Characterising Probabilistic Alternating Simulation for Concurrent Games

Chenyi Zhang
College of Information Science and Technology
Jinan University
Guangzhou, China
chenyi_zhang@jnu.edu.cn

Jun Pang
Faculty of Science, Technology and Medicine &
Interdisciplinary Centre for Security, Reliability and Trust
University of Luxembourg, Luxembourg
jun.pang@uni.lu

***Abstract*—Probabilistic game structures combine both nondeterminism and stochasticity, where players repeatedly take actions simultaneously to move to the next state of the concurrent game. Probabilistic alternating simulation is an important tool to compare the behaviour of different probabilistic game structures. In this paper, we present a sound and complete modal characterisation of this simulation relation by proposing a new logic based on probability distributions. The logic enables a player to enforce a property in the next state or distribution. Its extension with fixpoints, which also characterises the simulation relation, can express a lot of interesting properties in practical applications.**

***Index Terms*—Concurrent game, Probabilistic alternating simulation, Logic characterisation**

I. Introduction

Simulation relations and bisimulation relations [1] are important research topics in concurrency theory. In the classical model of labelled transition systems (LTS), simulation and bisimulation have been proved useful for comparing the behaviour of concurrent systems. The modal characterisation problem has been studied both in classical and in probabilistic systems, i.e., the Hennessy-Milner logic (HML) [2] that characterises image-finite LTS, and various modal logics have been proposed to characterise strong and weak probabilistic (bi)simulation in the model of probabilistic automata [3]–[5]. To study multi-player games, the concurrent game structure (GS) [6] is a model that defines a system that evolves while interacting with outside *players*. As a player's behaviour is not fully specified within a system, GS are often also known as *open* systems. Alternating simulation (A-simulation) is defined in GS focusing on players' ability to enforce temporal properties specified in alternating-time temporal logic (ATL) [6], and A-simulation is shown to be sound and complete to a fragment of ATL [7].

In this paper, we work on the model of probabilistic game structure (PGS) which has probabilistic transitions. PGS also allows probabilistic (or mixed) choices of players. In this setting, we assume all players have complete information about system states. The simulation relation in PGS, called probabilistic alternating simulation (PA-simulation), has been shown to preserve a fragment of probabilistic alternating-time temporal logic (PATL) under *mixed strategies* [8]. Given the classical results of modal characterisations for (non-probabilistic) LTS, probabilistic automata, as well as for (non-probabilistic) game structures, we investigate if a similar correspondence exists for processes and modal logics in the domain of concurrent games with probabilistic transitions and mixed strategies. We find that such a correspondence still holds by adapting a modal logic with nondeterministic distributions extended from the work of [4]. As a by-product, we extend that modal logic with fixpoint operators and study its expressiveness. Notably, similar to the fixpoint logics in [9], [10], the least fixpoint modality in our logic only expresses finite reachability, a property in line with the original μ-calculus [11], which somehow in the probabilistic setting may not be powerful enough to express certain reachability properties that require infinite accumulation of moves in a play.

Contributions. This paper studies modal characterisation of the probabilistic alternation simulation relation in probabilistic concurrent game structures, which defines a novel modal logic based on probability distributions. This new logic expresses a player's power to enforce a property in the next state or distribution. The logic also incorporates both probabilistic and nondeterminstic features that need to be considered during the two-player interplay. The second contribution is the introduction of a fixpoint logic, which also characterises the simulation relation, extended from that modal logic. The expressive power of the logic has been illustrated by examples.

Structure of the paper. We present some necessary preliminaries in Section II, and give the formal definition of probabilistic alternating simulation in Section III. A modal logic with nondeterministic distributions is introduced in Section IV, and we show that it has a correct characterisation of probabilistic alternating simulation for concurrent games. This logic is then extended with variables and fixpoint operators, resulting in a Probabilistic Alternating-time μ-calculus (PAMu), and the logic characterisation of probabilistic alternating simulation can be extended to PAMu as well. We discuss related work in Section VI and conclude the paper with possible future work in Section VII.

II. Preliminaries

A *discrete probabilistic distribution* Δ over a set S is a function of type $S \rightarrow [0,1]$, where $\sum_{s \in S} \Delta(s) = 1$. We write $\mathcal{D}(S)$ for the set of all such distributions, ranged over by

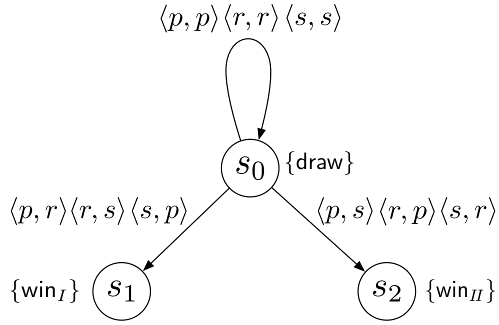

Fig. 1. The PGS for the repeated rock-paper-scissors game.

symbols $\Delta, \Theta, \ldots$. Given a set $T \subseteq S$, $\Delta(T) = \sum_{s \in T} \Delta(s)$, i.e., the probability for the given set T. Given an index set I, a list of distributions $\langle \Delta_i \rangle_{i \in I}$ and a list of values $\langle p_i \rangle_{i \in I}$ where $p_i \in [0,1]$ for all $i \in I$ and $\sum_{i \in I} p_i = 1$, we have that $\sum_{i \in I} p_i \Delta_i$ is also a distribution. If $|I| = 2$ we may also write $\Delta_1 \oplus_\alpha \Delta_2$ for the distribution $p_1 \Delta_1 + p_2 \Delta_2$ where $p_1 = \alpha$ and $p_2 = 1 - \alpha$. For $s \in S$, $s_\star$ represents a *point (or Dirac) distribution* satisfying $s_\star(s) = 1$ and $s_\star(t) = 0$ for all $t \neq s$. Given $\Delta \in \mathcal{D}(S)$, we define $\lceil \Delta \rceil$ as the set $\{s \in S \mid \Delta(s) > 0\}$, which is the *support* of Δ.

We work on a model called probabilistic (concurrent) game structure (PGS) with two players $\mathcal{I}$ and $\mathcal{II}$ (though we believe our results can be straightforwardly extended to handle a finite set of players as in the standard concurrent game structures [6]). Each player has complete information about the PGS at any time during a play. Let $\texttt{Prop}$ be a finite set of propositions.

Definition 1: A probabilistic game structure (PGS) $\mathcal{G}$ is a tuple $\langle S, s_0, L, \texttt{Act}, \delta \rangle$, where

- S is a finite set of states, with s_0 the initial state;
- $L : S \to 2^{\texttt{Prop}}$ is the labelling function which assigns to each state $s \in S$ a set of propositions true in s;
- $\texttt{Act} = \texttt{Act}_{\mathcal{I}} \times \texttt{Act}_{\mathcal{II}}$ is a finite set of joint actions, where $\texttt{Act}_{\mathcal{I}}$ and $\texttt{Act}_{\mathcal{II}}$ are, respectively, the sets of actions for players $\mathcal{I}$ and $\mathcal{II}$;
- $\delta : S \times \texttt{Act} \to \mathcal{D}(S)$ is a transition function.

If in state s player $\mathcal{I}$ performs action a_1 and player $\mathcal{II}$ performs action a_2, then $\delta(s, \langle a_1, a_2 \rangle)$ is the distribution for the next states. During each step the players choose their next moves simultaneously.

Example 1: Figure 1 presents the PGS of two players repeatedly playing the rock-paper-scissors game.[1] It has three states s_0, s_1, and s_2, with s_0 being the initial state. Each state is labelled with an atomic proposition indicating the result of a round of the game (which player wins or there is a draw). For instance, in state s_1 player $\mathcal{I}$ wins the game. Actions of the players are r (representing playing rock), p (representing playing paper), s (representing playing scissors). The joint actions $\langle a_1, a_2 \rangle$ with $a_1, a_2 \in \{r, p, s\}$ are depicted along with the transitions. The function δ describes the transition function as shown in Figure 1. The winning states s_1 and s_2 are absorbing, i.e., all actions from there make self-transitions, and the game effectively terminates there.

We define a *mixed action* of player i as a function from states to distributions on $\texttt{Act}_i$, ranged over by $\pi, \pi_1, \sigma \ldots$, and write Π_i for the set of mixed actions from player i. In particular, $a_\star$ is a *deterministic* mixed action which always chooses a with probability 1 in all states.

Example 2: In the rock-paper-scissors game (see Figure 1), for both players, the mixed action with probability $\frac{1}{3}$ for each of the actions (r, p and s) seems optimal, in the sense that it guarantees the chance to eventually win the game with probability at least $\frac{1}{2}$.

