \delta, \lambda, q^0)$, where $\mathsf{Ag} = [n+m]$ with $i \in [n]$ denotes the gene i, and $n + i \in \{n + 1, ...n + m\}$ denotes the control input i; $Q = \{0, 1\}^n$ corresponding to states of the network; $\mathsf{Act}_i = F_i$ for $i \in [n]$ denoting the predictor functions and $\mathsf{Act}_{n+i} = \{0, 1\}$ for $n + i \in \{n + 1, ...n + m\}$ denoting treatment choices; Γ_i is the function such that $\Gamma_i(\vec{q}) = \mathsf{Act}_i$ for every agent i and state $\vec{q}$; q^0 is the initial state of the network which is determined by the patient's physiology; δ is the probability transition function such that for every $\vec{q} \in Q$, $f^i_{j_i} \in F_i$ and $\vec{v} = v_1, ..., v_m \in \{0, 1\}^m$:

$$\delta(\vec{q}, \langle f^1_{j_1}, ..., f^n_{j_n}, v_1, ..., v_m\rangle, [f^1_{j_1}(\vec{q}), ..., f^n_{j_n}(\vec{q})]) := \prod_{i=1}^{n} c^i_{j_i,\vec{v}}.$$

Note that we do not specify λ explicitly as this is usually application-specific.

With PCGS defined above, we can formulate the *achievement property* as: $\mu Z.(p \vee \langle\langle A\rangle\rangle^{\geq 0.8} Z)$ for some $A \subseteq \{n + 1, ..., n + m\}$. This formula expresses that some desired-region is reachable, and there is treatment strategy using at most treatments in A at each step such that it has at least probability 0.8 to go on with the right direction.

The *maintenance property* can be formulated as $\nu Z.(p \wedge \langle\langle A\rangle\rangle^{>0.9} Z)$ stating that there is a treatment strategy using at most treatments in A at each step such that the network has less than 0.1 probability to escape from the desired-region.

Deciding Satisfiability

In this section, we present a reduction from the satisfiability problem of a PAMC sentence ϕ to solving a (turned-based, two-player) parity game $\mathcal{G}_\phi$ such that ϕ is satisfiable iff Player-0 has a winning strategy in $\mathcal{G}_\phi$. Intuitively, as in the classic decision procedure of μ-calculus, the fixpoint is handled by parity games. In particular, the two players are model "constructor" and "spoiler". The constructor intends to show that there a model witnessing the satisfiability of ϕ, while the spoiler tries to defeat the constructor by demonstrating a path in the game to show that a model cannot exist. Each state controlled by the constructor consists of a set of sub-sentences of ϕ for which the constructor strives to show that it is satisfiable. In addition, the coalition strategies for the set of sub-sentences are tackled by the notions of intersection graph and maximally independent set (cf. Definition 5), which are used to make sure that the strategies of an agent in different coalitions are consistent. To handle probabilities, weighted covers are used which would guarantee that there exist distributions over covers satisfying constraints $\bowtie k$. To illustrate the reduction, the PAMC sentence $\Psi = \langle\langle 1, 2\rangle\rangle^{>0.5} p_1 \wedge \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2) \wedge [[1, 3]]^{>0.4}(\neg p_1 \wedge \neg p_2)$ will be used as an example. We start with some concepts.

Parity automata. A *parity automaton* (PA) is a tuple $\mathcal{P} = (S, \Sigma, \delta, S^0, F)$, where S is a finite set of states, Σ is the input alphabet, $\delta : S \times \Sigma \to 2^S$ is a transition function, $S^0 \subseteq S$ is the set of initial states and $F : S \to \{0, ..., k\}$ is a *rank* function. A *run* ρ of $\mathcal{P}$ over an ω-word $\alpha_0\alpha_1... \in \Sigma^\omega$ is an infinite sequence of states $s_0 s_1 \cdots$ such that $s_0 \in S^0$, and for every $i \geq 0$, $s_{i+1} \in \delta(s_i, \alpha_i)$. Let $\mathsf{inf}(\rho)$ be the set of states visited infinitely often in ρ. ρ is *accepting* iff $\min_{s\in\mathsf{inf}(\rho)} F(s)$ is *even*. An infinite word is *accepted* by $\mathcal{P}$ iff $\mathcal{P}$ has an accepting run over this word. We denote by $L(\mathcal{P}) \subseteq \Sigma^\omega$ the set of all infinite words accepted by $\mathcal{P}$. A *Büchi automaton* (BA) is a special PA $(S, \Sigma, \delta, S^0, F)$ in which $F : S \to \{0, 1\}$. A PA $\mathcal{P} = (S, \Sigma, \delta, S^0, F)$ is *deterministic* if $|S^0| = 1$ and for all $(s, \alpha) \in S \times \Sigma$, $|\delta(s, \alpha)| \leq 1$. δ and S^0 in a deterministic parity automaton (DPA for short) are simplified as the function $\delta : S \times \Sigma \to S$ and a state s^0.

Turned-based two-player parity games. A *(turned-based two-player) parity game* $\mathcal{G}$ is a tuple $(V = V_0 \uplus V_1, E, v_{init}, \Xi)$, where V_0 is a finite set of states (i.e., vertices) of Player-0, V_1 is a finite set of states (i.e., vertices) of Player-1, $E \subseteq V \times V$ is a finite the set of edges, $v_{init} \in V_0$ is the initial state, and $\Xi : V \to \{0, ..., k\}$ is a *rank* function. An *infinite play* ρ of $\mathcal{G}$ is an infinite sequence of states $v_0 v_1 ...$ such that $v_0 = v_{init}$, and for every $i \geq 0$, $(v_i, v_{i+1}) \in E$. A *finite play* ρ of $\mathcal{G}$ is a finite sequence of states $v_0 v_1 ... v_n$ such that $v_0 = v_{init}$, and for every $i \in [n]$, $(v_{i-1}, v_i) \in E$. A *strategy* of Player-i is a partial function $\theta : V^* V_i \to V$ such that for every $\rho \in V^*$ and $v \in V_i$, if $\theta(\rho \cdot v)$ is defined, then $(v, \theta(\rho \cdot v)) \in E$. Given a strategy θ_0 for Player-0 and a strategy θ_1 for Player-1, let $\mathcal{G}_{\theta_0,\theta_1}$ be the play such that Player-0 and Player-1 enforce their strategies θ_0 and θ_1 during the play. Player-0 *wins* on a finite play ρ iff the last state of ρ is controlled by Player-1 and it cannot move to the next state anymore. Player-0 *wins* on an infinite play ρ iff $\min_{s\in\mathsf{inf}(\rho)} F(s)$ is *even*. Player-0 has a winning strategy iff Player-0 has a strategy θ_0 such that Player-0 wins on the play $\mathcal{G}_{\theta_0,\theta_1}$, for each strategy θ_1 Player-1 chooses.

