

p-Automata: New Foundations for Discrete-Time Probabilistic Verification

Michael Huth, Nir Piterman, and Daniel Wagner
Department of Computing, Imperial College London, London, UK

Abstract

We develop a new approach to probabilistic verification by adapting notions and techniques from alternating tree automata to the realm of Markov chains. The resulting p-automata determine languages of Markov chains which are proved to be closed under Boolean operations, to subsume bisimulation equivalence classes of Markov chains, and to subsume the set of models of any PCTL formula.

Our acceptance game for an input Markov chain to a p-automaton is shown to be well-defined and to be in EXPTIME in general; but its complexity is that of PCTL model checking for automata that represent PCTL formulas. We also derive a notion of simulation between p-automata that approximates language containment in EXPTIME.

These foundations therefore enable abstraction-based probabilistic model checking for probabilistic specifications that subsume Markov chains, and LTL and CTL like logics.*

1 Introduction

Markov chains are a very important modeling formalism in many areas of science. In computing, Markov chains form the basis of central techniques such as performance modeling, and the design and correctness of randomized algorithms used in security and communication protocols. Recognizing this prominent role of Markov chains, the formal-methods community has devoted significant attention to these models, e.g., in developing model checking for *qualitative* [15, 5, 23] and *quantitative* [1] properties, logics for reasoning about Markov chains [14, 19], and probabilistic simulation and bisimulation [20, 19]. Model-checking tools such as PRISM [12] and LiQuor [3] support such reasoning about Markov chains and have users in many fields of computer science and beyond.

In the non-probabilistic setting, the automata-theoretic approach to verification unifies such reasoning support for systems modeled as Kripke structures. Automata furnish the foundations for reasoning about such models through logic, model checking, synthesis, and abstraction. Alternating tree automata [10] were introduced to prove the decidability of satisfiability for monadic, second-order logic and they provide a unifying framework for all branching-

time temporal logics such as μ -calculus, CTL, and CTL*. Of particular interest to us is the result that alternating tree automata afford a complete framework for abstraction with respect to branching-time logic [6, 7]. Thus, in this context, alternating automata form the right basis for abstraction, the technique that makes model checking scale to realistic designs in the hardware and software industry. For Markov chains, their aforementioned techniques lack such a unifying framework and the quest for robust notions of abstraction is an active line of research. Here, we define p-automata and show that they render such a framework.

A p-automaton reads a Markov chain as input and either accepts or rejects it. PCTL formulas, the de-facto standard for model checking Markov chains, can be expressed as p-automata, and PCTL model checking can be reduced to deciding the acceptance of Markov chains by p-automata. Similarly, probabilistic versions of LTL, CTL*, or probabilistic ω -regular extensions of these logics one might wish to develop, can also be expressed as p-automata. Further, one can embed a Markov chain as a p-automaton accepting the language of Markov chains that are bisimilar to it.

The definition of p-automata is motivated by PCTL and alternating tree automata: it combines the rich combinatorial structure of alternating automata with PCTL's ability to quantify the probabilities of regular sets of paths. Much like alternating tree automata, whose acceptance of Kripke structures is determined by solving games (cf. [10]), acceptance by p-automata is determined by solving stochastic games. We show that acceptance of finite Markov chains can be determined in exponential time. But the acceptance games for p-automata that arise from PCTL are simpler, their complexity matching that of PCTL model checking.

We show that the languages of p-automata are closed under Boolean operations. Being more expressive than PCTL, emptiness of p-automata generalizes the long-standing open problem of PCTL satisfiability, and is here left open. We also define simulation of p-automata, show that simulation approximates language containment, and that this approximation is exact for p-automata arising from Markov chains.

We also introduce a new probabilistic separation operator, written $*$ in the paper, that decomposes the witness path set for a probability threshold into disjoint subsets. This

novel operator for probabilistic specifications leads, e.g., to succinct p-automata whose language is the bisimulation equivalence class of a Markov chain. Use of this operator, however, has a certain price in the complexity of the resulting acceptance games.