We generalise the transition function δ by $\widetilde{\delta}$ to handle mixed actions. Given $\pi_1 \in \Pi_{\mathcal{I}}$ and $\pi_2 \in \Pi_{\mathcal{II}}$, for all $s, t \in S$, we have $\widetilde{\delta}(s, \langle \pi_1, \pi_2 \rangle)(t)$

$$= \sum_{a_1 \in \texttt{Act}_{\mathcal{I}}, a_2 \in \texttt{Act}_{\mathcal{II}}} \pi_1(s)(a_1) \cdot \pi_2(s)(a_2) \cdot \delta(s, \langle a_1, a_2 \rangle)(t).$$

We further extend the function $\widetilde{\delta}$ to handle transitions from distributions to distributions. Formally, given a distribution $\Delta \in \mathcal{D}(S)$, $\pi_1 \in \Pi_{\mathcal{I}}$ and $\pi_2 \in \Pi_{\mathcal{II}}$, for all $s \in S$, we have $\widetilde{\delta}(\Delta, \langle \pi_1, \pi_2 \rangle)(s) = \sum_{t \in \lceil \Delta \rceil} \Delta(t) \cdot \widetilde{\delta}(t, \langle \pi_1, \pi_2 \rangle)(s)$. For better readability, sometimes we write $\Delta \xrightarrow{\pi_1, \pi_2} \Theta$ if $\Theta = \widetilde{\delta}(\Delta, \langle \pi_1, \pi_2 \rangle)$.

Let $\leq \subseteq S \times S$ be a partial order, define $\leq_{Sm} \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$, by $P \leq_{Sm} Q$ if for all $t \in Q$ there exists $s \in P$ such that $s \leq t$. In the literature this definition is known as the 'Smyth order' [13], [14] regarding '$\leq$'.

Relations in probabilistic systems usually require a notion of *lifting* [15], which extends the relations to the domain of distributions.[2] Let S, T be two sets and $\mathcal{R} \subseteq S \times T$ be a relation, then $\overline{\mathcal{R}} \subseteq \mathcal{D}(S) \times \mathcal{D}(T)$ is a *lifted relation* defined by $\Delta \, \overline{\mathcal{R}} \, \Theta$ if there exists a *weight* function $w : S \times T \to [0,1]$ such that

- $\sum_{t \in T} w(s,t) = \Delta(s)$ 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 $w(s,t) > 0$.

The intuition behind the lifting is that each state in the support of one distribution may correspond to a number of states in the support of the other distribution, and vice versa. In the following section, we extend the notion of alternating simulation [7] to a probabilistic setting in the way of lifting. The next example is taken from [16] which shows how to lift a relation.

Example 3: In Figure 2, we have two sets of states $S = \{s_1, s_2\}$ and $T = \{t_1, t_2, t_3\}$, and a relation $\mathcal{R} = \{(s_1, t_1), (s_1, t_2), (s_2, t_2), (s_2, t_3)\}$. Suppose $\Delta(s_1) = \Delta(s_2) = \frac{1}{2}$ and $\Theta(t_1) = \Theta(t_2) = \Theta(t_3) = \frac{1}{3}$, we may establish $\Delta \, \overline{\mathcal{R}} \, \Theta$. To check this, we define a weight function w by: $w(s_1, t_1) = \frac{1}{3}$, $w(s_1, t_2) = w(s_2, t_2) = \frac{1}{6}$, and

[1] A similar example was used in [12]. A concurrent stochastic structure (CSG), as defined in [12], is a PGS, in the sense that all the probability distributions involved in the CSG are point distributions.

[2] In a probabilistic system without explicit user interactions, state s is simulated by state t if for every $s \xrightarrow{a} \Delta_1$ there exists $t \xrightarrow{a} \Delta_2$ such that Δ_1 is simulated by Δ_2.

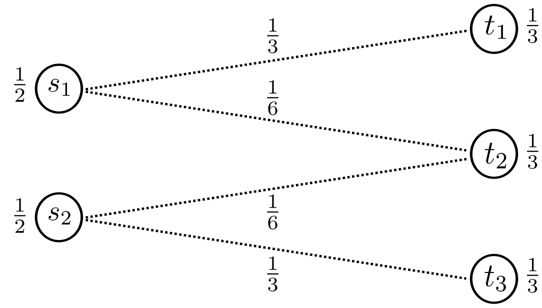

Fig. 2. An example showing how to lift one relation.

$w(s_2, t_3) = \frac{1}{3}$. The dotted lines in the graph indicate the allocation of weights that is required to relate Δ to Θ via $\overline{\mathcal{R}}$.

We present some properties of lifted relations. First we show that, by combining distributions that are lift-related with the same weight on both sides, we get the resulting distributions lift-related. The proofs of the following Lemmas regarding lifting are straightforward, which may be found in the literature (e.g. [17], [18]).

Lemma 1: Let $\mathcal{R} \subseteq S \times S'$ and $\langle p_i \rangle_{i \in I}$ be a list of values satisfying $\sum_{i \in I} p_i = 1$, and $\Delta_i \, \overline{\mathcal{R}} \, \Delta'_i$ for $\Delta_i \in \mathcal{D}(S)$ and $\Delta'_i \in \mathcal{D}(S')$ for all i, then $\sum_{i \in I} p_i \Delta_i \, \overline{\mathcal{R}} \, \sum_{i \in I} p_i \Delta'_i$.

Lemma 2 states that, given two related distributions, if we split a distribution on one side, then there exists a split on the other side by the same index set, so that the corresponding (sub)distributions are related by the lifted relation.

Lemma 2: Let $\Delta \in \mathcal{D}(S)$, $\Delta' \in \mathcal{D}(S')$, $\mathcal{R} \subseteq S \times S'$, $\langle p_i \rangle_{i \in I}$ be a list of values satisfying $\sum_{i \in I} p_i = 1$. If $\Delta \, \overline{\mathcal{R}} \, \Delta'$, then

1) for all lists of distributions $\langle \Delta_i \rangle_{i \in I}$ with $\Delta_i \in \mathcal{D}(S)$ for all $i \in I$, satisfying $\Delta = \sum_{i \in I} p_i \Delta_i$, there exists $\langle \Delta'_i \rangle_{i \in I}$ with $\Delta'_i \in \mathcal{D}(S')$ such that $\Delta' = \sum_{i \in I} p_i \Delta'_i$ and $\Delta_i \, \overline{\mathcal{R}} \, \Delta'_i$ for all $i \in I$;
2) for all lists of distributions $\langle \Delta'_i \rangle_{i \in I}$ with $\Delta'_i \in \mathcal{D}(S')$ for all $i \in I$, satisfying $\Delta' = \sum_{i \in I} p_i \Delta'_i$, there exists $\langle \Delta_i \rangle_{i \in I}$ with $\Delta_i \in \mathcal{D}(S)$ such that $\Delta = \sum_{i \in I} p_i \Delta_i$, and $\Delta_i \, \overline{\mathcal{R}} \, \Delta'_i$ for all $i \in I$.

III. Probabilistic Alternating Simulation

In concurrency models, simulation is used to relate states with respect to their behaviours. For example, in a labelled transition system (LTS) $\langle S, A, \rightarrow \rangle$, where S is a set of states, A is a set of actions and $\rightarrow \subseteq S \times A \times S$ is the transition relation, we say state s is simulated by state t, written $s \leq t$, if for every $s \xrightarrow{a} s'$ there exists $t \xrightarrow{a} t'$ such that $s' \leq t'$. In this coinductive definition, state t is able to simulate state s by performing the same action a, with their destination states still related. Simulation is a useful tool in abstraction and refinement based verification, as informally, in the above case, t contains at least as much behaviour as s does. If the relation '$\leq$' is symmetric, then it is also a bisimulation.

In two-player non-probabilistic game structures (GS), alternating simulation (A-simulation) is used to describe a player's ability to enforce certain temporal requirements (regardless of the other player's behaviours) [7]. In this paper we only focus on the ability of player $\mathcal{I}$ in a two-player game. Since in a game structure a transition requires the participation of both parties, fixing player $\mathcal{I}$'s input leaves a set of possible next states depending on player $\mathcal{II}$'s inputs. A-simulation is defined in the model of non-probabilistic game structures $\langle S, s_0, \texttt{Act}, L, \delta \rangle$, which has a set of states S with $s_0 \in S$ the initial state, $\texttt{Act} = \texttt{Act}_{\mathcal{I}} \times \texttt{Act}_{\mathcal{II}}$ the set of actions from players $\mathcal{I}$ and $\mathcal{II}$, L the labelling function and $\delta : S \times \texttt{Act} \rightarrow S$ the transition function. An A-simulation $\leq^A \subseteq S \times S$ is defined as follows. Let $s, t \in S$, s is A-simulated by t, i.e., $s \leq^A t$, if

- $L(s) = L(t)$, and
- for all $a \in \texttt{Act}_{\mathcal{I}}$ there exists $a' \in \texttt{Act}_{\mathcal{I}}$ such that $\delta(s, a) \leq^A_{Sm} \delta(t, a')$, where $\delta(s, a)$ is the "curried" transition function defined by $\{s' \in S \mid \exists b \in \texttt{Act}_{\mathcal{II}} : \delta(s, \langle a, b \rangle) = s'\}$.

