

On Probabilistic Alternating Simulations

Chenyi Zhang, Jun Pang

► **To cite this version:**

Chenyi Zhang, Jun Pang. On Probabilistic Alternating Simulations. 6th IFIP TC 1/WG 2.2 International Conference on Theoretical Computer Science (TCS) / Held as Part of World Computer Congress (WCC), Sep 2010, Brisbane, Australia. pp.71-85, 10.1007/978-3-642-15240-5_6 . hal-01054443

HAL Id: hal-01054443

<https://hal.inria.fr/hal-01054443>

Submitted on 6 Aug 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



On Probabilistic Alternating Simulations

Chenyi Zhang^{1,2} and Jun Pang¹

¹ Faculty of Sciences, Technology and Communication
University of Luxembourg, Luxembourg

² School of Computer Science and Engineering
University of New South Wales, Australia

Abstract. This paper presents simulation-based relations for probabilistic game structures. The first relation is called probabilistic alternating simulation, and the second called probabilistic alternating forward simulation, following the naming convention of Segala and Lynch. We study these relations with respect to the preservation of properties specified in probabilistic alternating-time temporal logic.

1 Introduction

Simulation relations [15] have proved to be useful for comparing the behavior of concurrent systems, which can be formally interpreted as labeled transition systems. The study of logic characterization of simulation is to build its connection to a modal or temporal logic which can be used to formulate some interesting properties. Soundness of logic characterization requires simulation preserve the satisfaction of logic formulas, while completeness shows the relation has the same strength as the logic. Intuitively, the fact that one state s_1 simulates another state s_2 can be used to establish the relation that any possible behavior of s_1 is also possible on s_2 . Thus it can preserve certain desirable properties formulated in temporal logics like CTL [11]. Simulation relations have set up the foundations for constructing correct abstractions.

Related work. Segala and Lynch [21] extend the classical notions of simulation for probabilistic automata [20], a general extension of labeled transition systems which admits both probabilistic and nondeterministic behaviors. Their main idea is to relate probability distributions over states, instead of relating individual states. They show soundness of the logical characterization of probabilistic simulation, which preserves probabilistic CTL formulas [12] without negation and existential quantification. Segala introduces probabilistic forward simulation, which relates states to probability distributions over states and is sound and complete for trace distribution precongruence [19, 13]. Logic characterization of strong and weak probabilistic bisimulation has been studied in [10, 17].

Alur, Henzinger and Kupferman [1, 2] define ATL (alternating-time temporal logic) to generalize CTL for game structures by requiring each path quantifier to be parametrized with a set of agents. Game structures are more general than LTS, in the sense that they allow both collaborative and adversarial behaviors

of individual agents in a system, and ATL can be used to express properties like “a set of agents can enforce a specific outcome of the system”. Alternating refinement relations, in particular alternating simulation, are introduced later in [3]. Alternating simulation is a natural game-theoretic interpretation of the classical simulation in two-player games. Logic characterization of this relation concentrates on a subset of ATL* formulas where negations are only allowed at proposition level and all path quantifiers are parametrized by a predefined set of agents A . This sublogic of ATL* contains all formulas expressing the properties that agents in A can enforce no matter what the other agents do. Alur et al. [3] have proved both soundness and completeness of their characterization.

Our contribution. Extending game structures with probabilistic behaviors of players gives rise to a more expressive framework for modeling (open) systems. Mixed strategies, which allow for players to randomly select their actions, are often necessary for the players to achieve their expected rewards [16]. As the papers [3, 2] only focus on pure strategies, it is a natural step to study the corresponding notion of simulation in a probabilistic game-based setting.

In this paper, we introduce two notions of simulation for probabilistic game structures — probabilistic alternating simulation and forward simulation, following the aforementioned results [19, 21, 3]. We prove the soundness of logical characterization of probabilistic alternating simulation relations, by showing that they preserve a fragment of a probabilistic extension of ATL.

Outline. The rest of the paper is organized as follows. We briefly explain some basic notations that are used throughout the paper in Sect. 2. Sect. 3 introduces the notion of probabilistic game structures and the definition of probabilistic executions. In Sect. 4 we present PATL an extension of the alternating-time temporal logic [2] for probabilistic systems, and roughly discuss its model checking problem. We define probabilistic alternating simulation and forward simulation in Sect. 5, and show their soundness for preserving properties specified in PATL in Sect. 6. Probabilistic alternating bisimulation is shortly discussed in Sect. 7. We conclude the paper with some future research topics in Sect. 8.

2 Preliminaries

This section contains basic notions that are used in the technical part. Let S be a set, then a discrete probabilistic distribution Δ over S is a function of type $S \rightarrow [0, 1]$, satisfying $\sum_{s \in S} \Delta(s) = 1$. We write $\mathcal{D}(S)$ for the set of all such distributions. For a set $S' \subseteq S$, define $\Delta(S') = \sum_{s \in S'} \Delta(s)$. Given two distributions Δ_1, Δ_2 and $p \in [0, 1]$, $\Delta_1 \oplus_p \Delta_2$ is a function of type $S \rightarrow [0, 1]$ defined as $\Delta_1 \oplus_p \Delta_2(s) = p \cdot \Delta_1(s) + (1 - p) \cdot \Delta_2(s)$ for all $s \in S$. Obviously, $\Delta_1 \oplus_p \Delta_2$ is also a distribution. We further extend this notion by combining a set of distributions $\{\Delta_i\}_{i \in I}$ ordered by an indexed set $\{p_i\}_{i \in I}$ into a distribution $\sum_{i \in I} p_i \Delta_i$, where $p_i \in [0, 1]$ for all $i \in I$ and $\sum_{i \in I} p_i = 1$. \bar{s} is called a point distribution satisfying $\bar{s}(s) = 1$ and $\bar{s}(t) = 0$ for all $t \neq s$. Let $\Delta \in \mathcal{D}(S)$, write $[\Delta]$ for the *support* of Δ as the set $\{s \in S \mid \Delta(s) > 0\}$.

Let $S = S_1 \times S_2 \times \dots \times S_n$, then $\mathbf{s} \in S$ is a vector of length n . We may also write $\mathbf{s} = \langle s_1, s_2, \dots, s_n \rangle$, with $\mathbf{s}(i) = s_i \in S_i$. Given a finite sequence $\alpha = s_1 s_2 \dots s_n \in S^*$, write $\text{last}(\alpha)$ for s_n . Let $S' \subseteq S$, then $\alpha \mid S'$ is a subsequence of α with exactly the elements not in S' removed. Given $L \subseteq S^*$, write $L \mid S'$ for the set $\{(\alpha \mid S') \mid \alpha \in L\}$.

3 Probabilistic Game Structures

Assume a set of players $\Sigma = \{1, 2, \dots, \mathbf{k}\}$. A probabilistic game structure (PGS) \mathcal{G} is defined as a tuple $\langle S, s_0, \mathcal{L}, \mathbf{Act}, \delta \rangle$, where

- S is a finite set of states, with s_0 the initial state,
- $\mathbf{Act} = \mathbf{Act}_1 \times \mathbf{Act}_2 \times \dots \times \mathbf{Act}_{\mathbf{k}}$ is a set of joint actions, where \mathbf{Act}_i is the set of actions for player $i = 1, \dots, \mathbf{k}$,
- $\mathcal{L} : S \rightarrow 2^{\text{Prop}}$ is the labeling function,
- $\delta : S \times \mathbf{Act} \rightarrow \mathcal{D}(S)$ is a transition function.

A play ρ is a (finite or infinite) sequence $s_0 \mathbf{a}_1 s_1 \mathbf{a}_2 s_2 \dots$, such that $\mathbf{a}_i \in \mathbf{Act}$ and $\delta(s_{i-1}, \mathbf{a}_i)(s_i) > 0$ for all i . Write $|\rho|$ for the length of a run ρ , which is the number of transitions in ρ , and $|\rho| = \infty$ if ρ is infinite. We write $\rho(i)$ for the i -th state in ρ starting from 0, and $\rho[i, j]$ for the subsequence starting from i -th state and ending at the j -th state, provided $0 \leq i \leq j \leq |\rho|$. Note that the players choose their next moves simultaneously, but their moves may or may not be cooperative. If on state s each player i performs action a_i , then $\delta(s, \langle a_1, a_2, \dots, a_{\mathbf{k}} \rangle)$ is the distribution for the next reachable states. In the following discussion, we fix a probabilistic game structure \mathcal{G} .