We show that the framework of p-automata constitutes the first complete abstraction framework for PCTL model checking on Markov chains: if an infinite-state Markov chain satisfies a PCTL formula, there is a finite p-automaton that abstracts (i.e. simulates) this Markov chain and whose language is contained in that of the formula. Thus, p-automata are a suitable back end for future counter-example guided abstraction refinement of PCTL model checking.

Our framework also enables extensions to PCTL, similar to those of LTL developed for hardware model checking. Our framework of p-automata further suggests a new approach to understanding the open problem of decidability of PCTL satisfiability: one can mimick the algorithms for checking emptiness of alternating tree automata and solving satisfiability of monadic second-order logic, μ -calculus, CTL*, and dynamic logic by trying to define a suitable notion of non-deterministic p-automata (for which non-emptiness would be decidable for standard reasons) and to show that p-automata can be converted into such a form.

In Section 2 we fix notation and recall needed concepts. Our p-automata are introduced in Section 3, their acceptance games defined in Section 4, and our expressiveness results featured in Section 5. Simulation and its salient properties are presented in Section 6. In Section 7 we discuss related and future work. Section 8 contains our conclusions.

2 Background

A *countable labeled Markov chain* M over set of atomic propositions $\mathbb{A}\mathbb{P}$ is a tuple (S, P, L, s^{in}) , where S is a countable set of *locations*, $P: S \times S \rightarrow [0, 1]$ a stochastic matrix, $s^{\text{in}} \in S$ the *initial* location, and $L: S \rightarrow 2^{\mathbb{A}\mathbb{P}}$ a *labeling function* with $L(s)$ the set of propositions true in location s . Let $\text{succ}(s)$ be the set $\{s' \in S \mid P(s, s') > 0\}$ of *successors* of s . All Markov chains are assumed to be *finitely branching*, i.e. $\text{succ}(s)$ is finite for all $s \in S$. We write $\text{MC}_{\mathbb{A}\mathbb{P}}$ for the set of all (finitely branching) Markov chains over $\mathbb{A}\mathbb{P}$. A *path* π from location s in M is an infinite sequence of locations $s_0 s_1 \dots$ with $s_0 = s$ and $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$. For $Y \subseteq S$, let $P(s, Y)$ abbreviate $\sum_{s' \in Y} P(s, s')$.

For Markov chain $M = (S, P, L, s^{\text{in}})$, a *bisimulation* [20] is an equivalence relation $H \subseteq S \times S$ such that $(s, s') \in H$ implies (i) $L(s) = L(s')$ and (ii) $P(s, C) = P(s', C)$ for all equivalence classes $C \in S/H$. The union of all bisimulations for M is the greatest bisimulation \sim ; locations s and s' are called *bisimilar* iff $s \sim s'$. This definition extends to Markov chains M_1 and M_2 by considering bisimilarity of their initial locations in the disjoint union of M_1 and M_2 .

Without loss of generality [8], one may define the prob-

$\phi, \psi ::=$	<i>PCTL formulas</i>	$\alpha ::=$	<i>Path formulas</i>
$\mathbf{a}, \neg \mathbf{a}$	Atom	$\mathbf{X} \phi$	Next
$\phi \wedge \psi$	Conjunction	$\phi \cup \psi$	Until
$\phi \vee \psi$	Disjunction	$\phi \mathbf{W} \psi$	Weak Until
$[\alpha]_{\bowtie p}$	Path Probability		

Figure 1. Syntax of PCTL, where $\mathbf{a} \in \mathbb{A}\mathbb{P}$, $p \in [0, 1]$, and $\bowtie \in \{>, \geq\}$

abilistic temporal logic PCTL [11] in “Greater Than Negation Normal Form”: only propositions can be negated and probabilistic bounds are either \geq or $>$ – see Fig. 1.