Intuitively, on state t action a' enforces a more restrictive set than action a enforces on state s, as shown by the Smyth-ordered relation $\leq^A_{Sm}$: for every $b' \in \texttt{Act}_{\mathcal{II}}$ there exists $b \in \texttt{Act}_{\mathcal{II}}$ such that $\delta(s, \langle a, b \rangle) \leq^A \delta(t, \langle a', b' \rangle)$.

Zhang and Pang extend A-simulation to probabilistic alternating simulation (PA-simulation) in PGS [8] and propose an algorithm for computing the largest PA-simulation [19]. Their definition requires lifting of the simulation relation to derive a relation on distributions of states.

Definition 2: Given a PGS $\langle S, s_0, L, \texttt{Act}, \delta \rangle$, a *probabilistic alternating simulation* (PA-simulation) is a relation $\sqsubseteq \subseteq S \times S$ such that $s \sqsubseteq t$ if

- $L(s) = L(t)$, and
- for all $\pi_1 \in \Pi_{\mathcal{I}}$, there exists $\pi_2 \in \Pi_{\mathcal{I}}$, such that $\widetilde{\delta}(s, \pi_1) \, \overline{\sqsubseteq}_{Sm} \, \widetilde{\delta}(t, \pi_2)$, where $\widetilde{\delta}(s, \pi) = \{\Delta \in \mathcal{D}(S) \mid \exists \pi' \in \Pi_{\mathcal{II}} : \widetilde{\delta}(s, \langle \pi, \pi' \rangle) = \Delta\}$.

If state s PA-simulates state t and t PA-simulates s, we say s and t are *PA-simulation equivalent*, which is written $s \simeq t$.

IV. A Modal Logic for Probabilistic GS

In the literature different modal logics have been introduced to characterise process semantics at different levels. Hennessy-Milner logic (HML) [2] provides a classical example that has been proved to be equivalent to bisimulation semantics in image-finite LTS. In other words, two states (or processes) satisfy the same set of HML formulas iff they are bisimilar. For a more comprehensive survey we refer to [20].

In this section we propose a modal logic for PGS that characterises a player's ability to enforce temporal properties. We define a new logic $\mathcal{L}^{\oplus}$ in the spirit of the logic of Deng et al. [4], [17]. The syntax of the logic $\mathcal{L}^{\oplus}$ is presented below.

$$\varphi ::= p \mid \neg p \mid \bigwedge_{i \in I} \varphi_i \mid \bigvee_{i \in I} \varphi_i \mid \langle\!\langle \mathcal{I} \rangle\!\rangle \varphi \mid \bigoplus_{j \in J} p_j \varphi_j \mid \bigsqcap_{j \in J} \varphi_j$$

In particular, p is an atomic formula that belongs to the set $\texttt{Prop}$. Formula $\bigwedge_{i \in I} \varphi_i$ produces a conjunction, and $\bigvee_{i \in I} \varphi_i$ produces a disjunction, both via a (possibly infinite) index set I. We then derive $\top = \bigwedge_{i \in \emptyset} \varphi_i$ is a formula that is true everywhere, and $\bot = \bigvee_{i\in\emptyset}\varphi_i$ is a formula false everywhere. $\langle\langle\mathcal{I}\rangle\rangle\varphi$ specifies player $\mathcal{I}$'s ability to enforce φ in the next step. The probabilistic summation operator $\bigoplus_{j\in J} p_j\varphi_j$ explicitly specifies that a distribution satisfying such a formula should be split with pre-defined weights, each part with weight p_j satisfying sub-formula φ_j. For a summation formula with index set J, we may explicitly write down each component coupled by its weight, such as in the way of $[p_1,\varphi_1]\oplus[p_2,\varphi_2]\oplus\ldots\oplus[p_{|J|},\varphi_{|J|}]$. The operator $\bigsqcap_{j\in J}\varphi_j$ asserts the existence of a linear interpolation among formulas φ_j. We only allow negation of formulas on the propositional level. We use $\mathcal{L}^{\oplus}$ to denote the set of modal formulas defined by the above syntax.

The semantics of $\mathcal{L}^{\oplus}$ is presented as follows. The interpretation of each formula is defined as a set of distributions of states in a finite PGS $\mathcal{G} = \langle S, s_0, L, \texttt{Act}, \delta\rangle$.

- $[\![p]\!] = \{\Delta \in \mathcal{D}(S) \mid \forall s \in \lceil\Delta\rceil : p \in L(s)\}$;
- $[\![\neg p]\!] = \{\Delta \in \mathcal{D}(S) \mid \forall s \in \lceil\Delta\rceil : p \notin L(s)\}$;
- $[\![\bigwedge_{i\in I}\varphi_i]\!] = \bigcap_{i\in I}[\![\varphi_i]\!]$; $[\![\bigvee_{i\in I}\varphi_i]\!] = \bigcup_{i\in I}[\![\varphi_i]\!]$;
- $[\![\langle\langle\mathcal{I}\rangle\rangle\varphi]\!] = \{\Delta \in \mathcal{D}(S) \mid \exists\pi_1 \in \Pi_{\mathcal{I}} : \forall\pi_2 \in \Pi_{\mathcal{II}} : \Delta \xrightarrow{\pi_1,\pi_2} \Theta \Longrightarrow \Theta \in [\![\varphi]\!]\}$;
- $[\![\bigoplus_{j\in J} p_j\varphi_j]\!] = \{\Delta \in \mathcal{D}(S) \mid \Delta = \sum_{j\in J} p_j\Delta_j \wedge \forall j \in J : \Delta_j \in [\![\varphi_j]\!]\}$;
- $[\![\bigsqcap_{j\in J}\varphi_j]\!] = \{\Delta \in \mathcal{D}(S) \mid \exists\{p_j\}_{j\in J} : \sum_{j\in J} p_j = 1 \wedge \Delta = \sum_{j\in J} p_j\Delta_j \wedge \forall j \in J : \Delta_j \in [\![\varphi_j]\!]\}$;

Note here we say a distribution Δ satisfies a propositional formula if the formula holds in every state in the support of Δ. The rest of the semantics is self-explained. Formally, given a formula $\varphi \in \mathcal{L}^{\oplus}$ and a distribution Δ, we write $\Delta \models \varphi$ iff $\Delta \in [\![\varphi]\!]$.

Remark. The probabilistic modal logic proposed by Parma and Segala [3] and Hermanns et al. [5] uses a fragment operator $[\varphi]_\alpha$, such that a distribution $\Delta \models [\varphi]_\alpha$ iff there exists $\Delta_1, \Delta_2 \in \mathcal{D}(S)$ such that $\Delta = \Delta_1 \oplus_\alpha \Delta_2$ and $\Delta_1 \models \varphi$. Informally, it states that a fragment of Δ with weight at least α satisfies φ. Note that the summation operator of $\mathcal{L}^{\oplus}$ can be used to encode the fragment operator $[\varphi]_\alpha$, in the way that $\Delta \models [\varphi]_\alpha$ iff $\Delta \models [\alpha,\varphi]\oplus[1-\alpha,\top]$. Therefore, a straightforward adaptation of the logic by Parma and Segala [3] and Hermanns et al. [5] does not yield a more expressive logic than $\mathcal{L}^{\oplus}$.

The semantics of $\bigsqcap_{j\in J}\varphi_j$ allows arbitrary linear interpolation among formulas φ_j. Similar to the way treating probabilistic summations, one may write down $\bigsqcap_{j\in J}\varphi_j$ by $[\varphi_1]\sqcap[\varphi_2]\sqcap\ldots\sqcap[\varphi_{|J|}]$. The following lemma is straightforward.

Lemma 3: Let $\varphi = \bigoplus_{j\in J} p_j\varphi_j$ and $\varphi' = \bigsqcap_{j\in J}\varphi_j$, and $\Delta \in \mathcal{D}(S)$. We have $\Delta \models \varphi$ implies $\Delta \models \varphi'$.

Similar to most of the literature, given a PGS $\mathcal{G}$, we define preorders on the set of states in $\mathcal{G}$ with respect to satisfaction of the modal logic $\mathcal{L}^{\oplus}$.

Definition 3: Given states $s, t \in S$, $s \sqsubseteq_{\mathcal{L}^{\oplus}} t$ if for all $\varphi \in \mathcal{L}^{\oplus}$, $s \models \varphi$ implies $t \models \varphi$. If $s \sqsubseteq_{\mathcal{L}^{\oplus}} t$ and $t \sqsubseteq_{\mathcal{L}^{\oplus}} s$, we write $s \simeq_{\mathcal{L}^{\oplus}} t$.

Now we state the first main result of the paper, and we leave its proof to the following two subsections.

Theorem 1: Let $\mathcal{G} = \langle S, s_0, L, \texttt{Act}, \delta\rangle$ be a PGS, then for all $s, t \in S$, $s \sqsubseteq t$ iff $s \sqsubseteq_{\mathcal{L}^{\oplus}} t$.