Covers and weighted covers. Given a set of sentences Φ, a *cover* c of Φ is a set $c \subseteq 2^\Phi$ such that $\bigcup_{v\in c} v = \Phi$. A *weighted cover* of Φ is a pair (c, w) such that c is a cover of Φ and $w : c \to (0, 1]$ is a probability function such that $\sum_{v\in c} w(v) = 1$. The *width* of a cover c is the cardinality of c. Given a weighted cover (c, w) of $\Phi = \{\phi_1, \cdots, \phi_m\}$, let $\Phi(\phi_i) := \{v \in c \mid \phi_i \in v\}$ and $w(\phi_i) := \sum_{v\in\Phi(\phi_i)} w(v)$. We write $C^b(\Phi)$ for the set of covers of Φ with width at most b. Given b and a set Φ, $|C^b(\Phi)|$ is at most $\frac{2^{|\Phi|\cdot(b+1)} - 2^{|\Phi|}}{2^{|\Phi|} - 1}$ (Chakraborty and Katoen 2016). Intuitively, a weight cover (c, w) of Φ can be seen as a distribution w over the cover c.

Definition 4. *A set v of sentences is* transitive *if the following conditions hold: (1) If $\phi_1 \wedge \phi_2 \in v$, then $\phi_1, \phi_2 \in v$; (2) If*

$\phi_1 \vee \phi_2 \in v$, *then* $\phi_1 \in v$ *or* $\phi_2 \in v$*; (3) If* $\eta Z.\varphi \in v$ *for* $\eta \in \{\mu, \nu\}$*, then* $\varphi[\eta Z.\varphi/Z] \in v$*; and (4) There exists at least one sentence of the form* $\langle\langle A\rangle\rangle^{\bowtie k}\varphi$ *or* $[[A]]^{\bowtie k}\varphi$ *in* v*.*

Let V_t denote the set of transitive sets of sentences. For every $v \in V_t$, we denote by $\langle\langle v\rangle\rangle$ and $[[v]]$ the sets $\{\langle\langle A\rangle\rangle^{\bowtie k}\varphi \in v\}$ and $\{[[A]]^{\bowtie k}\varphi \in v\}$, respectively. Intuitively, consider a set v that is controlled by the constructor in the parity game $\mathcal{G}_\phi$. If v is transitive, then the constructor should give successor states to satisfy sentences in $\langle\langle v\rangle\rangle$ and $[[v]]$.

Example 2. *The set* $v = \{\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2), \langle\langle 1,2\rangle\rangle^{>0.5}p_1, [[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)\}$ *is a transitive vertex. We have:* $\langle\langle v\rangle\rangle = \{\langle\langle 1,2\rangle\rangle^{>0.5}p_1, \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)\}$*, and* $[[v]] = \{[[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)\}$*.*

Definition 5. *An* intersection graph $G_v = (\langle\langle v\rangle\rangle, E_v)$ *formed from* $v \in V_t$ *is an undirected graph such that* $(\langle\langle A_1\rangle\rangle^{\bowtie_1 k_1}\varphi_1, \langle\langle A_2\rangle\rangle^{\bowtie_2 k_2}\varphi_2) \in E_v$ *iff* $A_1 \cap A_2 \neq \emptyset$*.*

A maximal independent set *(MIS)* $\mathfrak{M}$ *of* G_v *is a maximal set of vertices in* G_v *such that there are no edges between each pairs of vertices from* $\mathfrak{M}$*.*

Given a set $B \subseteq \mathsf{Ag}$*, a* B-dominant MIS *(B-MIS)* $\mathfrak{M}$ *of* G_v *is a maximal set of vertices in* G_v *such that (1) there is no edge between each pairs of vertices from* $\mathfrak{M}$ *and (2) for all* $\langle\langle A\rangle\rangle^{\bowtie k}\varphi \in \mathfrak{M}$*,* $A \subseteq B$*.*

Let $\mathtt{mis}_v$ (resp. $\mathtt{mis}_v^B$) denote the set of MIS (resp. B-MIS) in G_v. Note that $\mathtt{mis}_v$ and $\mathtt{mis}_v^B$ are sets whose members are sets of sub-sentences. Let $\mathtt{MIS}_v := \Big(\mathtt{mis}_v \setminus \bigcup_{[[B]]^{\bowtie k}\varphi \in [[v]]} \mathtt{mis}_v^B\Big) \cup \Big\{\mathfrak{M} \cup \{[[B]]^{\bowtie k}\varphi\} \mid \mathfrak{M} \in \mathtt{mis}_v^B, [[B]]^{\bowtie k}\varphi \in [[v]]\Big\}$. As mentioned before, MIS and B-dominant MIS are used to ensure the consistency of the strategies of an agent in different coalitions. Intuitively, consider a transitive set v controlled by the constructor. To satisfy sentences in $\langle\langle v\rangle\rangle$ and $[[v]]$, probabilistic constraints in each set $\mathfrak{M}$ of $\mathtt{MIS}_v$ should be ensured simultaneously by one distribution, but this distribution need not satisfy the probabilistic constraints outside of $\mathfrak{M}$ by choosing proper actions.

Example 3. *Consider the vertex* v *defined in Example 2, we have:* $\mathtt{mis}_v = \Big\{\{\langle\langle 1,2\rangle\rangle^{>0.5}p_1, \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)\}\Big\}$, $\mathtt{mis}_v^{\{1,3\}} = \Big\{\{\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)\}\Big\}$, $\mathtt{MIS}_v = \Big\{\{\langle\langle 1,2\rangle\rangle^{>0.5}p_1, \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)\}, \{\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2), [[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)\}\Big\}$.