We assume that the transition function δ is total on the set \mathbf{Act} . Note that this does not pose any limitation on the expressiveness of the model. If an action $c \in \mathbf{Act}_i$ of player i is not supposed to be enabled on state s for player i , we may find another action $c' \in \mathbf{Act}_i$ and define c to have the same effect as c' on s . Since player i knows the current state, he also knows the set of actions available to him, so that as a rational player he will not choose actions that are not enabled. This allows such models to express systems in which on some states the available (joint) actions are proper subsets of \mathbf{Act} .³ We may even disable a particular player on a state. A player i is disabled on s if $\delta(s, \mathbf{a}) = \delta(s, \mathbf{a}')$ for all action vectors $\mathbf{a}, \mathbf{a}' \in \mathbf{Act}$ satisfying $\mathbf{a}(j) = \mathbf{a}'(j)$ for all $j \neq i$. A PGS is *turn-based* if all but one player is disabled on s for all $s \in S$.

A strategy of a player $i \in \Sigma$ is a function of type $S^+ \rightarrow \mathcal{D}(\mathbf{Act}_i)$. We write $\Pi_i^{\mathcal{G}}$ for the set of strategies of player i in \mathcal{G} . A play ρ is compatible with an i -strategy π_i , if $\mathbf{a}_k(i) \in [\pi_i(\rho[0, k-1] \mid S)]$ for all $k \leq |\rho|$. Given a vector of strategies $\boldsymbol{\pi} \in \Pi_1^{\mathcal{G}} \times \Pi_2^{\mathcal{G}} \times \dots \times \Pi_{|\Sigma|}^{\mathcal{G}}$, a run ρ is compatible with $\boldsymbol{\pi}$ if $\mathbf{a}_k(i) \in [\boldsymbol{\pi}(i)(\rho[0, k-1] \mid S)]$ for all $k \leq |\rho|$ and $i = 1, \dots, \mathbf{k}$. Write $\mathcal{G}(\boldsymbol{\pi}, s)$ for the set of infinite plays compatible with every strategy in $\boldsymbol{\pi}$ starting from $s \in S$, and $\mathcal{G}^*(\boldsymbol{\pi}, s)$ the set of finite plays in \mathcal{G} that are compatible with $\boldsymbol{\pi}$ starting from s .

³ In the literature some authors encode available actions for player i as a function of type $S \rightarrow 2^{\mathbf{Act}_i} \setminus \{\emptyset\}$.

The set of finite plays compatible to strategy vector $\boldsymbol{\pi}$ is also called a set of *cones* [20], with each finite play α representing the set of infinite plays prefixed by α . Given a state $s_0 \in S$, we can derive the probability for every member in S^+ compatible with $\boldsymbol{\pi}$, by recursively defining a function $Pr_{\mathcal{G}(\boldsymbol{\pi}, s_0)}$ from S^+ to $[0, 1]$ as follows. This function $Pr_{\mathcal{G}(\boldsymbol{\pi}, s_0)}$ can be further generalized as the probability measure to the σ -field $\mathcal{F}_{\mathcal{G}, \boldsymbol{\pi}, s_0} \subseteq \mathcal{G}(\boldsymbol{\pi}, s_0)$ which is a unique extension from the set of cones $\mathcal{G}^*(\boldsymbol{\pi}, s)$ closed by countable union and complementation, in a way similar to [20]:

$$\begin{aligned} & - Pr_{\mathcal{G}(\boldsymbol{\pi}, s_0)}(s_0) = 1, \\ & - Pr_{\mathcal{G}(\boldsymbol{\pi}, s_0)}(\alpha \cdot s) = Pr_{\mathcal{G}(\boldsymbol{\pi}, s_0)}(\alpha) \cdot \bar{\delta}(\text{last}(\alpha), \langle \boldsymbol{\pi}(1)(\alpha), \boldsymbol{\pi}(2)(\alpha), \dots, \boldsymbol{\pi}(\mathbf{k})(\alpha) \rangle)(s), \end{aligned}$$

where $\bar{\delta}(s, \langle \Delta_1, \Delta_2, \dots, \Delta_{\mathbf{k}} \rangle)$ is a distribution over states derived from δ and the vector of action distributions defined by

$$\bar{\delta}(s, \langle \Delta_1, \dots, \Delta_{\mathbf{k}} \rangle) = \sum_{i \in \{1, \dots, \mathbf{k}\}, a_i \in [\Delta_i]} \Delta_1(a_1) \cdot \dots \cdot \Delta_{\mathbf{k}}(a_{\mathbf{k}}) \cdot \delta(s, \langle a_1, \dots, a_{\mathbf{k}} \rangle).$$

Given $A \subseteq \Sigma$, sometimes we write $\boldsymbol{\pi}(A)$ for a vector of $|A|$ strategies $\{\pi_i\}_{i \in A}$, and $\Pi(A)$ for the set of all such strategy vectors. Write \bar{A} for $\Sigma \setminus A$. Given $A \cap A' = \emptyset$, strategy vectors $\boldsymbol{\pi} \in \Pi(A)$ and $\boldsymbol{\pi}' \in \Pi(A')$, $\boldsymbol{\pi} \cup \boldsymbol{\pi}'$ is the vector of strategies $\{\pi_i\}_{i \in A} \cup \{\pi'_j\}_{j \in A'}$ that combines $\boldsymbol{\pi}$ and $\boldsymbol{\pi}'$.

We also define strategies of *finite depth* by restricting the size of their domains, by writing $\pi \in \Pi_i^{\mathcal{G}, n}$ as a *level- n* strategy, i.e., π is a function from traces of states with length up to n (i.e., the set $\bigcup_{m \in \{1, 2, \dots, n\}} S^m$) to $\mathcal{D}(\text{Act}_i)$. Given a set of strategies $\{\pi_i\}_{i \in I}$ of the same domain, and $\{p_i\}_{i \in I}$ with $\sum_{i \in I} p_i = 1$, let $\pi = \sum_{i \in I} p_i \cdot \pi_i$ be a (combined) strategy, by letting $\pi(\gamma) = \sum_{i \in I} p_i \cdot \pi_i(\gamma)$ for all γ in the domain.

We overload the function $\bar{\delta}$ as from a state in S and a vector of strategies (of any depth n) $\boldsymbol{\pi} \in \Pi_1^{\mathcal{G}, n} \times \Pi_2^{\mathcal{G}, n} \times \dots \times \Pi_{|\Sigma|}^{\mathcal{G}, n}$ to $\mathcal{D}(S)$, by $\bar{\delta}(s, \boldsymbol{\pi}) = \bar{\delta}(s, \mathbf{a})$, where $\mathbf{a}(i) = \boldsymbol{\pi}(i)(s)$ for all $i \in \Sigma$. Note each $\mathbf{a}(i)$ is a distribution over Act_i . We further lift $\bar{\delta}$ to be a transition function from state distributions and strategy vectors to state distributions, by

$$\bar{\delta}(\Delta, \boldsymbol{\pi}) = \sum_{s \in [\Delta]} \Delta(s) \cdot \bar{\delta}(s, \boldsymbol{\pi})$$

Probabilistic Executions

We settle the nondeterminism in a probabilistic game structure by fixing the behaviors of all players represented as strategies. Let $\mathcal{G} = \langle S, s_0, \mathcal{L}, \text{Act}, \delta \rangle$ be a PGS, define a *probabilistic execution* \mathcal{E} as in the form of $\langle E, \Delta, \mathcal{L}^{\mathcal{E}}, \delta^{\mathcal{E}} \rangle$, where

- $E \subseteq S^+$ is the set of finite plays starting form a state in the initial distribution and compatible with $\delta^{\mathcal{E}}$, i.e., $s_0 s_1 \dots s_n \in E$ if $s_0 \in [\Delta]$, and $\delta^{\mathcal{E}}(s_0 \dots s_i)(s_0 \dots s_{i+1}) > 0$ for all $0 \leq i < n$,
- $\Delta \in \mathcal{D}(S)$ an (initial) distribution,

- $\mathcal{L}^{\mathcal{E}}$ is the labeling function defined as $\mathcal{L}^{\mathcal{E}}(e) = \mathcal{L}(last(e))$ for all $e \in E$,
- $\delta^{\mathcal{E}} : E \rightarrow \mathcal{D}(E)$ is a (deterministic) transition relation, satisfying for all $e \in E$ there exists a (level 1) strategy vector π_e , such that $\delta^{\mathcal{E}}(e)(e \cdot t) = \bar{\delta}(last(e), \pi_e)(t)$ if $t \in [\bar{\delta}(last(e), \pi_e)]$, and 0 otherwise.