Proof: By combining Theorem 2 and Theorem 3 and by treating s and t as point distributions $s_\star$ and $t_\star$. □

Corollary 1: Let $\mathcal{G} = \langle S, s_0, L, \texttt{Act}, \delta\rangle$ be a PGS, then for all $s, t \in S$, $s \simeq t$ iff $s \simeq_{\mathcal{L}^{\oplus}} t$.

A. A Soundness Proof

Since the notion of PA-simulation given in Definition 2 is defined as a relation on states, in the following we show that the lifted PA-simulation is also a *simulation* on distributions over the states, which is used as a stepping stone to our soundness result. Similar to the way of treating distributions, we also allow linear combination of mixed actions.

Definition 4: Given a list of mixed actions $\langle\pi_i\rangle_{i\in I}$ (of player $\mathcal{I}$), and $\langle p_i\rangle_{i\in I}$ satisfying $\sum_{i\in I} p_i = 1$, $\sum_{i\in I} p_i\pi_i$ is a mixed action defined by $\left(\sum_{i\in I} p_i\pi_i\right)(s)(a) = \sum_{i\in I} p_i\cdot(\pi_i(s)(a))$ for all $s \in S$ and $a \in \texttt{Act}_{\mathcal{I}}$.

Lemma 4: Let $s \in S$, $\pi \in \Pi_{\mathcal{I}}$ and $\sigma = \sum_{i\in I} p_i\sigma_i \in \Pi_{\mathcal{II}}$, then $\widetilde{\delta}(s,\langle\pi,\sigma\rangle) = \sum_{i\in I} p_i\cdot\widetilde{\delta}(s,\langle\pi,\sigma_i\rangle)$.

Lemma 5: Let $s \in S$, $\pi = \sum_{i\in I} p_i\pi_i \in \Pi_{\mathcal{I}}$ and $\sigma \in \Pi_{\mathcal{II}}$, then $\widetilde{\delta}(s,\langle\pi,\sigma\rangle) = \sum_{i\in I} p_i\cdot\widetilde{\delta}(s,\langle\pi_i,\sigma\rangle)$.

Lemma 4 and Lemma 5 show that we can distribute the distributions over actions out of a transition operator.

Lemma 6: Let $\Delta \in \mathcal{D}(S)$ with $\Delta = \sum_{i\in I} p_i\Delta_i$, $\pi \in \Pi_{\mathcal{I}}$ and $\sigma \in \Pi_{\mathcal{II}}$, then we have $\widetilde{\delta}(\Delta,\langle\pi,\sigma\rangle) = \sum_{i\in I} p_i\cdot\widetilde{\delta}(\Delta_i,\langle\pi,\sigma\rangle)$.

Proof: Let $t \in S$, then

$$\begin{array}{rl} & \sum_{i\in I} p_i\cdot\widetilde{\delta}(\Delta_i,\langle\pi,\sigma\rangle)(t) \\ = & \sum_{i\in I} p_i\cdot\sum_{s\in S}\Delta_i(s)\cdot\widetilde{\delta}(s,\langle\pi,\sigma\rangle)(t) \\ = & \sum_{s\in S}\sum_{i\in I} p_i\Delta_i(s)\cdot\widetilde{\delta}(s,\langle\pi,\sigma\rangle)(t) \\ = & \sum_{s\in S}\Delta(s)\cdot\widetilde{\delta}(s,\langle\pi,\sigma\rangle)(t) \\ = & \widetilde{\delta}(\Delta,\langle\pi,\sigma\rangle)(t) \end{array}$$

□

Lemma 7: Let $\Delta, \Theta \in \mathcal{D}(S)$ with $\Delta = \sum_{i\in I} p_i\Delta_i$ and $\Theta = \sum_{i\in I} p_i\Theta_i$, $\pi_1, \pi_2 \in \Pi_{\mathcal{I}}$ and $\sigma_1, \sigma_2 \in \Pi_{\mathcal{II}}$. If $\widetilde{\delta}(\Delta_i,\langle\pi_1,\sigma_1\rangle)\ \overline{\sqsubseteq}\ \widetilde{\delta}(\Theta_i,\langle\pi_2,\sigma_2\rangle)$ for all $i \in I$, then $\widetilde{\delta}(\Delta,\langle\pi_1,\sigma_1\rangle)\ \overline{\sqsubseteq}\ \widetilde{\delta}(\Theta,\langle\pi_2,\sigma_2\rangle)$.

Proof: By Lemma 1, we have $\sum_{i\in I} p_i\cdot\widetilde{\delta}(\Delta_i,\langle\pi_1,\sigma_1\rangle)\ \overline{\sqsubseteq}\ \sum_{i\in I} p_i\cdot\widetilde{\delta}(\Theta_i,\langle\pi_2,\sigma_2\rangle)$. Then by applying Lemma 6 on both sides, we have $\widetilde{\delta}(\Delta,\langle\pi_1,\sigma_1\rangle)\ \overline{\sqsubseteq}\ \widetilde{\delta}(\Theta,\langle\pi_2,\sigma_2\rangle)$. □

Lemma 7 allows to merge the simulation by component distributions on both sides of the relation.

The next auxiliary lemma states that given a PA-simulation on states, the lifted PA-simulation on distributions of states can be treated as a *simulation* via mixed actions of player $\mathcal{I}$ and player $\mathcal{II}$.

Lemma 8: Let $\mathcal{G} = \langle S, s_0, \mathcal{L}, \texttt{Act}, \delta\rangle$ be a PGS, and $\sqsubseteq$ be a PA-simulation relation for $\mathcal{G}$. Given $\Delta\overline{\sqsubseteq}\Theta$, for all player $\mathcal{I}$ mixed actions π_1, there exists a player $\mathcal{I}$ mixed action π_2, such that $\widetilde{\delta}(\Delta,\pi_1)\overline{\sqsubseteq}_{Sm}\widetilde{\delta}(\Theta,\pi_2)$.

This can be proved by splitting distributions on both sides, and then merge related components to form distributions on both sides of the lifted relation, applying previous lemmas.

Proof: By definition there exists a weight function w, such that for all states $s, t \in S$, we have $w(s,t) > 0$ implies $s \sqsubseteq t$. For each pair of states $s, t \in S$ with $w(s,t) > 0$, we have a mixed action $\pi_{s,t}$ such that $\delta(s, \pi_1) \ \overline{\sqsubseteq}_{Sm} \ \delta(t, \pi_{s,t})$. We construct a mixed action $\pi_2 = \bigoplus_{s,t \in S: w(s,t)>0} w(s,t)\pi_{s,t}$, and show that $\widetilde{\delta}(\Delta, \pi_1) \ \overline{\sqsubseteq}_{Sm} \ \widetilde{\delta}(\Theta, \pi_2)$.

Let $\sigma_2 \in \Pi_{\mathcal{II}}$, and we show there exists $\sigma_1 \in \Pi_{\mathcal{II}}$ such that $\widetilde{\delta}(\Delta, \langle \pi_1, \sigma_1 \rangle) \ \overline{\sqsubseteq} \ \widetilde{\delta}(\Theta, \langle \pi_2, \sigma_2 \rangle)$. Since for each pair of states $s, t \in S$ with $w(s,t) > 0$, we have $\widetilde{\delta}(s, \pi_1) \ \overline{\sqsubseteq}_{Sm} \ \widetilde{\delta}(t, \pi_{s,t})$, there exists $\sigma_{s,t} \in \Pi_{\mathcal{II}}$ such that $\widetilde{\delta}(s, \langle \pi_1, \sigma_{s,t} \rangle) \ \overline{\sqsubseteq} \ \widetilde{\delta}(t, \langle \pi_{s,t}, \sigma_2 \rangle)$. Define $\sigma_1 = \bigoplus_{s,t \in S: w(s,t)>0} w(s,t)\sigma_{s,t}$. By Lemma 1, we have

$$\bigoplus_{s,t \in S: w(s,t)>0} w(s,t) \cdot \widetilde{\delta}(s, \langle \pi_1, \sigma_{s,t} \rangle)$$
$$\overline{\sqsubseteq} \bigoplus_{s,t \in S: w(s,t)>0} w(s,t) \cdot \widetilde{\delta}(t, \langle \pi_{s,t}, \sigma_2 \rangle)$$

Applying Lemma 4 on the LHS and Lemma 5 on the RHS, we have $\widetilde{\delta}(s, \langle \pi_1, \sigma_1 \rangle) \ \overline{\sqsubseteq} \ \widetilde{\delta}(t, \langle \pi_2, \sigma_2 \rangle)$. Then by Lemma 7, we have $\widetilde{\delta}(\Delta, \langle \pi_1, \sigma_1 \rangle) \ \overline{\sqsubseteq} \ \widetilde{\delta}(\Theta, \langle \pi_2, \sigma_2 \rangle)$. □

Theorem 2: Given $\varphi \in \mathcal{L}^{\oplus}$ and $\Delta \overline{\sqsubseteq} \Theta$ where $\sqsubseteq$ is a PA-simulation, then $\Delta \models \varphi$ implies $\Theta \models \varphi$.