Proposition 2. *(Moon and Moser 1965) Given a set* v *of sentences, the size of* $\mathtt{mis}_v$ *(resp.* $\mathtt{mis}_v^B$*) is at most* $3^{\frac{|v|}{3}}$ *and can be computed in time* $\mathbf{O}(3^{\frac{|v|}{3}})$*.*

Definition 6. *Given a set* $\mathfrak{M} \in \mathtt{MIS}_v$*, the* objective *of* $\mathfrak{M}$ *is the set* $O_\mathfrak{M} := \{\varphi \mid \langle\langle A\rangle\rangle^{\bowtie k}\varphi \in \mathfrak{M}$ *or* $[[A]]^{\bowtie k}\varphi \in \mathfrak{M}\}$*.*

A weighted cover $(c_\mathfrak{M}, w_\mathfrak{M})$ *of* $O_\mathfrak{M}$ *such that* $c_\mathfrak{M} \in C^b(O_\mathfrak{M})$ *satisfies* $\mathfrak{M}$*, denoted by* $(c_\mathfrak{M}, w_\mathfrak{M}) \models \mathfrak{M}$*, if for all* $\langle\langle A\rangle\rangle^{\bowtie k}\varphi, [[A]]^{\bowtie k}\varphi \in \mathfrak{M}$*,* $w_\mathfrak{M}(\varphi) \bowtie k$*, where* $b = |\mathfrak{M}| + 1$*.*

Example 4. *Consider the vertex* v *from Example 2. We have* $\mathtt{MIS}_v = \{\mathfrak{M}, \mathfrak{M}'\}$*, where* $\mathfrak{M} = \{\langle\langle 1,2\rangle\rangle^{>0.5}p_1, \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)\}$ *and* $\mathfrak{M}' = \{\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2), [[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)\}$*. Furthermore,* $O_\mathfrak{M} = \{p_1, \neg p_1 \vee p_2\}$*, and* $O_{\mathfrak{M}'} = \{\neg p_1 \vee p_2, \neg p_1 \wedge \neg p_2\}$*.* $O_\mathfrak{M}$ *has one cover* $c_1 = \{\{p_1, \neg p_1 \vee p_2\}\}$ *such that* $(c_1, w_1) \models \mathfrak{M}$ *for some weight* w_1*, e.g.,* $w_1(\{p_1, \neg p_1 \vee p_2\}) = 1$*.* $\mathfrak{M}'$ *have two covers* $c_2 = \{\{\neg p_1 \vee p_2, \neg p_1 \wedge \neg p_2\}\}$ *and* $c_2' = \{\{\neg p_1 \vee p_2\}, \{\neg p_1 \wedge \neg p_2\}\}$ *such that* $(c_2, w_2) \models \mathfrak{M}'$ *and* $(c_2', w_2') \models \mathfrak{M}'$ *for some weights* w_2 *and* w_2'*, e.g.* $w_2(\{\neg p_1 \vee p_2, \neg p_1 \wedge \neg p_2\} = 1)$*,* $w_2'(\{\neg p_1 \vee p_2\}) = 0.55$ *and* $w_2'(\{\neg p_1 \wedge \neg p_2\}) = 0.45)$*.*

Definition 7. *Given a PAMC sentence* ϕ*, the two-player game* G_ϕ *is a triple* $(V = V_0 \uplus V_1, E, v_{\text{init}})$*, where* V_0 *and* V_1 *form a partition of* V*,* V_i *for* $i \in \{0,1\}$ *is the finite set of vertices for Player-i,* $v_{\text{init}} = \{\phi\} \in V_0$ *is the initial vertex of the game and* $E \subseteq V \times V$ *is a finite set of edges.*

V_0 *and* V_1 *are defined below.* $E \subseteq V \times V$ *is the least set of edges satisfying the following conditions.*

1. $(v, v \cup \{\phi_i\}) \in E$ *if* $\phi_1 \vee \phi_2 \in v$ *and* $\phi_i \notin v$*, for* $i \in \{1,2\}$*.*
2. $(v, v \cup \{\varphi[\eta Z.\varphi/Z]\}) \in E$ *if* $\eta Z.\varphi \in v$ *and* $\varphi[\eta Z.\varphi/Z] \notin v$*, for* $\eta \in \{\mu, \nu\}$*.*
3. $(v, v \cup \{\phi_1, \phi_2\}) \in E$ *if* $\phi_1 \wedge \phi_2 \in v$ *and* $\{\phi_1, \phi_2\} \nsubseteq v$*.*
4. $(v, \perp) \in E$ *if both* $p \in v$ *and* $\neg p \in v$ *for some* $p \in \mathsf{AP}$*.*
5. $(v, \top) \in E$ *if* $v \notin V_t$ *and Items 1-4 cannot be applied to* v*.*
6. *For each* $v \in V_t$*, let* $\mathtt{MIS}_v = \{\mathfrak{M}_v^1, \cdots, \mathfrak{M}_v^{k_v}\}$*,* $C_j := \{c_\mathfrak{M} \mid (c_\mathfrak{M}, w_\mathfrak{M}) \models \mathfrak{M}_v^j\}$*, and* $C = \bigcup_{j \in [k_v]} C_j$*. If* $C_j = \emptyset$ *for some* $j \in [k_v]$*, then* $(v, \perp) \in E$*; otherwise* $(v, C) \in E$*,* $(C, C_j) \in E$ *for* $j \in [k_v]$*,* $(C_j, c) \in E$ *for each* $c \in C_j$*, and* $(c, v') \in E$ *for each* $v' \in c$*.*

As a result, $V_0 := 2^{\mathsf{FL}(\phi)} \cup \{\perp\} \cup \{C_j \mid (C, C_j) \in E\}$ *and* $V_1 := V_{1,0} \cup V_{1,1} \cup \{\top\}$*, where* $V_{1,0} = \{C \mid (v, C) \in E\}$*,* $V_{1,1} = \{c \mid (c, v') \in E\}$*,* C, C_j *and* c *are defined as above.*

The intuition of Items 1-5 are straightforward. To see the intuition behind Item 6. Let us begin with $v = \{\langle\langle A_1\rangle\rangle^{\bowtie_1 k_1}\phi_1, \cdots, \langle\langle A_n\rangle\rangle^{\bowtie_n k_n}\phi_n\}$. If $A_i \cap A_j = \emptyset$ for some $i, j \in [n]$, then for each pair of possible strategies of A_i and A_j, there always exist some common decisions when A_i and A_j enforce their strategies, hence their probabilistic constraints should be *simultaneously* ensured by one distribution. Suppose $\mathtt{mis}_v = \{\mathfrak{M}_1, \cdots, \mathfrak{M}_m\}$ (cf. Definition 5), then, (1) for each pair of $\langle\langle A\rangle\rangle^{\bowtie k}\phi$ and $\langle\langle A'\rangle\rangle^{\bowtie' k'}\phi'$ in $\mathfrak{M}_i$, $A \cap A' = \emptyset$, (2) for each sentence $\langle\langle A\rangle\rangle^{\bowtie k}\phi$ from $\mathfrak{M}_i$, for every set $\mathfrak{M}_j \neq \mathfrak{M}_i$, there exists a sentence $\langle\langle A'\rangle\rangle^{\bowtie' k'}\phi'$ in $\mathfrak{M}_j$ with $A \cap A' \neq \emptyset$, and (3) each sentence in v must occur at least in one set $\mathfrak{M}_i$. This allows us to consider each set $\mathfrak{M}_i$ *individually* by assigning the action a_j to players in A_j. For each set $\mathfrak{M}_j$, we compute its weighted covers, which check existence of distributions that satisfying the probabilistic constraints in $\mathfrak{M}_j$.