A probabilistic execution of \mathcal{G} can be uniquely determined by a strategy vector π and a state distribution. Given $\Delta \in \mathcal{D}(S)$, define $\mathcal{E}(\mathcal{G}, \pi, \Delta)$ as the probabilistic execution $\langle E^\pi, \Delta, \mathcal{L}^\pi, \delta^\pi \rangle$, with $E^\pi = \bigcup_{s \in [\Delta]} \mathcal{G}^*(\pi, s) \mid S$ for the set of compatible finite plays, \mathcal{L}^π defined as $\mathcal{L}^\pi(e) = \mathcal{L}(last(e))$ for all $e \in E^\pi$, and $\delta^\pi(e) = \bar{\delta}(last(e), \pi_e)$ for all $e \in E^\pi$, where $\pi_e(i) = \pi(i)(e)$ for all $i \in \Sigma$. Intuitively, a probabilistic execution resembles the notion of the same name proposed by Segala and Lynch [20, 21], and in this case the strategies of the players altogether represent a single adversary of Segala and Lynch.

4 Probabilistic Alternating-Time Temporal Logic

In this section we introduce a probabilistic version of alternating-time temporal logic [2], which focuses on the players ability to enforce a property with an expected probability. Let **Prop** be a nonempty set of propositions. Probabilistic alternating-time temporal logic (PATL) was initially proposed by Chen and Lu [7]. Here we show its original syntax can be slightly simplified. PATL formulas are defined as follows.

$$\begin{aligned} \phi &:= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \langle\langle A \rangle\rangle^{\bowtie\alpha} \psi \\ \psi &:= \bigcirc\phi \mid \phi_1 \mathbf{U}^{\leq k} \phi_2 \end{aligned}$$

where $A \subseteq \Sigma$ is a set of players, $\bowtie \in \{<, >, \leq, \geq\}$, $k \in \mathbb{N} \cup \{\infty\}$, $p \in \mathbf{Prop}$, and $\alpha \in [0, 1]$. We also write $\psi_1 \mathbf{U} \psi_2$ for $\psi_1 \mathbf{U}^{\leq \infty} \psi_2$ as ‘unbounded until’. The symbols ϕ, ϕ_1, ϕ_2 are state formulas, and ψ is a path formula. We omit the syntactic sugars in our definition, such as *true* $\equiv p \vee \neg p$ and *false* $\equiv p \wedge \neg p$ for some $p \in \mathbf{Prop}$, $\phi_1 \vee \phi_2 \equiv \neg(\neg\phi_1 \wedge \neg\phi_2)$ for state formulas. The path modality **R** can be expressed by **U** without introducing negations into path formulas, as we will show later in this section. One may also define $\square^{\leq k} \psi \equiv \text{false} \mathbf{R}^{\leq k} \psi$, and $\diamond^{\leq k} \psi \equiv \text{true} \mathbf{U}^{\leq k} \psi$, where $k \in \mathbb{N} \cup \{\infty\}$. The set of PATL formulas \mathbb{L} are the set of state formulas as defined above. We have the semantics of the path formulas and the state formulas defined as follows.

- $\rho \models \phi$ iff $\mathcal{G}, \rho(0) \models \phi$ where ϕ is a state formula,
- $\rho \models \bigcirc\phi$ iff $\rho(1) \models \phi$,
- $\rho \models \phi_1 \mathbf{U}^{\leq k} \phi_2$ iff there exists $i \leq k$ such that $\rho(j) \models \phi_1$ for all $0 \leq j < i$ and $\rho(i) \models \phi_2$,
- $\mathcal{G}, s \models p$ iff $p \in \mathcal{L}(s)$,
- $\mathcal{G}, s \models \neg\phi$ iff $\mathcal{G}, s \not\models \phi$,
- $\mathcal{G}, s \models \phi_1 \wedge \phi_2$ iff $\mathcal{G}, s \models \phi_1$ and $\mathcal{G}, s \models \phi_2$,
- $\mathcal{G}, s \models \langle\langle A \rangle\rangle^{\bowtie\alpha} \psi$ iff there exists a vector of strategies $\pi \in \Pi(A)$, such that for all vectors of strategies $\pi' \in \Pi(\bar{A})$ for players in \bar{A} , we have $Pr_{\mathcal{G}(\pi \cup \pi', s)}(\{\rho \in \mathcal{G}(\pi \cup \pi', s) \mid \rho \models \psi\}) \bowtie \alpha$,

where ρ is an infinite play in \mathcal{G} , $\alpha \in [0, 1]$, ϕ, ϕ_1, ϕ_2 are state formulas, and ψ is a path formula. Equivalently, given S the state space of a probabilistic game structure \mathcal{G} , we write $\llbracket \phi \rrbracket$ for $\{s \in S \mid s \models \phi\}$ for all PATL (state) formulas ϕ . For $\Delta \in \mathcal{D}(S)$, we write $\Delta \models \phi$ iff $\llbracket \Delta \rrbracket \subseteq \llbracket \phi \rrbracket$. Intuitively, $\mathcal{G}, s \models \langle\langle A \rangle\rangle^{\geq \alpha} \psi$ ($\mathcal{G}, s \models \langle\langle A \rangle\rangle^{\leq \alpha} \psi$) describes the ability of players in A to cooperatively enforce ψ with probability at least (at most) α in s .

The following lemma is directly from the PATL semantics. If a group of users A can enforce a linear-time temporal logic formula ψ to hold with probability at least α with strategies $\pi \in \Pi(A)$, then at the same time π enforces the formula $\neg\psi$ to hold with probability at most $1 - \alpha$. To simplify the notation, we let ‘ \sim ’ denote changes on directions of the symbols in $\{<, >, \leq, \geq\}$, e.g., symbol $\tilde{\geq}$ for \leq , $\tilde{\leq}$ for \geq , $\tilde{>}$ for $<$, and $\tilde{<}$ for $>$.

Lemma 1. $\mathcal{G}, s \models \langle\langle A \rangle\rangle^{\geq \alpha} \psi$ iff $\mathcal{G}, s \models \langle\langle A \rangle\rangle^{\tilde{\leq} 1 - \alpha} \neg\psi$

Therefore, the path quantifier R (release) can be expressed by the existing PATL syntax, in the way that $\langle\langle A \rangle\rangle^{\geq \alpha} \phi_1 R^{\leq k} \phi_2 \equiv \langle\langle A \rangle\rangle^{\tilde{\leq} 1 - \alpha} (\neg\phi_1) U^{\leq k} (\neg\phi_2)$, where both $\neg\phi_1$ and $\neg\phi_2$ are state formulas.

On Model Checking of PATL

In this section we briefly survey the results in the literature related to PATL model checking. Given a PATL formula in the form of $\langle\langle A \rangle\rangle^{\geq \alpha} \psi(\phi_1, \dots, \phi_n)$, a standard way to solve this problem is to determine the maximal or minimal probability that the players in A can enforce the LTL formula $\psi(\phi_1, \dots, \phi_n)$. In the following we write ψ for $\psi(\phi_1, \dots, \phi_n)$ without further confusions.