The theorem can be proved by structural induction on φ. The base cases when $\varphi = p$ and $\neg p$ are straightforward. For the INDUCTION STEP, we only show the case of $\langle\langle \mathcal{I} \rangle\rangle \psi$, as the other cases are straightforward. If $\varphi = \langle\langle \mathcal{I} \rangle\rangle \psi$, then there exists a player $\mathcal{I}$ mixed actions π_1 such that for all player $\mathcal{II}$ mixed actions σ_1, $\Delta \xrightarrow{\pi_1, \sigma_1} \Delta'$ and $\Delta' \models \psi$. By Lemma 8, there exists a player $\mathcal{I}$ mixed action π_2 such that $\widetilde{\delta}(\Delta, \pi_1) \overline{\sqsubseteq}_{Sm} \widetilde{\delta}(\Theta, \pi_2)$. Therefore, for all player $\mathcal{II}$ mixed actions σ_2 there exists a player $\mathcal{II}$ strategy σ_1', such that $\Delta \xrightarrow{\pi_1, \sigma_1'} \Delta''$, $\Theta \xrightarrow{\pi_2, \sigma_2} \Theta'$, and $\Delta'' \overline{\sqsubseteq} \Theta'$. Since $\Delta'' \models \psi$, by I.H., $\Theta' \models \psi$. This shows that π_2 is the player $\mathcal{I}$ mixed action for $\Theta \models \langle\langle \mathcal{I} \rangle\rangle \psi$.

B. A Completeness Proof

The completeness is proved by approximating the relations $\sqsubseteq$ and $\sqsubseteq_{\mathcal{L}^{\oplus}}$. For PA-simulation we construct relations $\sqsubseteq_n$ for $n \in \mathbb{N}$, where n denotes the number of steps that are required to check for a state to simulate another. (Intuitively, the more steps to check, the harder for a pair of states to satisfy the relation.) Similarly we define $\sqsubseteq_n^{\mathcal{L}}$, restricting to formulas in $\mathcal{L}^{\oplus}$ with *size* up to n. Then we prove that the relation $\sqsubseteq_n^{\mathcal{L}}$ is contained in $\sqsubseteq_n$ for all $n \in \mathbb{N}$.

Definition 5: Let $\mathcal{L}_0^{\oplus}$ be the set of formulas constructed by using only p, $\neg p$ and $\bigwedge_{i \in I} \varphi_i$. For $n \in \mathbb{N}$, a formula $\varphi \in \mathcal{L}_{n+1}^{\oplus}$ if either $\varphi \in \mathcal{L}_n^{\oplus}$ or φ is a conjunction of formulas of the form $\langle\langle \mathcal{I} \rangle\rangle \bigsqcap_{i \in I} \bigoplus_{j \in J} p_j \varphi_{i,j}$, where each $\varphi_{i,j} \in \mathcal{L}_n^{\oplus}$.

Intuitively, formulas in $\mathcal{L}_n^{\oplus}$ require n steps of transitions (for player $\mathcal{I}$) to enforce. Given states $s, t \in S$, we write $s \sqsubseteq_n^{\mathcal{L}} t$, if for all $\varphi \in \mathcal{L}_n^{\oplus}$, $s_\star \models \varphi$ implies $t_\star \models \varphi$. Similarly we define approximating relations for PA-simulation.

Definition 6: Given $s, t \in S$, $s \sqsubseteq_0 t$ if $L(s) = L(t)$. For $n \in \mathbb{N}$, $s \sqsubseteq_{n+1} t$ if $s \sqsubseteq_n t$, and for all player $\mathcal{I}$ mixed actions π_1 there exists a player $\mathcal{I}$ mixed action π_2, such that $\widetilde{\delta}(s, \pi_1) \ (\overline{\sqsubseteq}_n)_{Sm} \widetilde{\delta}(t, \pi_2)$.

Before starting the completeness proof, we define formulas that characterise properties of the game states. Let $s \in S$, the 0-characteristic formula for s is $\phi_s^0 = \bigwedge\{p \mid p \in L(s)\} \wedge \bigwedge\{\neg q \mid q \in \mathtt{Prop} \setminus L(s)\}$. Plainly, the level 0-characterisation considers only propositional formulas. For a distribution, we specify the characteristic formulas for the states in its support proportional to weights. The 0-characteristic formula ϕ_Δ^0 for distribution Δ is $\bigoplus_{t \in \lceil \Delta \rceil} \Delta(t) \cdot \phi_t^0$. Given all n-characteristic formulas defined, the $(n+1)$-characteristic formula ϕ_s^{n+1} for state s is $\bigwedge_{\pi \in \mathcal{D}(\mathtt{Act}_{\mathcal{I}})} \langle\langle \mathcal{I} \rangle\rangle \bigsqcap_{b \in \mathtt{Act}_{\mathcal{II}}} \phi_{\Delta_{\pi,b}}^n$, where $s_\star \xrightarrow{\pi, b_\star} \Delta_{\pi,b}$. Similarly, an n-characteristic formula ϕ_Δ^{n+1} for distribution Δ is $\bigoplus_{t \in \lceil \Delta \rceil} \Delta(t) \cdot \phi_t^{n+1}$.

Obviously every state or distribution satisfies its own characteristic formula, and the following lemma can be proved straightforwardly by induction on $n \in \mathbb{N}$.

Lemma 9: For all $\Delta \in \mathcal{D}(S)$, $\Delta \models \phi_\Delta^n$ for all $n \in \mathbb{N}$.

Lemma 10: For all states $s, t \in S$ and $n \in \mathbb{N}$, $s \sqsubseteq_n^{\mathcal{L}} t$ implies $s \sqsubseteq_n t$.

Proof: For each $n \in \mathbb{N}$, we have $s_\star \models \phi_s^n$ by Lemma 9. Let $s \sqsubseteq_n^{\mathcal{L}} t$, then $t_\star \models \phi_s^n$. We proceed by induction on the level of approximation n to show that $s \sqsubseteq_n t$.

We show that the state-based relation can be carried over to distributions. Suppose for all $s, t \in S$, $s \sqsubseteq_n^{\mathcal{L}} t$ implies $t_\star \models \phi_s^n$. Given two distributions $\Delta, \Theta \in \mathcal{D}(S)$, assume that $\Delta \ \overline{\sqsubseteq}_n^{\mathcal{L}} \ \Theta$. Then there exists a weight function w, such that $\Delta = \sum_{s \in \lceil \Delta \rceil, t \in \lceil \Theta \rceil : w(s,t)>0} w(s,t) \cdot s_\star$, and $\Theta = \sum_{s \in \lceil \Delta \rceil, t \in \lceil \Theta \rceil : w(s,t)>0} w(s,t) \cdot t_\star$, and $s \sqsubseteq_n^{\mathcal{L}} t$ for all $w(s,t) > 0$. Since ϕ_Δ^n can be written as $\bigoplus_{s \in \lceil \Delta \rceil, t \in \lceil \Theta \rceil : w(s,t)>0} w(s,t) \cdot \phi_s^n$, we must have $\Theta \models \phi_\Delta^n$ as well.

BASE CASE: Trivial.

INDUCTION STEP: Let $t_\star \models \phi_s^{n+1}$, where $\phi_s^{n+1} = \bigwedge_{\pi \in \mathcal{D}(\mathtt{Act}_{\mathcal{I}})} \langle\langle \mathcal{I} \rangle\rangle \bigsqcap_{b \in \mathtt{Act}_{\mathcal{II}}} \phi_{\Delta_{\pi,b}}^n$. Then for each $\pi \in \mathcal{D}(\mathtt{Act}_{\mathcal{I}})$, $t_\star \models \langle\langle \mathcal{I} \rangle\rangle \bigsqcap_{b \in \mathtt{Act}_{\mathcal{II}}} \phi_{\Delta_{\pi,b}}^n$. By definition there exists a player $\mathcal{I}$ mixed action π', such that for every player $\mathcal{II}$ mixed action σ, we have $t_\star \xrightarrow{\pi', \sigma} \Theta$ and $\Theta \models \bigsqcap_{b \in \mathtt{Act}_{\mathcal{II}}} \phi_{\Delta_{\pi,b}}^n$. We need to show that $\widetilde{\delta}(s, \pi) \ (\overline{\sqsubseteq}_n^{\mathcal{L}})_{Sm} \ \widetilde{\delta}(t, \pi')$.