Our semantics of PCTL is as in [11]: path formulas α are interpreted as predicates over paths in M , and wrap PCTL formulas into “LTL” operators for Next, (strong) Until, and Weak Until. The semantics $\|\phi\| \subseteq S$ of PCTL formula ϕ lifts path formulas to state formulas: $s \in \|\alpha\|_{\bowtie p}$ iff $\text{Prob}_M(s, \alpha)$, the probability of the measurable set [17] $\text{Path}(s, \alpha)$ of paths $ss_1s_2\dots$ in M with $ss_1s_2\dots \models \alpha$, satisfies $\bowtie p$. M satisfies ϕ , denoted $M \models \phi$, if $s^{\text{in}} \in \|\phi\|$.

Weak Games. A tuple $G = ((V, E), (V_0, V_1, V_p), \kappa, \alpha)$ is a *stochastic weak game* if (V, E) is a directed graph, (V_0, V_1, V_p) a partition of V , function κ associates with every $v \in V_p$ a distribution $\kappa(v)$ of mass 1 over $E(v) = \{v' \mid (v, v') \in E\}$ such that $(v, v') \in E$ iff $\kappa(v)(v') \neq 0$; we write $\kappa(v, v')$ instead of $\kappa(v)(v')$. Set $\alpha \subseteq V$ is the winning condition. Set V_0 contains the Player 0 configurations, V_1 the Player 1 configurations, and V_p the probabilistic configurations of G . We work with *weak games*, i.e. for every maximal, strongly connected component (SCC) $V' \subseteq V$ in (V, E) either $V' \subseteq \alpha$ or $V' \cap \alpha = \{\}$. If $V_p = \{\}$, we call G simply a *weak game*. Markov chains can be thought of as stochastic weak games where $V_0 = V_1 = \{\}$ and $\alpha = V$.

A play in G is a maximal sequence $v_0v_1\dots$ of configurations with $(v_i, v_{i+1}) \in E$ for all $i \in \mathbb{N}$. A play is winning for Player 0 if it is finite and ends in a Player 1 configuration, or if it is infinite and ends in a suffix of states in α . Otherwise, that play is winning for Player 1. A (pure memoryless) strategy for Player 0 is a function $\sigma: V_0 \rightarrow V$ with $(v, \sigma(v)) \in E$ for all $v \in V_0$. Play $v_0v_1\dots$ is consistent with strategy σ if $v_{i+1} = \sigma(v_i)$ whenever $v_i \in V_0$. Strategies for Player 1 are defined analogously. Let Σ (resp. Π) be the set of all strategies for Player 0 (resp. Player 1).

Each $(\sigma, \pi) \in \Sigma \times \Pi$ from game G determines a Markov chain $M^{\sigma, \pi}$ (with sinks for dead-ends in G) whose paths are plays in G consistent with σ and π . The set of plays from $v \in V$ that Player 0 wins is measurable in $M^{\sigma, \pi}$. Let $\text{val}_0^{\sigma, \pi}(v)$ be that measure, and $\text{val}_1^{\sigma, \pi}(v) = 1 - \text{val}_0^{\sigma, \pi}(v)$. Then $\text{val}_0(v) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \text{val}_0^{\sigma, \pi}(v) \in [0, 1]$ and $\text{val}_1(v) = \sup_{\pi \in \Pi} \inf_{\sigma \in \Sigma} \text{val}_1^{\sigma, \pi}(v) \in [0, 1]$ are the game values. Strategies that achieve these values are *optimal*.

Theorem 1 [4, 22] *Let $G = ((V, \cdot), \dots)$ be a stochastic weak game and $v \in V$. Then $\text{val}_0(v) + \text{val}_1(v) = 1$. If*

$\llbracket Q \rrbracket_{>}$	$= \{\llbracket q \rrbracket_{\bowtie p} \mid q \in Q, \bowtie \in \{\geq, >\}, p \in [0, 1]\}$
$\llbracket Q \rrbracket^*$	$= \{*(t_1, \dots, t_n) \mid n \in \mathbb{N}, \forall i: t_i \in \llbracket Q \rrbracket_{>}\}$
$\llbracket Q \rrbracket^{\forall}$	$= \{\forall(t_1, \dots, t_n) \mid n \in \mathbb{N}, \forall i: t_i \in \llbracket Q \rrbracket_{>}\}$
$\llbracket Q \rrbracket$	$= \llbracket Q \rrbracket^* \cup \llbracket Q \rrbracket^{\forall}$

Figure 2. Derived term sets for set Q

G is finite, $\text{val}_0(v)$ is computable in $NP \cap \text{co-NP}$, and optimal strategies exist for both players. If G is a weak game, $\text{val}_0(v)$ is in $\{0, 1\}$ and linear-time computable.