For the general case $v = \{\langle\langle A_1\rangle\rangle^{\bowtie_1 k_1}\varphi_1, \cdots, \langle\langle A_m\rangle\rangle^{\bowtie_m k_m}\varphi_m, [[B_1]]^{\sim_1 h_1}\psi_1, \cdots, [[B_n]]^{\sim_n h_n}\psi_n\}$, we assume there exists a player $j \notin \bigcup_{i=1}^n B_i$ (cf. Remark 1). By assigning the action b_i to the player j for each sentence $[[B_i]]^{\sim_i h_i}\psi_i$, it suffices to check *individually* each of $\{[[B_1]]^{\sim_1 h_1}\psi_1, \cdots, [[B_n]]^{\sim_n h_n}\psi_n\}$. For each B_i, let $\mathtt{mis}_v^{B_i} = \{\mathfrak{M}_1^i, \cdots, \mathfrak{M}_{m_i}^i\}$, then (1) for each sentence $\langle\langle A\rangle\rangle^{\bowtie k}\varphi$ from $\mathtt{mis}_v^{B_i}$, $A \subseteq B_i$, and (2) $\mathfrak{M}_\ell^i$ for $1 \leq \ell \leq m_i$ must be a subset of a set in $\mathtt{mis}_v$. By assigning proper actions to players (cf. previous paragraph), for each $[[B_i]]^{\sim_i h_i}\psi_i$, it suffices to check existence of a distribution for sub-sentences in $\mathfrak{M}_\ell^i$ and $[[B_i]]^{\sim_i h_i}\psi_i$ *simultaneously*, which is done by computing weighted covers. Moveover, each MIS $\mathfrak{M} \in \mathtt{mis}_v$ such that $\mathfrak{M} \subseteq \mathfrak{M}'$ for some $\mathfrak{M}' \in \mathtt{mis}_v^{B_i}$ is omitted to avoid double checking.

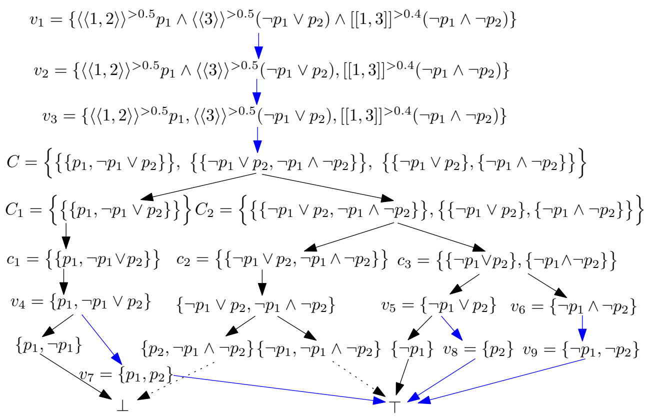

Figure 1: The game G_Ψ for Ψ, where $C, \top, c_1, c_2$, and c_3 are Player-1 vertices, others are Player-0 vertices.

Overall, for each $v \in V_t$ with $\mathtt{MIS}_v = \{\mathfrak{M}_v^1, \cdots, \mathfrak{M}_v^{k_v}\}$, if $C_j = \emptyset$ for some $j \in [k_v]$, then no distribution satisfies the probabilistic constraints of $\mathfrak{M}_v^j$, we added $(v, \perp)$ into E. i.e., Player-0 loses the game. Otherwise, the play goes to C and let Player-1 to choose one $\mathfrak{M}_v^j$ to dissatisfy, i.e., $(C, C_j) \in E$. At C_j, Player-0 chooses one cover $c \in C_j$ such that $(c, w) \models \mathfrak{M}_v^j$, i.e., the distribution w satisfies the probabilistic constraints of $\mathfrak{M}_v^j$. Next, Player-1 chooses one set of sub-sentences $v' \in c$ to dissatisfy.

Example 5. *Recalling the sentence* $\Psi = \langle\langle 1,2\rangle\rangle^{>0.5}p_1 \wedge \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2) \wedge [[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)$*, the corresponding game is shown in Figure 1, in which* $c_1, c_2, c_3, \top$ *and* C *are Player-1 vertices, others are Player-0 vertices.* $\langle\langle 1,2\rangle\rangle^{>0.5}p_1$ *and* $\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)$ *have some common decisions, as well as* $\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)$ *and* $[[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)$*. While,* $\langle\langle 1,2\rangle\rangle^{>0.5}p_1$ *and* $[[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)$ *can have disjoint decisions by assigning proper actions to Player-1. We have:* $\mathtt{MIS}_v = \{\mathfrak{M}, \mathfrak{M}'\}$*, where* $\mathfrak{M} = \{\langle\langle 1,2\rangle\rangle^{>0.5}p_1, \langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2)\}$ *and* $\mathfrak{M}' = \{\langle\langle 3\rangle\rangle^{>0.5}(\neg p_1 \vee p_2), [[1,3]]^{>0.4}(\neg p_1 \wedge \neg p_2)\}$*. If Player-1 chooses* C_1 *(resp.* C_2*) to dissatisfy, then Player-0 can choose* c_1 *(resp.* c_3*) as the witness of* $\mathfrak{M}$ *(resp.* $\mathfrak{M}'$*). It is easy to verify that Player-0 has a winning strategy in the game* G_Ψ *by choosing* c_1, c_3, v_7, v_8 *and* v_9*. The model construction from a winning strategy is given in Example 6.*

Finally, in order to prevent from regeneration sequences for each μ-sentence ψ (i.e., derivation sequences that derive ψ infinitely often), we construct a DPA which accepts all terminating regeneration sequences. By constructing the cross-product of the two-player game G_ϕ and the DPA, we get the resulting parity game $\mathcal{G}_\phi$ such that ϕ is satisfiable iff Player-0 has a winning strategy in the parity game $\mathcal{G}_\phi$.