It suffices to check each $b \in \mathtt{Act}_{\mathcal{II}}$ from t can be followed by a player $\mathcal{II}$ mixed action from s to establish such a simulation. Let b' be a player $\mathcal{II}$ action, and $t_\star \xrightarrow{\pi', b'_\star} \Theta$. Since $\Theta \models \bigsqcap_{b \in \mathtt{Act}_{\mathcal{II}}} \phi_{\Delta_{\pi,b}}^n$, there exists a list of probability values $\langle p_c \rangle_{c \in \mathtt{Act}_{\mathcal{II}}}$, such that $\Theta \models \sum_{c \in \mathtt{Act}_{\mathcal{II}}} p_c \phi_{\Delta_{\pi,c}}^n$. Then by definition, we have $\Theta = \sum_{c \in \mathtt{Act}_{\mathcal{II}}} p_c \cdot \Theta_c$, $\sum_{c \in \mathtt{Act}_{\mathcal{II}}} p_c = 1$ and $\Theta_c \models \phi_{\Delta_{\pi,c}}^n$ for all $c \in \mathtt{Act}_{\mathcal{II}}$. In state s, we define a player $\mathcal{II}$ mixed action σ satisfying $\sigma(s)(c) = p_c$ for all $c \in \mathtt{Act}_{\mathcal{II}}$. Then by Lemma 4, we have $\widetilde{\delta}(s, \langle \pi, \sigma \rangle) = \sum_{c \in \mathtt{Act}_{\mathcal{II}}} p_c \cdot \delta(s, \langle \pi, c \rangle) = \Delta_{\pi,\sigma}$ for all $c \in \mathtt{Act}_{\mathcal{II}}$. By Lemma 1, it suffices to show $\delta(s, \langle \pi, c \rangle) \overline{\sqsubseteq}_n \Theta_c$ for all $c \in \mathtt{Act}_{\mathcal{II}}$. Since $\Theta_c \models \phi_{\Delta_{\pi,c}}^n$, we have $\Delta_{\pi,c} \overline{\sqsubseteq}_n \Theta_c$ by I.H.. □

Intuitively, by fixing a mixed strategy from player $\mathcal{I}$, a transition in the PGS is bounded by deterministic actions from player $\mathcal{II}$, as mimicked in the structure of the characteristic formulas. The way of showing satisfaction of a characteristic formula thus mimics the PA-simulation in the proof of Lemma 10.

Theorem 3: For all $s, t \in S$, $s \sqsubseteq^{\mathcal{L}} t$ implies $s \sqsubseteq t$.

Proof: In a finite state PGS (i.e., the space $S \times S$ is finite) there exists $n \in \mathbb{N}$ such that $\sqsubseteq = \sqsubseteq_n$. Since $\sqsubseteq^{\mathcal{L}} \subseteq \sqsubseteq^{\mathcal{L}}_n$ for all n, and $\sqsubseteq^{\mathcal{L}}_n \subseteq \sqsubseteq_n$ by Lemma 10, we have $\sqsubseteq^{\mathcal{L}} \subseteq \sqsubseteq_n = \sqsubseteq$. □

V. Probabilistic Alternating-time Mu-Calculus

Modal logics of finite modality depth are not enough to express temporal requirements such as "something bad never happens". In this section, we extend the logic $\mathcal{L}^{\oplus}$ into a Probabilistic Alternating-time μ-calculus (PAMu), by adding variables and fixpoint operators.

$$\begin{aligned}\varphi ::= \quad & p \mid \neg p \mid \bigwedge_{i\in I} \varphi_i \mid \bigvee_{i\in I} \varphi_i \mid \langle\langle \mathcal{I} \rangle\rangle \varphi \mid \bigsqcap_{j\in J} \varphi_j \\ & \mid \bigoplus_{j\in J} p_j \varphi_j \mid Z \mid \mu Z.\varphi \mid \nu Z.\varphi\end{aligned}$$

Let the environment $\rho : \mathcal{V} \to \mathcal{P}(\mathcal{D}(S))$ be a mapping from variables in $\mathcal{V}$ to sets of distributions on states, and the semantics of the fixpoint operators of PAMu are defined in the standard way.

- $[\![p]\!]_\rho = \{\Delta \in \mathcal{D}(S) \mid \forall s \in \lceil \Delta \rceil : p \in L(s)\}$;
- $[\![\neg p]\!]_\rho = \{\Delta \in \mathcal{D}(S) \mid \forall s \in \lceil \Delta \rceil : p \notin L(s)\}$;
- $[\![\bigwedge_{i\in I} \varphi_i]\!]_\rho = \bigcap_{i\in I} [\![\varphi_i]\!]_\rho$; $[\![\bigvee_{i\in I} \varphi_i]\!] = \bigcup_{i\in I} [\![\varphi_i]\!]_\rho$;
- $[\![\langle\langle \mathcal{I} \rangle\rangle \varphi]\!]_\rho = \{\Delta \in \mathcal{D}(S) \mid \exists \pi_1 \in \Pi_{\mathcal{I}} : \forall \pi_2 \in \Pi_{\mathcal{II}} : \Delta \xrightarrow{\pi_1, \pi_2} \Theta$ implies $\Theta \in [\![\varphi]\!]_\rho\}$;
- $[\![\bigoplus_{j\in J} p_j \varphi_j]\!]_\rho = \{\Delta \in \mathcal{D}(S) \mid \Delta = \sum_{j\in J} p_j \Delta_j \wedge \forall j \in J : \Delta_j \in [\![\varphi_j]\!]_\rho\}$;
- $[\![\bigsqcap_{j\in J} \varphi_j]\!]_\rho = \{\Delta \in \mathcal{D}(S) \mid \exists \{p_j\}_{j\in J} : \sum_{j\in J} p_j = 1 \wedge \Delta = \sum_{j\in J} p_j \Delta_j \wedge \forall j \in J : \Delta_j \in [\![\varphi_j]\!]_\rho\}$;
- $[\![Z]\!]_\rho = \rho(Z)$;
- $[\![\mu Z.\varphi]\!]_\rho = \bigcap\{D \subseteq \mathcal{D}(S) \mid [\![\varphi]\!]_\rho [Z \mapsto D] \subseteq D\}$;
- $[\![\nu Z.\varphi]\!]_\rho = \bigcup\{D \subseteq \mathcal{D}(S) \mid D \subseteq [\![\varphi]\!]_\rho [Z \mapsto D]\}$.

The set of closed PAMu formulas are the formulas with all variables bounded, which form the set $\mathcal{L}^\mu$, and we can safely drop the environment ρ for those formulas.

Example 4: For the rock-paper-scissors game in Figure 1, the property describing that player $\mathcal{I}$ has a strategy to eventually win the game once can be expressed as $\mu Z.\mathsf{win}_{\mathcal{I}} \vee \langle\langle \mathcal{I} \rangle\rangle Z$. This property does not hold. However, player $\mathcal{I}$ has a strategy to eventually win the game with probability almost $\frac{1}{2}$, i.e., the system satisfies $\mu Z.([\frac{1}{2} - \epsilon, \mathsf{win}_{\mathcal{I}}] \oplus [\frac{1}{2} + \epsilon, \top]) \vee \langle\langle \mathcal{I} \rangle\rangle Z$ for arbitrarily small $\epsilon > 0$. We explain the reason why players can only enforce ϵ-optimal strategies in a later part of the section.

The logic characterisation of PA-Simulation can be extended to PAMu.

Theorem 4: Given $\Delta, \Theta \in \mathcal{D}(S)$, $\Delta \overline{\sqsubseteq} \Theta$ iff $\{\varphi \in \mathcal{L}^\mu \mid \Delta \models \varphi\} \subseteq \{\varphi \in \mathcal{L}^\mu \mid \Theta \models \varphi\}$.

Since $\mathcal{L}^{\oplus}$ is syntactically a sublogic of $\mathcal{L}^\mu$, we only need to show the soundness of PA-simulation to the logic $\mathcal{L}^\mu$. We apply the classical approach of approximants for Modal Mu-Calculus [21]. Given formulas $\mu Z.\varphi$ and $\nu Z.\varphi$, we define the following.

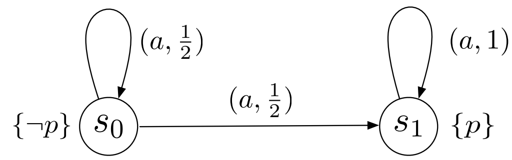

Fig. 3. An example for $\langle\langle \mathcal{I} \rangle\rangle^{\geq 1} \Diamond p$.

$$\begin{aligned}&\mu^0 Z.\varphi = \bot && \nu^0 Z.\varphi = \top \\ &\mu^{i+1} Z.\varphi = \varphi[Z \mapsto \mu^i Z.\varphi] && \nu^{i+1} Z.\varphi = \varphi[Z \mapsto \nu^i Z.\varphi] \\ &\mu^\omega Z.\varphi = \bigvee_{i\in\mathbb{N}} \mu^i Z.\varphi && \nu^\omega Z.\varphi = \bigwedge_{i\in\mathbb{N}} \nu^i Z.\varphi.\end{aligned}$$

Next we show approximants are semantically equivalent to the fixpoint formulas.

Lemma 11:

1) $[\![\mu^\omega Z.\varphi]\!] = [\![\mu Z.\varphi]\!]$;
2) $[\![\nu^\omega Z.\varphi]\!] = [\![\nu Z.\varphi]\!]$.