One can generalize these results to the setting in which some configurations have pre-seeded game values (in $[0, 1]$ for stochastic weak games, and in $\{0, 1\}$ for weak games).

3 Uniform Weak p-Automata

We introduce p-automata and their uniform weak variant. Traditional probabilistic automata [21] map an input to a probability of accepting it. Such an automaton A then gives rise to a mapping from words to the probability of their acceptance, and thus to probabilistic languages \mathcal{L}_μ of words which are accepted with probability above a fixed threshold μ . In contrast, our p-automata either accept or reject an entire Markov chain. In particular, a p-automaton determines a language of Markov chains.

We assume familiarity with basic notions of trees and (alternating) tree automata. For set T , let $B^+(T)$ be the set of positive Boolean formulas generated from elements $t \in T$, constants tt and ff, and disjunctions and conjunctions:

$$\varphi, \psi ::= t \mid \text{tt} \mid \text{ff} \mid \varphi \vee \psi \mid \varphi \wedge \psi \quad (1)$$

Formulas in $B^+(T)$ are finite even if T is not.

For set Q , the set of states of a p-automaton, we define term sets in Fig. 2. This uses n -ary operators $*_n$ and \forall_n for every $n \in \mathbb{N}$, which we write as $*$ and \forall throughout as n will be clear from context. Intuitively, a state $q \in Q$ of a p-automaton and its transition structure model a probabilistic path set. So $\llbracket q \rrbracket_{\bowtie p}$ holds in location s if the measure of paths that begin in s and satisfy q is $\bowtie p$. Now, $*(\llbracket q_1 \rrbracket_{>p_1}, \llbracket q_2 \rrbracket_{\geq p_2})$, e.g., means q_1 and q_2 hold with probability greater than p_1 and greater than or equal to p_2 , respectively; and that the sets supplying these probabilities are disjoint. Dually, $\forall(\llbracket q_1 \rrbracket_{\geq p_1}, \llbracket q_2 \rrbracket_{\geq p_2})$ means that either (i) there is $i \in \{1, 2\}$ such that q_i holds with probability at least p_i or (ii) the intersection of q_1 and q_2 holds with probability at least $\max(p_1 + p_2 - 1, 0)$. So $*$ and \forall model a “disjoint and” and “intersecting or” operator, respectively. We may write $\llbracket q \rrbracket_{\bowtie p}$ for $*(\llbracket q \rrbracket_{\bowtie p})$, and similarly for \forall .

An element of $Q \cup \llbracket Q \rrbracket$ is therefore either a state of the p-automaton, a $*$ composition of terms $\llbracket q_i \rrbracket_{\bowtie p_i}$, or a \forall composition of such terms. Given $\varphi \in B^+(Q \cup \llbracket Q \rrbracket)$, its closure $\text{cl}(\varphi)$ is the set of all subformulas of φ according to (1). In particular, $*(t_1, t_2) \in \text{cl}(\varphi)$ does not imply $t_1, t_2 \in \text{cl}(\varphi)$. For a set Φ of formulas, let $\text{cl}(\Phi) = \bigcup_{\varphi \in \Phi} \text{cl}(\varphi)$.




Figure 3. (a) Graph G_A of automaton A from Example 1 and (b) a Markov chain M

Definition 1 A p -automaton A is a tuple $\langle \Sigma, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$, where Σ is a finite input alphabet, Q a set of states (not necessarily finite), $\delta: Q \times \Sigma \rightarrow B^+(Q \cup \llbracket Q \rrbracket)$ the transition function, $\varphi^{\text{in}} \in B^+(Q \cup \llbracket Q \rrbracket)$ the initial condition, and $\alpha \subseteq Q$ an acceptance condition.