LTL properties are special cases of ω -regular winning objectives [22] in two-player concurrent (zero-sum) games [9, 6]. In such games one may group a set of players $A \subseteq \Sigma$ into a single protagonist and \bar{A} into a single antagonist. Given an ω -regular winning objective ξ and starting from a state $s \in S$, the protagonist plays with a strategy trying to maximize the probability for a play to satisfy ξ while the antagonist tries to minimize the probability. In such a game there always exists a unique value in $[0, 1]$, on which both players have strategies to guarantee (or infinitely approach) their best performances, regardless of the strategies played by their opponents. Such a supremum value (or infimum value, as for the antagonist) is called the *value of the game* [14, 9]. In a probabilistic multi-player game, we let a group of players $A \subseteq \Sigma$ be a single player, and \bar{A} be the other, and the supremal probability for A to enforce an LTL formula ψ starting from a given state $s \in S$ can be uniquely determined, as defined by

$$\langle A \rangle \psi(s) = \bigsqcup_{\pi \in \Pi(A)} \bigsqcap_{\pi' \in \Pi(\bar{A})} Pr_{\mathcal{G}(\pi \cup \pi', s)}(\{\rho \in \mathcal{G}(\pi \cup \pi', s) \mid \rho \models \psi\})$$

A vector of strategies, which does not necessarily exist, is *optimal* for a group of players, if it enforces the *value of the game* for that group.




Fig. 1. An example showing that player I can guarantee to satisfy $\diamond\phi$ with probability α for all $0 \leq \alpha < 1$, but he cannot ensure that property with probability 1.

Example 1. Fig. 1 gives a PGS with two players $\{\mathbf{I}, \mathbf{II}\}$, initial state s_0 , $\text{Act}_{\mathbf{I}} = \{a_1, a_2\}$ and $\text{Act}_{\mathbf{II}} = \{b_1, b_2\}$. Note that this PGS is deterministic, i.e. no probabilities in its transitions. We assume that the only available transitions from s_1 and s_2 are self-loops, and the other transition relations are as depicted in the graph. Suppose player I wants to maximize the probability to enforce the property $\diamond\phi$, and player II aims to minimize it.

Since the strategies applied on s_1 and s_2 do not matter, we focus on the choices of actions from both players on s_0 . We first focus on memoryless strategies, and let player I's strategy π_1 gives $\pi_1(\gamma)(a_1) = p$ and $\pi_1(\gamma)(a_2) = 1 - p$ for all $\gamma \in S^+$. Similarly we let II assign probability q to b_1 and $1 - q$ to b_2 all the time. This produces an infinite tree, on which we write $x_{s_0}(\mathbf{I})$ for the actual probability I can achieve $\diamond\phi$ from s_0 , given the above memoryless strategies. (Note that $x_{s_1}(\mathbf{I}) = 0$ and $x_{s_2}(\mathbf{I}) = 1$ in all cases.) This establishes an equation which further derives $x_{s_0}(\mathbf{I}) = \frac{(1-p) + (2p-1)q}{(1-p) + pq}$. A simple analysis shows that when p approaches 1, the minimal value of $x_{s_0}(\mathbf{I})$ approaches 1 as well, for all choices of q . That is, there exists a strategy for player I to enforce $\diamond\phi$ with probability $1 - \varepsilon$ for all $\varepsilon > 0$. However, if player I chooses $p = 1$, player II may set $q = 0$ so that a play will be trapped in s_0 for ever that yields $x_{s_0}(\mathbf{I}) = 0$. The result of [9] shows that in this case player I cannot do better even with general (history dependent) strategies. In fact there are no strategies for player I to enforce $\diamond\phi$ with probability 1. \square

Indeed, $\langle A \rangle \psi(s)$ can be *almost* the best, i.e., we have $\mathcal{G}, s \models \langle\langle A \rangle\rangle \geq \langle A \rangle \psi(s) - \varepsilon \psi$ for all $\varepsilon > 0$ [8]. Nevertheless, the quantitative version of determinacy [14] ensures that for all LTL formulas ψ and $s \in S$, we have

$$\langle A \rangle \psi(s) + \langle \bar{A} \rangle \neg \psi(s) = 1$$

The PATL model checking problems can be solved by calculating the values $\langle A \rangle \psi_s(s)$ for each state s , where each local objective ψ_s related to s might be distinct. The algorithms of [9] define monotonic functions of type $(S \rightarrow [0, 1]) \rightarrow (S \rightarrow [0, 1])$ to arbitrarily approach a vector $\{\langle A \rangle \psi_s(s)\}_{s \in S}$ in a game structure with finite state space S with respect to an ω -regular winning objective ψ . Within each step one has to go through $\mathcal{O}(|S|)$ matrix games, and each iteration produces a unique fixed point. The algorithms on safety and reachability objectives are special cases of solving stochastic games [18]. More complex properties

can be expressed as nested fixed points [9]. Therefore, the upper bound complexities become exponential to the size of the winning objectives translated from LTL formulas. More recently, alternative algorithms proposed in [6] prove that for quantitative games with ω -regular winning objectives expressed as parity conditions, whether the values of a game is within $[r - \epsilon, r + \epsilon]$ can be decided in $NP \cap coNP$ for all rational $r \in [0, 1]$ and $\epsilon > 0$, which improves the theoretical upper bound for estimating the optimal values.

It has been shown in [9] that for *safety* games there always exist optimal strategies for the protagonists, however for *reachability* games it is not always the case. We generalise results on the *existence* of optimal strategies for PATL path formulas as follows. Note that the result for unbounded release R has not been studied in the literature, to the authors' knowledge, and the rest of the lemma is derivable from [9].

Lemma 2. *Let s be a state, ψ be a path formula, and A the set of protagonists.*

1. *If ψ is of the form $\bigcirc\phi$, $\phi_1 U^{\leq k} \phi_2$, $\phi_1 R^{\leq k} \phi_2$, or $\phi_1 R \phi_2$ with $k \in \mathbb{N}$, there always exists a joint optimal strategy for A that enforces ψ on s with probability at least $\langle A \rangle \psi(s)$.*
2. *If ψ is of the form $\phi_1 U \phi_2$, there always exists a joint ϵ -optimal strategy for A that enforces ψ on s with probability at least $\langle A \rangle \psi(s) - \epsilon$, for all $\epsilon > 0$.*

The next results prove the existence of a joint A strategy to enforce an PATL path formula with probability greater than α if there exists a joint strategy to enforce that formula with probability greater than α against an optimal \bar{A} strategy. These two lemmas are essential for the proof of the main result (Theorem 1).

Lemma 3. *Let ψ be a PATL path formula and π' be a joint optimal strategy for the antagonists \bar{A} on state s , if there exists a joint strategy π for the protagonists A such that $Pr_{\mathcal{G}(\pi \cup \pi', s)}(\{\rho \in \mathcal{G}(\pi \cup \pi', s) \mid \rho \models \psi\}) > \alpha$, then $\mathcal{G}, s \models \langle\langle A \rangle\rangle^{>\alpha} \psi$.*

Proof. Since π' is the optimal strategy for the antagonists, we have for all joint strategies π'' , $Pr_{\mathcal{G}(\pi'' \cup \pi', s)}(\{\rho \in \mathcal{G}(\pi'' \cup \pi', s) \mid \rho \models \psi\}) \leq \langle A \rangle \psi(s)$, then we have $\langle A \rangle \psi(s) > \alpha$. If there exists an optimal joint strategy for A then we have $s \models \langle\langle A \rangle\rangle^{\geq \langle A \rangle \psi(s)} \psi$, which implies $s \models \langle\langle A \rangle\rangle^{>\alpha} \psi$. Otherwise by Lemma 2 there exists an ϵ -optimal joint strategy for A with small $\epsilon > 0$ to enforce ψ with probability at least $\langle A \rangle \psi(s) - \epsilon > \alpha$. This also gives us $s \models \langle\langle A \rangle\rangle^{>\alpha} \psi$. \square

This result does not hold if we replace the operator “ $>$ ” by “ \geq ” for unbounded until U . This is because if there does not exist a joint optimal strategy for A to enforce $\phi_1 U \phi_2$ with probability $\geq \alpha$, we have no space to insert a tiny $\epsilon > 0$ as we did in the above proof. For the fragment of path formulas without unbounded until, we extend the results for \geq , by the fact that optimal joint strategies for A always exist for these path modalities, as stated by the following lemma.