Definition 8. *Given a PAMC sentence* ϕ*, we define a Büchi automaton* $\mathcal{P}_\mu = (S_\mu, \Sigma, \delta_\mu, S_\mu^0, F_\mu)$ *where* $S_\mu = S_\mu^0 = \{\psi \in \mathsf{FL}(\phi) \mid \psi$ *contains some* μ*-sentence*$\}$, $\Sigma = 2^{\mathsf{FL}(\phi)}$, $\delta_\mu(q, \alpha) = q'$*, if* q' *is derived from* q *and* $q' \in \alpha$*, and* $F_\mu(s) = 0$ *for all* $s \in S_\mu$*.*

$\mathcal{P}_\mu$ precisely accepts all the sequences in which some μ-sub-sentence of ϕ is derived infinitely often. Let $\mathcal{P}_\phi = (S_\phi, \Sigma, \delta_\phi, s_\phi^0, F_\phi)$ be a DPA such that $L(\mathcal{P}_\phi) = \Sigma^\omega \setminus L(\mathcal{P}_\mu)$. Therefore, $\mathcal{P}_\phi$ accepts all the sequences for which all the μ-sentences in $\mathsf{FL}(\phi)$ are regenerated finitely often.

Definition 9. *Based on Definition 7 and Definition 8, we define a turn-based two-player parity game* $\mathcal{G}_\phi = (V' = V_0' \uplus V_1', E', v_{\text{init}}', \Xi)$ *where*

- $V_0' = V_0 \times (S_\phi \cup \{\diamond\})$, $V_1' = V_1 \times (S_\phi \cup \{\diamond\})$, $v_{\text{init}}' = (v_{\text{init}}, s_\phi^0)$,
- *for every* $v \in V$ *such that* $v \notin V_t$, $s \in S_\phi$ *and* $(v, v') \in E$*: then* $((v, s), (v', \delta_\phi(s, v))) \in E'$*;*
- *for every* $v \in V_t$ *and* $s \in S_\phi$*:*
  - $((v, s), (\perp, \diamond)) \in E'$ *if* $(v, \perp) \in E$*;*
  - $((v, s), (C, \diamond)) \in E'$ *if* $(v, C) \in E$*;*
  - $((C, \diamond), (C_j, \diamond)) \in E'$ *for every* $(C, C_j) \in E$*;*
  - $((C_j, \diamond), (c, \diamond)) \in E'$ *for every* $c \in C_j$*;*
  - $((c, \diamond), (v', \delta_\phi(s, v))) \in E'$ *for every* $v' \in c$*.*
- $\Xi : V' \to \{0, \cdots, k\}$ *such that* $\Xi(v, s) = F_\phi(s)$ *for all* $s \in S_\phi$*,*

where $\diamond$ *is a placeholder,* C, C_j *and* c *are the vertices as in Definition 7.*

Lemma 1. *Player-0 has a winning strategy in the game* $\mathcal{G}_\phi$ *for every satisfiable sentence* ϕ*.*

To show the other direction of Lemma 1, we shall show how to construct a model from a winning strategy of Player-0. W.l.o.g., we assume that $\{1, \cdots, g\}$ is the set of all players appeared in ϕ. It is well-known that, for parity games, if Player-0 has some winning strategies, then Player-0 has a pure memoryless winning strategy (Gurevich and Harrington 1982). Let $\zeta' : V_0' \to V'$ be a winning strategy of Player-0 in $\mathcal{G}_\phi$. Since the DPA $\mathcal{P}_\phi$ is deterministic, we can directly extract from $\zeta' : V_0' \to V'$ a winning strategy $\zeta : V_0 \to V$ for Player-0 by projecting from V' onto V.

The winning strategy ζ and the game $\mathcal{G}_\phi$ together yield the digraph $\mathcal{G}_\phi^\zeta$ in which only the edges specified by the strategy ζ are retained. Let Π be the set of all finite paths $\vec{v} = v_1 \cdots v_r \in V_0^+(V_{1,0} \cup \{\top\})$ in $\mathcal{G}_\phi^\zeta$ such that v_1 is either the initial vertex v_{init} or there is an edge (v, v_1) in $\mathcal{G}_\phi^\zeta$ such that $v \in V_{1,1}$. Note that Π is a finite set. We denote by $\mathsf{fst}(\vec{v})$ and $\mathsf{lst}(\vec{v})$ the vertex v_1 and v_r, respectively. Given a vertex $v \in V_0$, let $\Pi_v \subseteq \Pi$ be the set of finite paths starting from v.

Let AP be the set of atomic propositions appearing in ϕ, $\mathcal{M}_\phi = (\mathsf{Ag}, \mathsf{Act}, Q, \Gamma, \delta, \lambda, q^0)$ be a PCGS obtained from $\mathcal{G}_\phi^\zeta$, where $\mathsf{Ag} := \{1, \cdots, g+1\}$, $\mathsf{Act} := \bigcup_{i \in \mathsf{Ag}} \mathsf{Act}_i$, $\mathsf{Act}_i := \bigcup_{q \in Q} \Gamma_i(q)$, $Q := \{q_{\vec{v}} \mid \vec{v} \in \Pi\} \cup \{q_\top\}$, $q^0 := q_{\vec{v}}$ such that $\mathsf{fst}(\vec{v}) = v_{\text{init}}$, $\lambda(q_{\vec{v}}) := (\bigcup_{i \in [r-1]} v_i) \cap \mathsf{AP}$, for $\vec{v} = v_1 \cdots v_r \in \Pi$ and $\lambda(q_\top) := \emptyset$. $\Gamma := (\Gamma_i)_{i \in \mathsf{Ag}}$ and δ are defined as follows.

Consider a path $\vec{v} = v_1 \cdots v_r \in \Pi$ such that $v_r \neq \top$, then $v_{r-1} \in V_t$, we assume that

$$\langle\langle v_{r-1}\rangle\rangle := \{\langle\langle A_1\rangle\rangle^{\bowtie_1 k_1}\varphi_1, \cdots, \langle\langle A_m\rangle\rangle^{\bowtie_m k_m}\varphi_m\}$$
$$[[v_{r-1}]] := \{[[B_1]]^{\sim_1 h_1}\psi_1, \cdots, [[B_n]]^{\sim_n h_n}\psi_n\}.$$

Let $\iota(\langle\langle A_j\rangle\rangle^{\bowtie_j k_j}\varphi_k)$ (resp. $\iota([[B_j]]^{\sim_j h_j}\psi_j)$) denote the index j.