Below, p-automata have states, Markov chains have locations, and weak stochastic games have configurations.

Example 1 Let $A = \langle 2^{\{a,b\}}, \{q_1, q_2\}, \delta, \llbracket q_1 \rrbracket_{\geq 0.5}, \{q_2\} \rangle$ be a p-automaton where δ is defined by

$$\begin{aligned} \delta(q_1, \{a, b\}) &= \delta(q_1, \{a\}) = q_1 \vee \llbracket q_2 \rrbracket_{\geq 0.5} \\ \delta(q_2, \{b\}) &= \delta(q_2, \{a, b\}) = \llbracket q_2 \rrbracket_{\geq 0.5} \\ \delta(q_1, \{\}) &= \delta(q_1, \{b\}) = \delta(q_2, \{\}) = \delta(q_2, \{a\}) = \text{ff} \end{aligned}$$

Term $\llbracket q_2 \rrbracket_{\geq 0.5}$ represents the recursive property ϕ , that atomic proposition b holds at the location presently read by q_2 , and that ϕ will hold with probability at least 0.5 in the next locations. State q_1 asserts that it is possible to get to a location that satisfies $\llbracket q_2 \rrbracket_{\geq 0.5}$ along a path that satisfies atomic proposition a . The initial condition $\llbracket q_1 \rrbracket_{\geq 0.5}$ means the set of paths satisfying a $\cup \phi$ has probability at least 0.5.

In order to be able to decide acceptance of input for p-automata through the solution of weak stochastic games, we restrict the cycles in the transition graph of p-automata. In doing so, we differentiate states q' appearing within a term in $\llbracket Q \rrbracket$ (bounded transition) from q' appearing “free” in the transition of a state q (unbounded transition). In this way, a p-automaton $A = \langle \Sigma, Q, \delta, \dots \rangle$ determines a labeled, directed graph $G_A = \langle Q', E, E_b, E_u \rangle$:

$$\begin{aligned} Q' &= Q \cup \text{cl}(\delta(Q, \Sigma)) \\ E &= \{(\varphi_1 \wedge \varphi_2, \varphi_i), (\varphi_1 \vee \varphi_2, \varphi_i) \mid \varphi_i \in Q' \setminus Q, \\ &\quad i \in \{1, 2\}\} \cup \{(q, \delta(q, \sigma)) \mid q \in Q, \sigma \in \Sigma\} \\ E_u &= \{(\varphi \wedge q, q), (q \wedge \varphi, q), (\varphi \vee q, q), (q \vee \varphi, q) \mid \\ &\quad \varphi \in Q', q \in Q\} \\ E_b &= \{(\varphi, q) \mid \varphi \in \llbracket Q \rrbracket \text{ and } q \in \text{gs}(\varphi)\} \end{aligned}$$

where $\text{gs}(\varphi)$ is the set of guarded states of φ : all $q \in Q$ occurring in some term in φ . Elements $(\varphi, q) \in E_u$ are unbounded transitions; elements $(\varphi, q) \in E_b$ are bounded transitions; and elements of E are called simple transitions. We mark $(\varphi, q) \in E_b$ with $*$ (and respectively, with \forall) if $\varphi \in \llbracket Q \rrbracket^*$ (respectively, $\varphi \in \llbracket Q \rrbracket^{\forall}$). Note that E, E_u , and E_b are pairwise disjoint. Let $\varphi \preceq_A \tilde{\varphi}$ iff there is a finite path from φ to $\tilde{\varphi}$ in $E \cup E_b \cup E_u$. Let \equiv be $\preceq_A \cap \preceq_A^{-1}$ and (φ) the equivalence class of φ with respect to \equiv . Each (φ) is an SCC in the directed graph G_A .