Lemma 4. *For path formulas ψ in the form of $\bigcirc\phi$ or $\phi_1 U^{\leq k} \phi_2$ and optimal strategies π' of \bar{A} for the antagonists \bar{A} on state s , if there exists a joint strategy π for the protagonists A such that $Pr_{\mathcal{G}(\pi \cup \pi', s)}(\{\rho \in \mathcal{G}(\pi \cup \pi', s) \mid \rho \models \psi\}) \bowtie \alpha$, then $\mathcal{G}, s \models \langle\langle A \rangle\rangle^{\bowtie \alpha} \psi$, where $k \in \mathbb{N}$ and $\bowtie \in \{>, \geq\}$.*

A-PATL

We define a sublogic of PATL by focusing on a particular set of players. Similar to the approach of [3], we only allow negations to appear on the level of propositions. Let $A \subseteq \Sigma$, an A -PATL formula ϕ is a state formula defined as follows:

$$\phi := p \mid \neg p \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \langle\langle A' \rangle\rangle^{\bowtie\alpha} \phi \mid \langle\langle A' \rangle\rangle^{\bowtie\alpha} \phi_1 \mathbf{U}^{\leq k} \phi_2 \mid \langle\langle A' \rangle\rangle^{>\alpha} \phi_1 \mathbf{U} \phi_2$$

where $k \in \mathbb{N}$, $\bowtie \in \{>, \geq\}$ and $A' \subseteq A$. Write \mathbb{L}_A for the set of A -PATL formulas. An A -PATL formula describes a property that players in A are able to ensure with a minimal expectation by their joint strategies. Note that we only allow ' $> \alpha$ ' in the construction of unbounded until.

5 Probabilistic Alternating Simulation Relations

We define probabilistic versions of alternating simulation [3]. An alternating simulation is a two-step simulation. For a sketch, suppose state s is simulated by state t . In the first step the protagonists choose their actions on t to simulate the behavior of the protagonists on s , and in the second step the antagonists choose actions on s to respond to the behavior of the antagonists on t . This somehow results in a simulation-like relation, so that for a certain property the protagonists can enforce on s , they can also enforce it on t . To this end we split Σ into two groups of players — one group of protagonist and the other group of antagonist. Subsequently, we consider only the two-player case in a probabilistic game structure — player \mathbf{I} for the protagonist and player \mathbf{II} for the antagonist, since what we can achieve in the two-player case naturally extends to a result in systems with two complementary sets of players, i.e., $A \cup \bar{A} = \Sigma$. For readability we also write the transition functions as $\delta(s, a_1, a_2)$ and $\bar{\delta}(s, \pi_1, \pi_2)$ for $\delta(s, \langle a_1, a_2 \rangle)$ and $\bar{\delta}(s, \langle \pi_1, \pi_2 \rangle)$, respectively.

Let S, T be two sets and $\mathcal{R} \subseteq S \times T$ be a relation, then $\bar{\mathcal{R}} \subseteq \mathcal{D}(S) \times \mathcal{D}(T)$ is defined by $\Delta \bar{\mathcal{R}} \Theta$ if there exists a weight function $w : S \times T \rightarrow [0, 1]$ satisfying

- $\sum_{t \in T} w(s, t) = \Delta(s)$ for all $s \in S$,
- $\sum_{s \in S} w(s, t) = \Theta(t)$ for all $t \in T$,
- $s \mathcal{R} t$ for all $s \in S$ and $t \in T$ with $w(s, t) > 0$.

Based on the notion of lifting, we define the probabilistic alternating simulation relation for player \mathbf{I} that extends the alternating simulation relation of [3]. The definition for player \mathbf{II} can be made in a similar way.

Definition 1. Consider $\mathcal{G}, \mathcal{G}'$ as two probabilistic game structures. A probabilistic alternating \mathbf{I} -simulation $\sqsubseteq \subseteq S \times S'$ is a relation satisfying if $s \sqsubseteq s'$, then

- $\mathcal{L}(s) = \mathcal{L}'(s')$,
- for all $\pi_1 \in \Pi_{\mathbf{I}}^{\mathcal{G}, 1}$, there exists $\pi'_1 \in \Pi_{\mathbf{I}}^{\mathcal{G}', 1}$, such that for all $\pi'_2 \in \Pi_{\mathbf{II}}^{\mathcal{G}', 1}$, there exists $\pi_2 \in \Pi_{\mathbf{II}}^{\mathcal{G}, 1}$, such that $\bar{\delta}(s, \pi_1, \pi_2) \sqsubseteq \bar{\delta}'(s', \pi'_1, \pi'_2)$.

Note we use level-1 strategies instead of actions (or distributions on actions) on establishing simulations, as in a game structure it is more natural to define simulation in a *behavior*-based way. Also note that a distribution on level-1 strategies yields a level-1 strategy.

Next we propose the notion of probabilistic alternating forward simulation, as per Segala [19], which relates a state to a distribution of states. This requires a different way of lifting. Let $\mathcal{R} \subseteq S \times \mathcal{D}(S)$ be a relation, write $\overline{\mathcal{R}}$ for the smallest relation satisfying $\Delta \overline{\mathcal{R}} \Theta$ if there exists an index set $\{p_i\}_{i \in I}$ satisfying $\sum_{i \in I} p_i = 1$, such that $\Delta = \sum_{i \in I} p_i \cdot \overline{s}_i$, $\Theta = \sum_{i \in I} p_i \cdot \Theta_i$ and $s_i \mathcal{R} \Theta_i$ for all i . Now we define the probabilistic alternating forward simulation relation for player I, and the definition for player II can be made in a similar way.

Definition 2. Consider two probabilistic game structures $\mathcal{G} = \langle S, s_0, \mathcal{L}, \text{Act}, \delta \rangle$ and $\mathcal{G}' = \langle S', s'_0, \mathcal{L}', \text{Act}', \delta' \rangle$. A probabilistic alternating forward I-simulation $\sqsubseteq_f \subseteq S \times \mathcal{D}(S')$ is a relation satisfying if $s \sqsubseteq_f \Delta'$, then

- $\mathcal{L}(s) = \mathcal{L}'(s')$ for all $s' \in [\Delta']$,
- for all $\pi_1 \in \Pi_I^{\mathcal{G},1}$, there exists $\pi'_1 \in \Pi_I^{\mathcal{G}',1}$, such that for all $\pi'_2 \in \Pi_{II}^{\mathcal{G}',1}$, there exists $\pi_2 \in \Pi_{II}^{\mathcal{G},1}$, such that $\delta(s, \pi_1, \pi_2) \overline{\sqsubseteq}_f \overline{\delta'}(\Delta', \pi'_1, \pi'_2)$.

Lemma 5. $s \sqsubseteq t$ implies $s \sqsubseteq_f \bar{t}$.

This lemma says that every probabilistic alternating simulation is a probabilistic alternating forward simulation with a point distribution on the right hand side of the relation. The other way does not hold, i.e., probabilistic alternating forward simulation relates strictly more game structures than probabilistic alternating simulation. In Fig. 2 (which is essentially of [20]), we assume Act_I and Act_{II} are both singleton sets. One may find that there are no states in the set $\{s'_2, s'_3, s'_4, s'_5\}$ in Fig. 2(b) that can simulate states s_3 and s_5 in Fig. 2(a). Therefore, we cannot establish a probabilistic alternating simulation from s_1 to s'_1 . However, s_1 is related to s'_1 by probabilistic alternating forward simulation, since s_3 (s_5) can be related to a uniform distribution over s'_2 and s'_3 (s'_4 and s'_5). The next result shows that the definition of forward simulation also works on the lifted relation.

Lemma 6. If $\Delta \overline{\sqsubseteq}_f \Theta$, then for all $\pi_1 \in \Pi_I^{\mathcal{G},1}$, there exists $\pi_2 \in \Pi_I^{\mathcal{G}',1}$, such that for all $\pi'_2 \in \Pi_{II}^{\mathcal{G}',1}$, there exists $\pi'_1 \in \Pi_{II}^{\mathcal{G},1}$, such that $\overline{\delta}(\Delta, \pi_1, \pi'_1) \overline{\sqsubseteq}_f \overline{\delta'}(\Theta, \pi_2, \pi'_2)$.