Case 1: If $[[v_{r-1}]] = \bigcup_{i \in [m]} A_i = \emptyset$, then we use one action a and let $\Gamma_i(q_{\vec{v}}) := \{a\}$ for every $i \in \mathsf{Ag}$.

Case 2: Otherwise, we consider $m + n$ distinct actions $a_1, ..., a_m, b_1, ..., b_n$. Let $\Gamma_{g+1}(q_{\vec{v}}) := \{b_1, \cdots, b_n\}$ and $\Gamma_i(q_{\vec{v}})$ be the least set of actions satisfying the following conditions:

1. for every $j \in [m]$ and $i \in A_j$, $a_j \in \Gamma_i(q_{\vec{v}})$
2. for every $j \in [n]$ and $i \in \overline{B_j}$, $b_j \in \Gamma_i(q_{\vec{v}})$.

In the sequel, we define δ. We first extract essential information from $\vec{v}$. Suppose $\mathtt{mis}^{B_j}_{v_{r-1}} := \{\mathfrak{M}^j_1, \cdots, \mathfrak{M}^j_{\ell_j}\}$, for $j \in [n]$ and $\mathtt{mis}_{v_{r-1}} \setminus \bigcup_{j\in[n]} \mathtt{mis}^{B_j}_{v_{r-1}} := \{\mathfrak{M}^0_1, \cdots, \mathfrak{M}^0_{\ell_0}\}$. Note that in **Case 1**, there is only one MIS that is $\langle\langle v_{r-1}\rangle\rangle$.

By the construction of $\mathcal{G}_\phi$, for every $\mathfrak{M} \in \mathtt{MIS}_{v_{r-1}}$, there is a weighted cover (c, w) such that $|c| \leq |\mathfrak{M}| + 1$, $(c, w) \models \mathfrak{M}$. Let $w_{\mathfrak{M}}$ be the distribution such that $w_{\mathfrak{M}}(q_{\vec{v'}}) = w(\mathsf{fst}(\vec{v'}))$ for all $\vec{v'} \in \bigcup_{v\in c} \Pi_v$, $w_\top$ be the distribution such that $w_{\mathfrak{M}}(q_\top) = 1$. Moreover, $v_r \equiv \{c \mid (c, w) \models \mathfrak{M},\ \mathfrak{M} \in \mathtt{MIS}_{v_{r-1}}\}$.

To define δ, we introduce some auxiliary notations.

- For a sentence $\psi = \langle\langle A\rangle\rangle^{\bowtie k}\varphi \in \langle\langle v_{r-1}\rangle\rangle$, let
$$\mathsf{D}[\psi]_\exists := \{\mathsf{d} \in \mathsf{D}(q_{\vec{v}}) \mid \forall i \in A,\ \mathsf{d}(i) = a_{\iota(\psi)}\}.$$
- For an MIS $\mathfrak{M} \in \{\mathfrak{M}^0_1, \cdots, \mathfrak{M}^0_{\ell_0}\}$, let
$$\mathsf{D}[\mathfrak{M}]_\exists := \bigcap_{\psi\in\mathfrak{M}} \mathsf{D}[\psi]_\exists \text{ and } \mathsf{D}_\exists := \bigcup_{\mathfrak{M}\in\{\mathfrak{M}^0_1,\cdots,\mathfrak{M}^0_{\ell_0}\}} \mathsf{D}[\mathfrak{M}]_\exists.$$
- For a sentence $\varphi = [[B]]^{\sim h}\psi \in [[v_{r-1}]]$, let
$$\mathsf{D}[\varphi]_\forall := \{\mathsf{d} \in \mathsf{D}(q_{\vec{v}}) \mid \forall i \in \overline{B},\ \mathsf{d}(i) = b_{\iota(\varphi)}\}.$$
- For a sentence $\varphi = [[B]]^{\sim h}\psi \in [[v_{r-1}]]$ and a B-MIS $\mathfrak{M} \in \mathtt{mis}^B_{v_{r-1}}$, let
$$\mathsf{D}[\mathfrak{M}, \varphi]_\forall := \mathsf{D}[\mathfrak{M}]_\exists \cap \mathsf{D}[\varphi]_\forall \text{ and}$$
$$\mathsf{D}_\forall := \bigcup_{[[B]]^{\sim h}\psi\in[[v_{r-1}]],\mathfrak{M}\in\mathtt{mis}^B_{v_{r-1}}} \mathsf{D}[\mathfrak{M}, [[B]]^{\sim h}\psi]_\forall.$$

The following proposition reveals the property of action assignments of agents.

Proposition 3. *The following statements hold:*

1. *For every* $\mathfrak{M} \in \{\mathfrak{M}^0_1, \cdots, \mathfrak{M}^0_{\ell_0}\}$, $[[B]]^{\sim h}\psi$ *and* $\mathfrak{M}' \in \mathtt{mis}^B_{v_{r-1}}$*:*
$$\mathsf{D}[\mathfrak{M}]_\exists \cap \mathsf{D}[\mathfrak{M}', [[B]]^{\sim h}\psi]_\forall = \emptyset.$$
2. *For every* $\mathfrak{M}_1, \mathfrak{M}_2 \in \{\mathfrak{M}^0_1, \cdots, \mathfrak{M}^0_{\ell_0}\}$*, if* $\mathfrak{M}_1 \neq \mathfrak{M}_2$*, then*
$$\mathsf{D}[\mathfrak{M}_1]_\exists \cap \mathsf{D}[\mathfrak{M}_2]_\exists = \emptyset.$$
3. *For every* $[[B]]^{\sim h}\psi, [[B']]^{\sim' h'}\psi'$, $\mathfrak{M} \in \mathtt{mis}^B_{v_{r-1}}$ *and* $\mathfrak{M}' \in \mathtt{mis}^{B'}_{v_{r-1}}$*, if* $\mathfrak{M} \neq \mathfrak{M}'$ *or* $[[B]]^{\sim h}\psi \neq [[B']]^{\sim' h'}\psi'$*, then*
$$\mathsf{D}[\mathfrak{M}, [[B]]^{\sim h}\psi]_\forall \cap \mathsf{D}[\mathfrak{M}', [[B']]^{\sim' h'}\psi']_\forall = \emptyset.$$
4. *For every* $\varphi, \varphi' \in \langle\langle v_{r-1}\rangle\rangle$*, if* $\varphi \neq \varphi'$*, then*
$$(\mathsf{D}[\varphi]_\exists \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)) \cap (\mathsf{D}[\varphi']_\exists \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)) = \emptyset.$$
5. *For every* $\varphi, \varphi' \in [[v_{r-1}]]$*, if* $\varphi \neq \varphi'$*, then*
$$(\mathsf{D}[\varphi]_\forall \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)) \cap (\mathsf{D}[\varphi']_\forall \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)) = \emptyset.$$
6. *For every* $\varphi \in \langle\langle v_{r-1}\rangle\rangle$ *and* $\varphi' \in [[v_{r-1}]]$*,*
$$(\mathsf{D}[\varphi]_\exists \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)) \cap (\mathsf{D}[\varphi']_\forall \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)) = \emptyset.$$

We now define δ as follows.