Definition 2 A p -automaton A is called uniform if:

- For each cycle in G_A , its set of transitions is either in $E \cup E_b$ or in $E \cup E_u$.
- For each cycle in $\langle Q, E \cup E_b \rangle$, its set of markings is either $\{\}$, $\{*\}$ or $\{\checkmark\}$, and so cannot be $\{*, \checkmark\}$.
- There are only finitely many equivalence classes $((\varphi))$ with $\varphi \in Q \cup \text{cl}(\delta(Q, \Sigma))$.

A (not necessarily uniform) p -automaton A is called weak if for all $q \in Q$, either $((q)) \cap Q \subseteq \alpha$ or $((q)) \cap \alpha = \{\}$.

Then, A is uniform, if the full subgraph of every equivalence class in \preceq_A contains only one type of non-simple transitions and at most one kind of marking $*$ or \checkmark . Also, all states $q' \in Q$ or formulas φ occurring in $\delta(q, \sigma)$ for some $q \in Q$ and $\sigma \in \Sigma$ can be classified as unbounded, bounded with $*$, bounded with \checkmark , or simple – according to SCC $((q))$.

Example 2 Figure 3(a) depicts the graph G_A for A of Example 1. p -Automaton A is uniform: $((q_1)) = \{q_1, q_1 \vee [q_2]_{\geq 0.5}\}$ and $((q_2)) = \{q_2, [q_2]_{\geq 0.5}\}$; in $((q_1))$ there are no bounded edges, in $((q_2))$ there are no unbounded edges; and G_A has no markings for $*$ or \checkmark . The SCC $(([q_1]_{\geq 0.5})) = \{[q_1]_{\geq 0.5}\}$ is trivial. In addition, A is weak as $\alpha = \{q_2\}$.

Intuitively, the cycles in the structure of a uniform p -automaton A take either no bounded edges or no unbounded edges, and cycles that take bounded edges don't have both markings $*$ and \checkmark . Subsequently, all p -automata are uniform weak unless mentioned otherwise. Uniformity allows to define acceptance of input for p -automata through the solution of weak stochastic games. But, a more relaxed notion of uniformity is what really drives the proof of well-definedness: that any chain in the partial order on SCCs on the graph of a p -automaton has only finitely many alternations between bounded and unbounded SCCs.

The requirement of weakness is made merely to simplify the presentation. Using a parity condition instead, e.g., would still allow us to decide acceptance of input for uniform p -automata, by solving stochastic parity games.

4 Acceptance Games

For any $\mathbb{A}\mathbb{P}$, p -automata $A = \langle 2^{\mathbb{A}\mathbb{P}}, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$ have $\text{MC}_{\mathbb{A}\mathbb{P}}$ as set of inputs. For $M = (S, P, L, s^{\text{in}}) \in \text{MC}_{\mathbb{A}\mathbb{P}}$, we exploit the uniform structure of A to reduce the decision of whether A accepts M to solving a sequence of weak games and stochastic weak games. Intuitively, unbounded cycles in G_A correspond to weak stochastic games and bounded cycles to weak games. The weak acceptance of A implies that these games are weak. Then the language of A is $\mathcal{L}(A) = \{M \in \text{MC}_{\mathbb{A}\mathbb{P}} \mid A \text{ accepts } M\}$. Just like acceptance games of alternating tree automata, all states of A and all subformulas appearing in its transitions form part of acceptance games. For A as above, let $T = Q \cup \text{cl}(\delta(Q, 2^{\mathbb{A}\mathbb{P}}))$.