Consequently, we are able to show that lifted probabilistic alternating forward simulations are transitive.

Corollary 1. (Transitivity of alternating forward simulation) Let \sqsubseteq_f be a probabilistic alternating forward I-simulation, then $\Delta_1 \overline{\sqsubseteq}_f \Delta_2$ and $\Delta_2 \overline{\sqsubseteq}_f \Delta_3$ implies $\Delta_1 \overline{\sqsubseteq}_f \Delta_3$.

6 Forward I-Simulation is Sound for I-PATL

This section establishes the main result of the paper: a relationship between probabilistic forward I-simulation and I-PATL formulas. Recall that a I-PATL




Fig. 2. An example showing that probabilistic alternating forward simulation is strictly weaker than probabilistic alternating simulation.

formula has only strategy modalities $\langle\langle I \rangle\rangle$ and $\langle\langle \emptyset \rangle\rangle$, and negations are only allowed to appear immediately before the propositions. For readability we write $\langle\langle I \rangle\rangle$ for $\langle\langle \{I\} \rangle\rangle$. Let \mathcal{G} and \mathcal{G}' be two PGSs, $\Delta \in \mathcal{D}(S)$ and $\Delta' \in \mathcal{D}(S')$ such that $\Delta \sqsubseteq_f \Delta'$ by a probabilistic alternating forward I-simulation. We need to show that $\Delta \models \phi$ implies $\Delta' \models \phi$ for all I-PATL formula ϕ .

Our proof relies on the existence of player Π 's optimal strategies for path formulas as winning objectives (as shown in Sect. 4). Suppose π_1 is a I strategy that enforces ϕ , we construct another I strategy π'_1 that simulates π all along the way, in the sense that provided the optimal Π strategy π'_2 there exists another Π strategy π_2 such that the probabilistic execution $\mathcal{E}(\mathcal{G}, \langle \pi_1, \pi_2 \rangle, \Delta)$ will be “simulated” by the probabilistic execution $\mathcal{E}(\mathcal{G}', \langle \pi'_1, \pi'_2 \rangle, \Delta')$. Since π_1 enforces ϕ , $\mathcal{E}(\mathcal{G}, \langle \pi_1, \pi_2 \rangle, \Delta)$ satisfies ϕ , and we show that it is also the case of $\mathcal{E}(\mathcal{G}', \langle \pi'_1, \pi'_2 \rangle, \Delta')$.

Let $\mathcal{E} = \langle E, \Delta, \mathcal{L}^{\mathcal{E}}, \delta^{\mathcal{E}} \rangle$ and $\mathcal{E}' = \langle E', \Delta', \mathcal{L}^{\mathcal{E}'}, \delta^{\mathcal{E}'} \rangle$ be probabilistic executions of \mathcal{G} and \mathcal{G}' , respectively. Also let $\sqsubseteq_f \subseteq S \times \mathcal{D}(S')$ be a probabilistic alternating forward I-simulation. We say the pair $(\mathcal{E}, \mathcal{E}')$ is an *instance of simulation*, by writing $\mathcal{E} \sqsubseteq \mathcal{E}'$, if there exists a (simulation) relation $\sqsubseteq' \subseteq E \times \mathcal{D}(E')$, such that

- $\Delta \sqsubseteq' \Delta'$,
- if $e \sqsubseteq' \theta$ then $last(e) \sqsubseteq_f last(\theta)$,
- if $e \sqsubseteq' \theta$ then $\delta^{\mathcal{E}}(e) \sqsubseteq' \delta^{\mathcal{E}'}(\theta)$,

where $last(\theta)$ is a distribution satisfying $last(\theta)(s) = \sum_{last(e)=s} \theta(e)$. A few properties of the relation \sqsubseteq' are as follows.

Lemma 7. 1. $\Delta \sqsubseteq' \theta$ implies $\delta^{\mathcal{E}}(\Delta) \sqsubseteq' \delta^{\mathcal{E}'}(\theta)$.

2. $\Delta \sqsubseteq' \theta$ and $\Delta = \Delta_1 \oplus_{\alpha} \Delta_2$ with $\alpha \in [0, 1]$, then there exist θ_1, θ_2 such that $\Delta_1 \sqsubseteq' \theta_1$, $\Delta_2 \sqsubseteq' \theta_2$, and $\theta = \theta_1 \oplus_{\alpha} \theta_2$.

Let Δ be a state distribution of \mathcal{G} , Δ' be a state distribution of \mathcal{G}' , and $\Delta \sqsubseteq_f \Delta'$. Suppose π_1 is a I strategy in \mathcal{G} that enforces ϕ with probability at least α , and π'_2

is a Π strategy in \mathcal{G}' , step-by-step we establish a I strategy π'_1 and a Π strategy π_2 , so that the probabilistic execution decided by π_1 and π_2 from Δ will be simulated by the probabilistic execution decided by π'_1 and π'_2 from Δ' .

Lemma 8. *Let $\mathcal{G} = \langle S, s_0, \mathcal{L}, \text{Act}, \delta \rangle$ and $\mathcal{G}' = \langle S', s'_0, \mathcal{L}', \text{Act}', \delta' \rangle$ be two PGSSs. If $\Delta \sqsubseteq_f \Delta'$, then for all $\pi_1 \in \Pi_I^{\mathcal{G}}$ and $\pi'_2 \in \Pi_{II}^{\mathcal{G}'}$, there exists $\pi'_1 \in \Pi_I^{\mathcal{G}'}$ and $\pi_2 \in \Pi_{II}^{\mathcal{G}}$, such that $\mathcal{E}(\mathcal{G}, \langle \pi_1, \pi_2 \rangle, \Delta) \sqsubseteq \mathcal{E}'(\mathcal{G}', \langle \pi'_1, \pi'_2 \rangle, \Delta')$.*

In order to measure the probability of a path formula to be satisfied when the strategies from both player I and player II are fixed, we define a relation $\models^{\bowtie \alpha}$ for probabilistic executions.

Definition 3. *Let \mathcal{G} be a probabilistic game structure, $\mathcal{E}(\Delta) = \langle E, \Delta, \mathcal{L}^{\mathcal{E}}, \delta^{\mathcal{E}} \rangle$ a probabilistic execution determined by a vector $\pi_{\mathcal{E}}$, and ψ a path formula, define*

$$\mathcal{E}(\Delta) \models^{\bowtie \alpha} \psi \quad \text{iff} \quad Pr_{\mathcal{E}}^{\Delta}(\{\rho \in \bigcup_{s \in [\Delta]} \mathcal{G}(\pi_{\mathcal{E}}, s) \mid \rho \models \psi\}) \bowtie \alpha$$

It is conceivable that in a probabilistic execution every finite or infinite trace in $E^* \cup E^{\omega}$ maps to a trace in \mathcal{G} , in the way that $\rho = e_1 e_2 e_3 \dots$ is a trace in \mathcal{E} implies that $proj(\rho) = last(e_1)last(e_2)last(e_3) \dots$ is a play in \mathcal{G} , where the function $proj$ projects every finite sequence of states in E into its last state in S . Consequently, we let $Pr_{\mathcal{E}}^{\Delta}$ be a probabilistic measure over E^{ω} , such that for the cone sets (of finite traces), we have $Pr_{\mathcal{E}}^{\Delta}(e) = \Delta(last(e))$, and $Pr_{\mathcal{E}}^{\Delta}(\gamma \cdot e_1 \cdot e_2) = Pr_{\mathcal{E}}^{\Delta}(\gamma \cdot e_1) \cdot \delta^{\mathcal{E}}(e_1)(e_2)$, for $\gamma \in E^*$ and $e_1, e_2 \in E$. Let ρ be an infinite trace in \mathcal{E} , we write $\rho \models \psi$ iff $proj(\rho) \models \psi$. Similarly, for a state formula ϕ and $e \in E$, write $e \in \llbracket \phi \rrbracket$ iff $last(e) \in \llbracket \phi \rrbracket$. In the following we study the properties of the satisfaction relation for a probabilistic execution to satisfy a I-PATL path formula by means of unfolding.