We briefly sketch a proof of Lemma 11(1), and the proof for the other part of the lemma is similar. To show $[\![\mu^\omega Z.\varphi]\!] \subseteq [\![\mu Z.\varphi]\!]$, we initially have $[\![\mu^0 Z.\varphi]\!] = \emptyset \subseteq [\![\mu Z.\varphi]\!]$, then by the monotonicity of φ, given $[\![\mu^i Z.\varphi]\!] \subseteq [\![\mu Z.\varphi]\!]$, we prove $[\![\mu^{i+1} Z.\varphi]\!] \subseteq [\![\mu Z.\varphi]\!]$ by applying φ on both sides of $\subseteq$. Therefore, $[\![\mu^i Z.\varphi]\!] \subseteq [\![\mu Z.\varphi]\!]$ for all $i \in \mathbb{N}$, thus $[\![\bigvee_{i\in\mathbb{N}} \mu^i Z.\varphi]\!] \subseteq [\![\mu Z.\varphi]\!]$. To show $[\![\mu Z.\varphi]\!] \subseteq [\![\mu^\omega Z.\varphi]\!]$, it is straightforward to see that $\mu^\omega Z.\varphi$ is a prefixpoint, therefore it contains $\mu Z.\varphi$, the intersection of all prefixpoints.

From Lemma 11, and by the soundness of PA-simulation to $\mathcal{L}^{\oplus}$ (Theorem 2), we get the the soundness of PA-simulation to the logic $\mathcal{L}^\mu$, as required.

Expressiveness of PAMu. There exist game-based extensions of probabilistic temporal logics, such as the logic PAMC [10] that extends the Alternating-time Mu-Calculus [6], and PATL [22] that extends PCTL [23]. The semantics of both logics are sets of states, rather than sets of distributions. It has also been shown in [10] that PAMC and PATL are incomparable on probabilistic game structures, based on a result showing that PCTL and PμTL are incomparable on Markov chains [9]. Here we make a short comparison between PAMu and those logics.

Distribution formulas of PAMu cannot be expressed by state-based logics. For example, the formula $\langle\langle \mathcal{I} \rangle\rangle [\frac{1}{2}, p] \oplus [\frac{1}{2}, q]$, expressing that player $\mathcal{I}$ has a strategy to enforce in the next move a distribution which has half of its weight satisfying p and the other half satisfying q, cannot be expressed by PATL or PAMC. As the latter two logics have probability values bundled with strategy modalities, a formula such as $\langle\langle \mathcal{I} \rangle\rangle^{\geq \frac{1}{2}} p \wedge \langle\langle \mathcal{I} \rangle\rangle^{\geq \frac{1}{2}} q$ denotes that player $\mathcal{I}$ has a strategy to enforce p with at least probability $\frac{1}{2}$ in the next step and player $\mathcal{I}$ also has a *possibly different* strategy to enforce q with at least probability $\frac{1}{2}$ in the next step. However, the resulting states (or distributions) that satisfy p and q may overlap.

The PATL formula $\langle\langle \mathcal{I} \rangle\rangle^{\geq 1} \Diamond p$ is not expressible by PAMu. Given the PGS in Figure 3 where player $\mathcal{I}$ has action set $\{a\}$ and player $\mathcal{II}$ has action set $\emptyset$. Then it is straightforward to see both s_0 and s_1 satisfies $\langle\langle \mathcal{I} \rangle\rangle^{\geq 1} \Diamond p$. The closest formula in PAMu is $\mu Z.p \vee \langle\langle \mathcal{I} \rangle\rangle Z$, but $s_0 \not\models \mu Z.p \vee \langle\langle \mathcal{I} \rangle\rangle Z$. More precisely, $s_0 \models \mu Z.([\alpha, p] \oplus [1-\alpha, \top]) \vee \langle\langle \mathcal{I} \rangle\rangle Z$ for all $0 \leq \alpha < 1$. Intuitively, the semantics of the least fixpoint operator in PAMu only track finite number of probabilistic transitions, as starting from s_0, player $\mathcal{I}$ can only reach distributions that satisfy p with probability strictly less than 1 with finite number of steps. Intuitively, $s_{0\star} \xrightarrow{a} [\frac{1}{2}, s_0] \oplus [\frac{1}{2}, s_1] \xrightarrow{a} [\frac{1}{4}, s_0] \oplus [\frac{3}{4}, s_1] \xrightarrow{a} \ldots \xrightarrow{a} [\frac{1}{2^i}, s_0] \oplus [1-\frac{1}{2^i}, s_1] \ldots$. We shall see that in a finite number of transitions one never reaches $s_{1\star}$ from $s_{0\star}$ with strict probability 1. However, such a restriction may be alleviated in practice, as implemented in PRISM-game [12], ϵ-optimal strategies are synthesized for unbounded reachability properties. PAMC formulas that contain the $[\![\mathcal{I}]\!]^{\bowtie\alpha}$ modalities do not seem expressible in PAMu. For instance, the PAMC formula $[\![\mathcal{I}]\!]^{\geq\alpha}\varphi$ [3] is semantically equivalent to $\neg\langle\langle \mathcal{I} \rangle\rangle^{<\alpha}\varphi$, which is not expressible by PAMu as negation is only allowed at the propositional level in PAMu.

Example 5: The authors of [12] proposed a CGS variant of a futures market investor model [24], which studies the interactions between an investor and a stock market. The investor and the market take their decisions simultaneously in the CGS model, and the authors show that this does not give any additional benefits to the investor by analysing his or her maximum expected value over a fixed period of time.[4] We take this example to demonstrate the expressiveness of PAMu. For instance, the property "it is always possible for the investor to cash in" can be specified with two nested fixpoints as

$$\nu X.(\mu Y.\mathsf{cashin} \vee \langle\langle \mathit{investor} \rangle\rangle\, Y) \wedge \langle\langle \mathit{investor} \rangle\rangle X.$$

Here the greatest fixpoint νX asserts that the investor is able to enforce the system to keep in a set of states (or a distribution of states taken from this set), such that from any state it is possible to cash in within finite number of transitions which is enforced by the inner least fixpoint μY.

Another interesting property is to check whether the investor has a strategy to ensure a good chance to make a profit. This can be formulated in PAMu with $\frac{1}{2} < \alpha \leq 1$, as

$$\mu Z.(\mathsf{cashin} \wedge [\alpha, \mathsf{profit}] \oplus [1-\alpha, \top]) \vee \langle\langle \mathit{investor} \rangle\rangle Z$$

By using the least fixpoint μZ, this formula asserts that the investor is able to reach a position within finite number of steps where the investor's cash in behaviour is accompanied by a profitable state with at least probability α, where $\alpha > \frac{1}{2}$.

VI. Related Work

Segala and Lynch [25] introduce a probabilistic simulation relation which preserves probabilistic computation tree logic (PCTL) formulas without negation and existential quantification. Segala introduces the notion of probabilistic forward simulation, which relates states to probability distributions over states and is sound and complete for trace distribution precongruence [26], [27]. Parma and Segala [3] use a probabilistic extension of the Hennessy-Milner logic which allows countable conjunction and admits a new operator $[\phi]_p$ – a distribution satisfies $[\phi]_p$ if the probability on the set of states satisfying ϕ is at least p, with a sound and complete logic characterisation. Hermanns et al. [5] further extend this result for image-infinite probabilistic automata. Deng et al. [4], [17] introduce a few probabilistic operators to derive a probabilistic modal mu-calculus (pMu). A fragment of pMu is proved to characterise (strong) probabilistic simulation in finite-state probabilistic automata. Our work extends the above work by enriching with the concurrent game semantics that are initiated in [6], which is discussed in the following paragraph.

Alur, Henzinger and Kupferman [6] define alternating-time temporal logic (ATL) to generalise CTL for game structures by requiring each path quantifier to be parameterised by a set of agents. GS are more general than LTS, in the sense that they allow both collaborative and adversarial behaviours 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". The alternating simulation, which is a natural game-theoretic interpretation of the classical simulation in (deterministic) multi-player games, is introduced in [7]. Logic characterisation of this relation concentrates on a subset of ATL* formulas where negations are only allowed at propositional level and all path quantifiers are parameterised by a predefined set of agents. Comparing with the standard alternating simulation and its logic characterisation, PA-simulation focuses on the extension which allows mixed strategies in probabilistic game structures (PGS).

Game structures deal well with systems in which the players execute a *pure strategy*, i.e., a strategy in which the moves are chosen deterministically. However, *mixed strategies*, which are formed by combining pure strategies, are necessary for a player to achieve *optimal* rewards [28]. Zhang and Pang [8] extend the notion of game structures to probabilistic game structures (PGS) and introduce notions of simulation that are sound for a fragment of probabilistic alternating-time temporal logic (PATL), a probabilistic extension of ATL.