1. For $\mathfrak{M} \in \{\mathfrak{M}^0_1, \cdots, \mathfrak{M}^0_{\ell_0}\}$ and $\mathsf{d} \in \mathsf{D}[\mathfrak{M}]_\exists$, $\delta(q_{\vec{v}}, \mathsf{d}) := w_{\mathfrak{M}}$.
2. For all $\varphi = [[B]]^{\sim h}\psi \in [[v_{r-1}]]$, $\mathfrak{M} \in \mathtt{mis}^B_{v_{r-1}}$ and $\mathsf{d} \in \mathsf{D}[\mathfrak{M}, \varphi]_\forall$, $\delta(q_{\vec{v}}, \mathsf{d}) := w_{\mathfrak{M}}$.
3. For all $\varphi \in \langle\langle v_{r-1}\rangle\rangle$ and $\mathsf{d} \in \mathsf{D}[\varphi]_\exists \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)$, for some $\mathfrak{M} \in \mathtt{MIS}_{v_{r-1}}$ such that $\varphi \in \mathfrak{M}$, $\delta(q_{\vec{v}}, \mathsf{d}) := w_{\mathfrak{M}}$. Note that such $\mathfrak{M}$ always exists.
4. For all $\varphi = [[B]]^{\sim h}\psi \in [[v_{r-1}]]$ and $\mathsf{d} \in \mathsf{D}[\varphi]_\forall \setminus (\mathsf{D}_\exists \cup \mathsf{D}_\forall)$, for some $\mathfrak{M} \in \mathtt{MIS}_{v_{r-1}}$ such that $\varphi \in \mathfrak{M}$, $\delta(q_{\vec{v}}, \mathsf{d}) := w_{\mathfrak{M}}$. Note that such $\mathfrak{M}$ always exists as well.
5. For all the other $\mathsf{d} \in \mathsf{D}(q_{\vec{v}})$, $\delta(q_{\vec{v}}, \mathsf{d}) = w_\top$.

Example 6. *Consider the game in* Figure 1 *and the strategy* ζ *with* $\zeta(C_2) = c_3$, $\zeta(v_4) = v_7$ *and* $\zeta(v_5) = v_8$*. Then,* ζ *is a winning strategy for Player-*0 *from which we can construct the digraph* $\mathcal{G}^\zeta_\Psi$*.* $\mathcal{G}^\zeta_\Psi$ *is the subgraph of the game shown in Figure 1 consisting of the vertices* $\{v_1, \cdots, v_9, C, C_1, C_2, c_1, c_3, \top\}$ *and related edges between them.* Π *consists of four finite paths* $\{v_1v_2v_3C, v_4v_7\top, v_5v_8\top, v_6v_9\top\}$ *(blue paths in Figure 1). From* Π*, using the weights defined in Example 4, we get the model of* ϕ *as shown in Figure 2, where players are* $\{1, 2, 3, 4\}$ *and actions are* $\{a_1, a_2, b_1\}$ *with* $\Gamma_1(q_0) = \{a_1\}$, $\Gamma_2(q_0) = \{a_1, b_1\}$ *and* $\Gamma_3(q_0) = \{a_2\}$, $\lambda(q_1) = \{p_1, p_2\}$, $\lambda(q_2) = \{p_2\}$ *and* $\lambda(q_3) = \emptyset$*.*

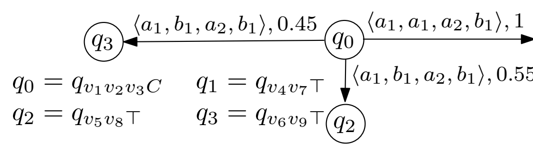

Figure 2: The model of Ψ.

By the construction (in particular, the function δ and finite path set Π), we can verify that

Lemma 2. $\mathcal{M}_\phi$ *satisfies* ϕ*.*

Theorem 3. *The satisfiability problem for PAMC is* EXPTIME*-complete. Moreover, if* ϕ *is satisfiable, we can construct a model of* ϕ *in exponential size of* $|\phi|$ *such that*

- *the number of players is bounded by* $k + 1$*, where* k *is the number of the players occurring in* ϕ*,*
- *the number of actions is bounded by* $|\mathsf{FL}_\exists(\phi)| + |\mathsf{FL}_\forall(\phi)| + 1$*,*
- *the out-degree is bounded by* $|\mathsf{FL}_\exists(\phi)| + 2$*.*

Remark 1. *Without adding the player* $g + 1$*, our reduction still works if* $\mathsf{Ag} \neq (B \cup B')$ *for each distinct pair of* $[[B]]^{\sim h}\psi$ *and* $[[B']]^{\sim' h'}\psi'$ *in* $[[v]]$*. The case* $\mathsf{Ag} = B \cup B'$ *could be solved by adapting the reduction by adding new MISs* $\mathfrak{M}$ *into* $\mathtt{MIS}_v$*, which takes into account the common decisions for* $[[B]]^{\sim h}\psi$ *and* $[[B']]^{\sim' h'}\psi'$*. We leave it as future work.*

Related Work

Probabilistic temporal logics such as PCTL, PCTL*, PLTLK, μ-calculi and their variants (Vardi 1985; Morgan and McIver 1997; Huth and Kwiatkowska 1997; de Alfaro and Majumdar 2001; McIver and Morgan 2002; 2007; Mio 2012; Cleaveland, Iyer, and Narasimha 2005; Huang and Kwiatkowska 2016; Fu et al. 2018) can express probabilistic properties, rather than coalitions of players.