Lemma 9. *Let ϕ, ϕ_1 and ϕ_2 be I-PATL (state) formulas, and $\bowtie \in \{>, \geq\}$ then*

1. $\mathcal{E}(\Delta) \models^{\bowtie \alpha} \bigcirc \phi$ iff there exists $\alpha' \bowtie \alpha$, such that $\delta^{\mathcal{E}}(\Delta) = \Delta_1 \oplus_{\alpha'} \Delta_2$ with $[\Delta_1] \cap [\Delta_2] = \emptyset$, and $\Delta_1 \models \phi$.
2. $\mathcal{E}(\Delta) \models^{\bowtie \alpha} \phi_1 \mathcal{U}^{\leq k} \phi_2$ iff there exists a finite sequence of triples $\{ \langle (\Delta_{i,0}, \alpha_{i,0}), (\Delta_{i,1}, \alpha_{i,1}), (\Delta_{i,2}, \alpha_{i,2}) \rangle \}_{0 \leq i \leq j}$ for some $j \leq k$, with $[\Delta_{i,\ell}] \cap [\Delta_{i,\ell'}] = \emptyset$ for all distinct $\ell, \ell' \in \{0, 1, 2\}$ and $0 \leq i \leq j$, such that

$$(1) \quad \sum_{i \in [0..j]} \left(\alpha_{i,1} \cdot \prod_{i' \in [0..i-1]} \alpha_{i',0} \right) \bowtie \alpha,$$

(2) $\Delta = \sum_{\ell \in \{0,1,2\}} \alpha_{0,\ell} \cdot \Delta_{0,\ell}$, and $\delta^{\mathcal{E}}(\Delta_{i,0}) = \sum_{\ell \in \{0,1,2\}} \alpha_{i+1,\ell} \cdot \Delta_{i+1,\ell}$ for all $0 \leq i < j$, (3) $\Delta_{i,0} \models \phi_1$ and $\Delta_{i,1} \models \phi_2$ for all $0 \leq i \leq j$.

3. $\mathcal{E}(\Delta) \models^{\bowtie \alpha} \phi_1 \mathcal{U} \phi_2$ iff there exists a finite or infinite sequence of triples $\{ \langle (\Delta_{i,0}, \alpha_{i,0}), (\Delta_{i,1}, \alpha_{i,1}), (\Delta_{i,2}, \alpha_{i,2}) \rangle \}_{0 \leq i < j}$ for some $j \in \mathbb{N}^+ \cup \{\infty\}$, with $[\Delta_{i,\ell}] \cap [\Delta_{i,\ell'}] = \emptyset$ for all distinct $\ell, \ell' \in \{0, 1, 2\}$ and $0 \leq i < j$, such that

$$(1) \quad \sum_{0 \leq i < j} \left(\alpha_{i,1} \cdot \prod_{i' \in [0..i-1]} \alpha_{i',0} \right) \bowtie \alpha,$$

(2) $\Delta = \sum_{\ell \in \{0,1,2\}} \alpha_{0,\ell} \cdot \Delta_{0,\ell}$, and $\delta^{\mathcal{E}}(\Delta_{i,0}) = \sum_{\ell \in \{0,1,2\}} \alpha_{i+1,\ell} \cdot \Delta_{i+1,\ell}$ for all $0 \leq i < j$, (3) $\Delta_{i,0} \models \phi_1$ and $\Delta_{i,1} \models \phi_2$ for all $0 \leq i < j$.

Theorem 1. *Let $\mathcal{G} = \langle S, s_0, \mathcal{L}, \text{Act}, \delta \rangle$ and $\mathcal{G}' = \langle S', s'_0, \mathcal{L}', \text{Act}', \delta' \rangle$ be two PGSSs, $\sqsubseteq_f \subseteq S \times \mathcal{D}(S')$ a probabilistic alternating forward I-simulation. If $\Delta \sqsubseteq_f \Delta'$, then $\mathcal{G}, \Delta \models \phi$ implies $\mathcal{G}', \Delta' \models \phi$ for all $\phi \in \mathbb{L}_{\text{I}}$.*

Proof. (sketch) We prove by induction on the structure of a I-PATL formula ϕ . Base case: suppose $\Delta \models p$, then $s \models p$ for all $s \in [\Delta]$. By $\Delta \sqsubseteq_f \Delta'$, there exists an index set $\{q_i\}_{i \in I}$ satisfying $\sum_{i \in I} q_i = 1$, $\Delta = \sum_{i \in I} q_i \bar{s}_i$, $\Delta' = \sum_{i \in I} q_i \Delta_i$, and $s_i \sqsubseteq_f \Delta_i$. Therefore $\mathcal{L}(s_i) = \mathcal{L}'(t)$ for all $t \in [\Delta_i]$. So $t \models p$ for all $t \in [\Delta_i]$ for all i . Therefore $\Delta' \models p$. The case of $\neg p$ is similar.

We show the case when $\phi = \langle\langle \text{I} \rangle\rangle^{>\alpha} \phi_1 \text{U} \phi_2$, and the proof methods for the other PATL path constructors are similar. Since for all $t \in [\Delta']$ there exists an optimal strategy π^t for the winning objective $\neg \phi_1 \mathcal{R} \neg \phi_2$ by Lemma 2(1), and we combine these strategies into a single strategy π'_2 satisfying $\pi'_2(t \cdot \alpha) = \pi^t(t \cdot \alpha)$ for all $t \in [\Delta']$ and $\alpha \in S^*$. Then π'_2 is optimal for $\neg \phi_1 \mathcal{R} \neg \phi_2$ on Δ' . Then by Lemma 8, there exist $\pi_2 \in \Pi_{\Pi}^{\mathcal{G}}$ and $\pi'_1 \in \Pi_{\text{I}}^{\mathcal{G}'}$ such that $\mathcal{E}(\mathcal{G}, \langle \pi_1, \pi_2 \rangle, \Delta) \sqsubseteq \mathcal{E}'(\mathcal{G}', \langle \pi'_1, \pi'_2 \rangle, \Delta')$. Since π_1 enforces $\phi_1 \text{U} \phi_2$ with probability greater than α , we have $\mathcal{E}(\Delta) \models^{>\alpha} \phi_1 \text{U} \phi_2$. By Lemma 9(3) there exists a sequence of triples $\{(\Delta_{i,0}, \alpha_{i,0}), (\Delta_{i,1}, \alpha_{i,1}), (\Delta_{i,2}, \alpha_{i,2})\}_{0 \leq i < j}$ for some $j \in \mathbb{N}^+ \cup \{\infty\}$ satisfying the properties as stated in Lemma 9(3). By repetitively applying Lemma 7 we establish another sequence of triples $\{(\Delta'_{i,0}, \alpha_{i,0}), (\Delta'_{i,1}, \alpha_{i,1}), (\Delta'_{i,2}, \alpha_{i,2})\}_{0 \leq i < j}$, such that (1) $\sum_{0 \leq i < j} (\alpha_{i,1} \cdot \prod_{i' \in [0..i-1]} \alpha_{i',0}) > \alpha$, (2) $\Delta' = \sum_{\ell \in \{0,1,2\}} \alpha_{0,\ell} \cdot \Delta'_{0,\ell}$, and $\delta^{\mathcal{E}}(\Delta'_{i,0}) = \sum_{\ell \in \{0,1,2\}} \alpha_{i+1,\ell} \cdot \Delta'_{i+1,\ell}$ for all $0 \leq i < j$, (3) $\Delta_{i,0} \sqsubseteq_f \Delta'_{i,0}$ and $\Delta_{i,1} \sqsubseteq_f \Delta'_{i,1}$ for all $0 \leq i < j$. By induction hypothesis we have $\Delta'_{i,0} \models \phi_1$ and $\Delta'_{i,1} \models \phi_2$ for all $0 \leq i < j$. Therefore $\mathcal{E}(\Delta') \models^{>\alpha} \phi_1 \text{U} \phi_2$ by Lemma 9(3). Since π'_2 is an optimal strategy of Π , we have $\Delta' \models \langle\langle \text{I} \rangle\rangle^{>\alpha} \phi_1 \text{U} \phi_2$ by Lemma 3.