Fixpoint logics for sets of states in Markov chains and PGS have been studied more recently in [9], [10], and a short comparison is given in Section V.

Metric-based simulation on game structures have been studied by de Alfaro et al. [29] regarding the probability of winning games whose goals are expressed in quantitative μ-calculus (qMu) [24]. Two states are equivalent if the players can win the same games with the same probability from both states, and *similarity* among states can thus be measured. Algorithmic verification complexities are further studied for MDP and turn-based games [30]. Metric-based approaches allow to analyze similarity with a quantitative measure, in which sense our approach is more strict. However our definition of PA-simulation is purely by actions and strategies, while metric-based approach is more target-based as it defines similarity on states by the ability to achieve same outcomes with similar probabilities.

[3] $[\![\mathcal{I}]\!]^{\geq\alpha}\varphi$ can be interpreted as for all player $\mathcal{I}$ strategies π, there exists player $\mathcal{II}$ strategy σ, such that the combined effect of π and σ enforces φ with probability at least α.

[4] For details of the model, we refer to [24] and the website https://www.prismmodelchecker.org.

More recently, algorithmic verification of turn-based and concurrent games have been implemented in an extension of PRISM [12], [31]. The properties can be specified as state formulas, path formulas and reward formulas. The verification procedure requires solving matrix games for concurrent game structures, and it applies value iteration algorithms to approach the goal (similar to [29], [32]). For unbounded properties, the synthesised strategy is memoryless (but only ϵ-optimal strategies). Finite-memory strategies are synthesised for bounded properties. The model checking algorithms for PAMu in PGS may be extended from the existing algorithms implemented for PRISM.

VII. Conclusions and Future Work

In this work, we have presented sound and complete modal characterisations of PA-simulation for concurrent games by introducing two new logics $\mathcal{L}^{\oplus}$ and PAMu (with fixpoints). Both logics incorporate nondeterministic and probabilisitic features and express the ability of the players to enforce a property in the current state. Since $\mathcal{L}^{\oplus}$ and PAMu restrict negation at the propositional level, the logics discussed in this paper cannot characterize bisimulation-like relations which require a logic to specify not only what is enabled but also what is not enabled from a game state. We leave the problem of PA-bisimulation characterisation as future work. In the future, we also aim to study verification complexities for these two logics.

References

[1] R. Milner, *Communication and Concurrency*. Prentice Hall, 1989.

[2] M. Hennessy and R. Milner, "Algebraic laws for nondeterminism and concurrency," *Journal of the ACM*, vol. 32, no. 1, pp. 137–161, 1985.

[3] A. Parma and R. Segala, "Logical characterizations of bisimulations for discrete probabilistic systems," in *Proc. 10th Conference on Foundations of Software Science and Computational Structures*, ser. LNCS, vol. 4423. Springer, 2007, pp. 287–301.

[4] Y. Deng and R. J. van Glabbeek, "Characterising probabilistic processes logically (extended abstract)," in *Proc. 17th Conference on Logic for Programming, Artificial Intelligence, and Reasoning*, ser. LNCS, vol. 6397. Springer, 2010, pp. 278–293.

[5] H. Hermanns, A. Parma, R. Segala, B. Wachter, and L. Zhang, "Probabilistic logical charaterization," *Information and Computation*, vol. 209, no. 2, pp. 154–172, 2011.

[6] R. Alur, T. A. Henzinger, and O. Kupferman, "Alternating-time temporal logic," *Journal of ACM*, vol. 49, no. 5, pp. 672–713, 2002.

[7] R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi, "Alternating refinement relations," in *Proc. 9th Conference on Concurrency Theory*, ser. LNCS, vol. 1466. Springer, 1998, pp. 163–178.

[8] C. Zhang and J. Pang, "On probabilistic alternating simulations," in *Proc. 6th IFIP Conference on Theoretical Computer Science*, ser. IFIP AICT, vol. 323, 2010, pp. 71–85.

[9] W. Liu, L. Song, J. Wang, and L. Zhang, "A simple probabilistic extension of modal mu-calculus," in *Proc. 24th International Joint Conference on Artificial Intelligence*. AAAI Press, 2015, pp. 882–888.

[10] F. Song, Y. Zhang, T. Chen, Y. Tang, and Z. Xu, "Probabilistic alternating-time μ-calculus," in *Proc. 33rd AAAI Conference on Artificial Intelligence*. AAAI Press, 2019.

[11] D. C. Kozen, "Results on the propositional μ-calculus," *Theoretical Computer Science*, vol. 27, pp. 333–354, 1983.

[12] M. Kwiatkowska, G. Norman, D. Parker, and G. Santos, "Automated verification of concurrent stochastic games," in *Proc. 15th Conference on Quantitative Evaluation of Systems*, ser. LNCS, vol. 11024. Springer, 2018, pp. 223–239.

[13] M. B. Smyth, "Power domains," *Journal of Computer and System Sciences*, vol. 16, no. 1, pp. 23–36, 1978.

[14] B. v. Karger, "Plotkin, Hoare and Smyth order: On observational models for CSP," in *Proc. of the IFIP TC2/WG2.1/WG2.2/WG2.3 Working Conference on Programming Concepts, Methods and Calculi*, ser. PROCOMET'94, 1994, pp. 383–402.

[15] B. Jonsson and K. G. Larsen, "Specification and refinement of probabilistic processes," in *Proc. 6th IEEE Symposium on Logic in Computer Science*. IEEE CS, 1991, pp. 266–277.

[16] R. Segala, "Modeling and verification of randomized distributed real-time systems," Ph.D. dissertation, Massachusetts Institute of Technology, 1995.

[17] Y. Deng, R. van Glabbeek, M. Hennessy, C. Morgan, and C. Zhang, "Characterising testing preorders for finite probabilistic processes," in *Proc. 22nd IEEE Symposium on Logic in Computer Science*. IEEE CS, 2007, pp. 313–325.

[18] Y. Deng, R. J. van Glabbeek, M. Hennessy, and C. Morgan, "Testing finitary probabilistic processes," in *Proc. 20th Conference on Concurrency Theory*, ser. LNCS, vol. 5710. Springer, 2009, pp. 274–288.

[19] C. Zhang and J. Pang, "An algorithms for probabilistic alternating simulation," in *Proc. 38th Conference on Current Trends in Theory and Practice of Computer Science*, ser. LNCS, vol. 7147. Springer, 2012, pp. 431–442.

[20] R. van Glabbeek, "The linear time-branching time spectrum I. The semantics of concrete, sequential processes." in *Handbook of Process Algebra*. Elsevier, 2001, pp. 3–99.

[21] J. Bradfield and C. Stirling, "Modal mu-calculi," in *The Handbook of Modal Logic*. Elsevier, 2006, p. 721–756. [Online]. Available: https://homepages.inf.ed.ac.uk/jcb/Research/MLH-bradstir.pdf

[22] T. Chen and J. Lu, "Probabilistic alternating-time temporal logic and model checking algorithm," in *Proc. 4th Conference on Fuzzy Systems and Knowledge Discovery*. IEEE CS, 2007, pp. 35–39.

[23] H. Hansson and B. Jonsson, "A logic for reasoning about time and reliability," *Formal Aspects of Computing*, vol. 6, no. 5, pp. 512–535, 1994.

[24] A. McIver and C. Morgan, "Results on the quantitative μ-calculus qMu," *ACM Transactions on Computational Logic*, vol. 8, no. 1, 2007.

[25] R. Segala and N. A. Lynch, "Probabilistic simulations for probabilistic processes," *Nordic Journal of Computing*, vol. 2, no. 2, pp. 250–273, 1995.

[26] R. Segala, "A compositional trace-based semantics for probabilistic automata," in *Proc. 6th Conference on Concurrency Theory*, ser. LNCS, vol. 962. Springer, 1995, pp. 234–248.

[27] N. A. Lynch, R. Segala, and F. W. Vaandrager, "Observing branching structure through probabilistic contexts," *SIAM Journal of Computing*, vol. 37, no. 4, pp. 977–1013, 2007.

[28] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior*. Princeton University Press, 1947.

[29] L. de Alfaro, R. Majumdar, V. Raman, and M. Stoelinga, "Game refinement relations and metrics," *Logic Methods in Computer Science*, vol. 4, no. 3:7, pp. 1–28, 2008.

[30] K. Chatterjee, L. de Alfaro, R. Majumdar, and V. Raman, "Algorithms for game metrics (full version)," *Logical Methods in Computer Science*, vol. 6, no. 3:13, pp. 1–27, 2010.

[31] M. Kwiatkowska, D. Parker, and C. Wiltsche, "PRISM-games: Verification and strategy synthesis for stochastic multi-player games with multiple objectives," *International Journal on Software Tools for Technology Transfer*, vol. 20, no. 2, pp. 195–210, 2018.

[32] L. de Alfaro and R. Majumdar, "Quantitative solution of omega-regular games," *Journal of Computer and System Sciences*, vol. 68, no. 2, pp. 374–397, 2004.