The most relevant work includes PATL and PATL* (Chen and Lu 2007; Huang, Su, and Zhang 2012), their variants (Chen et al. 2013; Huang and Luo 2013) and SGL (Baier et al. 2007). These papers mainly focus on the model checking problem (on different game structures or interpretations of strategies). The satisfiability problem, which is the main focus of the current paper, is a long-standing open problem for PATL and PATL*, notwithstanding with some progress (Brázdil et al. 2008; Chakraborty and Katoen 2016).

Satisfiability checking of (non-probabilistic) alternating-time temporal logics has been considered (Goranko and Vester 2014). Our decision procedure shares some similarity, specifically, MIS plays a similar role as the distributed control therein. Modulo the probabilistic aspects and fixpoints, the main difference is that, to handle coalition strategies we are considering *maximally* independent sets rather than all subsets of distributed controls. As a result, our construction leads to potentially smaller models, which is important for implementation, but is at the cost of a considerably more complicated construction.

Much research has been dedicated to algorithms of solving stochastic games, e.g., (Chatterjee 2007; Chatterjee and Henzinger 2012; Brázdil et al. 2006) in which objectives are usually given by ω-regular properties or PCTL. Stochastic games are useful for synthesis and verification of reactive systems, but hard to express cooperation of players.

Conclusion and Future Work

We proposed a temporal logic PAMC for reasoning about strategic abilities of agents in stochastic multi-agent systems. We showed that the complexities of its model checking and satisfiability checking problems lie in the same complexity classes of the respective problems of modal μ-calculus and AMC. We also implemented a satisfiability solver for PAMC.

Several questions are left open for PAMC, such as (complete) axiomatization and extensions with epistemic operators. We leave them as future work.

References

Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *J. ACM* 49(5):672–713.

Baier, C.; Brázdil, T.; Größer, M.; and Kucera, A. 2007. Stochastic game logic. In *QEST*, 227–236.

Brázdil, T.; Brozek, V.; Forejt, V.; and Kucera, A. 2006. Stochastic games with branching-time winning objectives. In *LICS*, 349–358.

Brázdil, T.; Forejt, V.; Kretínský, J.; and Kucera, A. 2008. The satisfiability problem for probabilistic CTL. In *LICS*, 391–402.

Chakraborty, S., and Katoen, J. P. 2016. On the satisfiability of some simple probabilistic logics. In *LICS*, 56–65.

Chatterjee, K., and Henzinger, T. A. 2012. A survey of stochastic ω-regular games. *J. Comput. Syst. Sci.* 78(2):394–413.

Chatterjee, K.; Henzinger, T. A.; and Piterman, N. 2010. Strategy logic. *Information and Computation* 208(6):677–693.

Chatterjee, K. 2007. *Stochastic Omega-Regular Games*. Ph.D. Dissertation, University of California, Berkeley.

Chen, T., and Lu, J. 2007. Probabilistic alternating-time temporal logic and model checking algorithm. In *FSKD*, 35–39.

Chen, T.; Forejt, V.; Kwiatkowska, M. Z.; Parker, D.; and Simaitis, A. 2013. Automatic verification of competitive stochastic systems. *Form. Method. Syst. Des.* 43(1):61–92.

Cleaveland, R.; Iyer, S.; and Narasimha, M. 2005. Probabilistic temporal logics via the modal μ-calculus. *Theor. Comput. Sci.* 342(2-3):316–350.

de Alfaro, L., and Majumdar, R. 2001. Quantitative solution of omega-regular games. In *STOC*, 675–683.

Foster, H.; Krolnik, A.; and Lacey, D. 2004. *Assertion-based design (2. ed.)*. Kluwer.

Fu, C.; Turrini, A.; Huang, X.; Song, L.; Feng, Y.; and Zhang, L. 2018. Model checking probabilistic epistemic logic for probabilistic multiagent systems. In *IJCAI*, 4757–4763.

Goranko, V., and Vester, S. 2014. Optimal decision procedures for satisfiability in fragments of alternating-time temporal logics. In *AIML*, 234–253.

Gurevich, Y., and Harrington, L. 1982. Trees, automata, and games. In *STOC*, 60–65.

Huang, X., and Kwiatkowska, M. 2016. Model checking probabilistic knowledge: A PSPACE case. In *AAAI*, 2516–2522.

Huang, X., and Luo, C. 2013. A logic of probabilistic knowledge and strategy. In *AAMAS*, 845–852.

Huang, X.; Su, K.; and Zhang, C. 2012. Probabilistic alternating-time temporal logic of incomplete information and synchronous perfect recall. In *AAAI*.

Huth, M., and Kwiatkowska, M. Z. 1997. Quantitative analysis and model checking. In *LICS*, 111–122.

Liu, W.; Song, L.; Wang, J.; and Zhang, L. 2015. A simple probabilistic extension of modal mu-calculus. In *IJCAI*, 882–888.

McIver, A., and Morgan, C. 2002. Games, probability and the quantitative μ-calculus qmμ. In *LPAR*, 292–310.

McIver, A., and Morgan, C. 2007. Results on the quantitative μ-calculus qmμ. *ACM Trans. Comput. Log.* 8(1):3.

Mio, M. 2012. *Game semantics for probabilistic modal μ-calculi*. Ph.D. Dissertation, The University of Edinburgh.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. Y. 2014. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.* 15(4):34:1–34:47.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. Y. 2017. Reasoning about strategies: on the satisfiability problem. *Log. Meth. Comput. Sci.* 13(1).

Moon, J., and Moser, L. 1965. On cliques in graphs. *Israel J. Math.* 3(1):23–28.

Morgan, C., and McIver, A. 1997. A probabilistic temporal calculus based on expectations. In *FMP*, 4–22.

Pauly, M. 2011. *Logic for Social Software*. Ph.D. Dissertation, University of Amsterdam.

Rozier, K. Y., and Vardi, M. Y. 2010. LTL satisfiability checking. *STTT* 12(2):123–137.

Schewe, S., and Finkbeiner, B. 2006. Satisfiability and finite model property for the alternating-time *mu*-calculus. In *CSL*, 591–605.

Schewe, S. 2008. ATL* satisfiability is 2EXPTIME-complete. In *ICALP*, 373–385.

Shmulevich, I., and Dougherty, E. R. 2010. *Probabilistic Boolean Networks - The Modeling and Control of Gene Regulatory Networks*. SIAM.

Vardi, M. Y. 1985. Automatic verification of probabilistic concurrent finite-state programs. In *FOCS*, 327–338.

Walther, D.; Lutz, C.; Wolter, F.; and Wooldridge, M. 2006. ATL satisfiability is indeed EXPTIME-complete. *J. Logic Comput.* 16(6):765–787.