For a formula $\langle\langle \emptyset \rangle\rangle^{\bowtie \alpha} \psi$ we apply the same proof strategies as for $\langle\langle \text{I} \rangle\rangle^{\bowtie \alpha} \psi$, except that player I does not need to enforce ψ with a certain probability $\bowtie \alpha$ since every probabilistic execution generated by a pair of I and Π strategies will enforce ψ with that probability. \square

7 Probabilistic Alternating Bisimulation

If a probabilistic alternating simulation is symmetric, we call it a probabilistic alternating bisimulation.

Definition 4. *Consider two probabilistic game structures $\mathcal{G} = \langle S, s_0, \mathcal{L}, \text{Act}, \delta \rangle$ and $\mathcal{G}' = \langle S', s'_0, \mathcal{L}', \text{Act}', \delta' \rangle$. A probabilistic alternating I-bisimulation $\simeq_{\subseteq} \subseteq S \times S'$ is a symmetric relation satisfying if $s \simeq s'$, then*

- $\mathcal{L}(s) = \mathcal{L}'(s')$,
- for all $\pi_1 \in \Pi_{\text{I}}^{\mathcal{G},1}$, there exists $\pi'_1 \in \Pi_{\text{I}}^{\mathcal{G}',1}$, such that for all $\pi'_2 \in \Pi_{\Pi}^{\mathcal{G}',1}$, there exists $\pi_2 \in \Pi_{\Pi}^{\mathcal{G},1}$, such that $\delta(s, \pi_1, \pi_2) \simeq \delta'(s', \pi'_1, \pi'_2)$,

where \cong is a lifting of \simeq by weight functions.

Since every probabilistic alternating I-simulation is also a probabilistic alternating forward I-simulation by treating the right hand side state as a point distribution (Lemma 5), the lifted probabilistic alternating I-simulation is also a lifted probabilistic alternating forward I-simulation. This fact extends for bisimulation. A probabilistic alternating I-bisimulation also preserves formulas in \mathbb{L}_I . Moreover, we write \mathbb{L}_I^+ for the set of formulas defined as follows, which allows negations to appear anywhere in a formula, and further we are able to show that probabilistic alternating bisimulation preserves all properties expressed in \mathbb{L}_I^+ .

$$\phi := p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \langle\langle A' \rangle\rangle^{\bowtie\alpha} \phi \mid \langle\langle A' \rangle\rangle^{\bowtie\alpha} \phi_1 \mathbf{U}^{\leq k} \phi_2 \mid \langle\langle A' \rangle\rangle^{>\alpha} \phi_1 \mathbf{U} \phi_2$$

Theorem 2. *Let $\mathcal{G} = \langle S, s_0, \mathcal{L}, Act, \delta \rangle$ and $\mathcal{G}' = \langle S', s'_0, \mathcal{L}', Act', \delta' \rangle$ be two PGSs, $\simeq \subseteq S \times S'$ is a probabilistic alternating I-bisimulation. For all $s \in S$ and $s' \in S'$ with $s \simeq s'$ and $\phi \in \mathbb{L}_I^+$, we have $\mathcal{G}, s \models \phi$ iff $\mathcal{G}', s' \models \phi$.*

The proof methodology basically follows that of Theorem 1, besides that whenever $\Delta \cong \Delta'$ and $\Delta \models \neg\phi$, we show that if there were $s' \in [\Delta]'$ such that $\mathcal{G}', s' \models \phi$ then we would also have $\mathcal{G}, s \models \phi$ for some $s \in [\Delta]$, which is a contradiction. And from that we have $\Delta' \models \neg\phi$ as well.

8 Conclusion and Future Work

We report our first results on probabilistic alternating simulation relations. We have introduced two notions of simulation for probabilistic game structures — probabilistic alternating simulation and probabilistic alternating forward simulation, following the seminal works of Segala and Lynch [19, 21] on probabilistic simulation relations and the work of Alur et al. [3] on alternating refinement relations for non-probabilistic game structures. Our main effort has been devoted to a logical characterization for probabilistic alternating simulation relations, by showing that they preserve a fragment of PATL formulas. It is worth noting that on our way to the main result, we find that the proof strategy accommodated in [3] no longer applies, due to the failure in reconstructing a strategy from sub-strategies when the system transitions become probabilistic. We circumvent this problem by incorporating the results of probabilistic determinacy [14] and the existence of optimal strategies [9] in stochastic games. A full version of the paper is available as a technical report [23].

There are several ways to proceed. We want to study the *completeness* of logical characterization for probabilistic alternating forward simulation. It is also of our interest to investigate the *complexity* for checking probabilistic alternating simulation relations by studying the results in the literature [3, 5]. Our work was partially motivated by the paper [4], where PATL is used to formalize a *balanced* property for a probabilistic contract signing protocol. Here, a balanced protocol means that a dishonest participant never has a strategy to unilaterally determine the outcome of the protocol. It is interesting to see how much the development of simulation relations for probabilistic game structures can help the verification of such kind of security protocols.

References

1. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. In: *Proc. FOCS*, IEEE CS (1997) 100–109
2. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5) (2002) 672–713
3. R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In *Proc. CONCUR*, LNCS 1466., Springer (1998) 163–178
4. M. Aizatulin, H. Schnoor, and T. Wilke. Computationally sound analysis of a probabilistic contract signing protocol. In *Proc. ESORICS*, LNCS 5789., Springer (2009) 571–586
5. C. Baier, B. Engelen, and M. E. Majster-Cederbaum. Deciding bisimilarity and similarity for probabilistic processes. *J. Comput. Syst. Sci.*, 60(1) (2000) 187–231
6. K. Chatterjee, L. de Alfaro, and T. A. Henzinger. The complexity of quantitative concurrent parity games. In *Proc. SODA*, ACM (2006) 678–687
7. T. Chen and J. Lu. Probabilistic alternating-time temporal logic and model checking algorithm. In *Proc. FSKD*, IEEE CS (2007) 35–39
8. L. de Alfaro, T. A. Henzinger, and O. Kupferman. Concurrent reachability games. In *Proc. FOCS*, IEEE CS (1998) 564–575
9. L. de Alfaro and R. Majumdar. Quantitative solution of omega-regular games. *J. Comput. Syst. Sci.*, 68(2) (2004) 374–397
10. J. Desharnais, V. Gupta, R. Jagadeesan, and P. Panangaden. Weak bisimulation is sound and complete for PCTL*. In *Proc. CONCUR*, LNCS 2421., Springer (2002) 355–370
11. E. A. Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science (B)*, MIT Press (1990) 955–1072
12. H. Hansson. *Time and Probability in Formal Design of Distributed Systems*. Elsevier, 1994.
13. N. A. Lynch, R. Segala, and F. W. Vaandrager. Observing branching structure through probabilistic contexts. *SIAM J. Comput.*, 37(4) (2007) 977–1013
14. D. A. Martin. The determinacy of Blackwell games. *J. Symb. Log.*, 63(4) (1998) 1565–1581
15. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
16. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1947.
17. A. Parma and R. Segala. Logical characterizations of bisimulations for discrete probabilistic systems. In *Proc. FoSSaCS*, LNCS 4423., Springer (2007) 287–301
18. T. E. S. Raghavan and J. A. Filar. Algorithms for stochastic games – A survey. *Mathematical Methods of Operations Research*, 35(6) (1991) 437–472
19. R. Segala. A compositional trace-based semantics for probabilistic automata. In *Proc. CONCUR*, LNCS 962., Springer (1995) 234–248
20. R. Segala. *Modeling and Verification of Randomized Distributed Real-Time Systems*. PhD thesis, MIT, 1995.
21. R. Segala and N. A. Lynch. Probabilistic simulations for probabilistic processes. *Nord. J. Comput.*, 2(2) (1995) 250–273
22. W. Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science (B)*, MIT Press (1990) 133–192
23. C. Zhang and J. Pang. On probabilistic alternating simulations. Tech. Rep. (2010) <http://arxiv.org/abs/1003.0788>