Finite partial order $(T/\equiv, \leq_A)$ has set $\{((t)) \mid t \in T\}$ ordered by $((\tilde{t})) \leq_A ((t))$ iff $\tilde{t} \preceq_A t$. For M as above, each $((t))$ determines a game $G_{M,((t))} = ((V, E), (V_0, V_1, V_p), \kappa, \tilde{\alpha})$. Most of its configurations are in $S \times T$. The construction is such that $(s^{\text{in}}, \varphi^{\text{in}})$ occurs as configuration in exactly one of these games $G_{M,((t))}$, and $\text{val}(s^{\text{in}}, \varphi^{\text{in}}) \in [0, 1]$. Then A accepts M iff $\text{val}(s^{\text{in}}, \varphi^{\text{in}}) = 1$. We define these games as follows. Since A is uniform weak, each $((t))$ is of one of three types and each type determines a weak game or weak stochastic game as detailed in the three cases below. All game values already computed for games $G_{M,((\tilde{t}))}$ of SCCs $((\tilde{t}))$ higher up with respect to \leq_A (i.e. by induction) are used as pre-seeded values in $G_{M,((t))}$. Below, we write $\text{val}(s, \varphi) = \perp$ for configurations (s, φ) in $G_{M,((t))}$ whose game value has not been pre-seeded.

Case 1: For an SCC $((t))$ such that none of its transitions are in E_b , game $G_{M,((t))}$ is a stochastic weak game with

$$\begin{aligned} V &= \{(s, \tilde{t}) \mid s \in S \text{ and } \tilde{t} \preceq_A \tilde{t}\} & V_0 &= \{(s, \varphi_1 \vee \varphi_2) \in V\} \\ V_1 &= \{(s, \varphi_1 \wedge \varphi_2) \in V\} & V_p &= (S \times Q) \cap V \\ \kappa((s, q), (s', \delta(q, L(s)))) &= P(s, s') & \tilde{\alpha} &= \{\} \text{ or } V \end{aligned}$$

$$\begin{aligned} E &= \{((s, \varphi_1 \wedge \varphi_2), (s, \varphi_i)) \in V \times V \mid i \in \{1, 2\}\} \cup \\ &\quad \{((s, \varphi_1 \vee \varphi_2), (s, \varphi_i)) \in V \times V \mid i \in \{1, 2\}\} \cup \\ &\quad \{((s, q), (s', \delta(q, L(s)))) \in V \times V \mid P(s, s') > 0\} \end{aligned}$$

where $\tilde{\alpha}$ equals V iff some state q in $((t))$ is in α . By Theorem 1, for every configuration $c \in V$ we have $\text{val}_0(c) \in [0, 1]$. We set $\text{val}(c) = \text{val}_0(c)$.

Case 2: Let $((t))$ be an SCC such that none of its transitions are in E_u and none have \checkmark markings. For each formula $\varphi \in ((t)) \cap [Q]^*$ of form $*([q_1]_{\bowtie_1 p_1}, \dots, [q_n]_{\bowtie_n p_n})$ we define, for each $s \in S$, sets $V_0^{s, \varphi}$, $V_1^{s, \varphi}$, and $E^{s, \varphi}$. Then

$$\begin{aligned} V_0 &= \bigcup_{s, \varphi} V_0^{s, \varphi} & V_1 &= \bigcup_{s, \varphi} V_1^{s, \varphi} & V_p &= \{\} \\ E &= \bigcup_{s, \varphi} E^{s, \varphi} & \tilde{\alpha} &= \{\} \text{ or } V \end{aligned}$$

defines the weak game $G_{M,((t))}$ – where $\tilde{\alpha}$ is V if some $q \in ((t))$ is in α , and is empty otherwise. It remains to define $V_0^{s, \varphi}$, $V_1^{s, \varphi}$, and $E^{s, \varphi}$, for which we use pre-seeded values $\text{val}(s, \tilde{t})$ for all $s \in S$ and all $\tilde{t} \notin ((t))$ with $((\tilde{t})) \leq_A ((t))$.

As $\text{succ}(s)$ and $\delta(q_i, L(s))$ are finite, so are

$$\begin{aligned} R_{s, \varphi} &= \bigcup_{i=1}^n \{(s', \varphi') \mid s' \in \text{succ}(s), \varphi' \in \text{cl}(\delta(q_i, L(s)))\} \\ \text{Val}_{s, \varphi} &= \{0, 1, \text{val}(s', \varphi') \mid (s', \varphi') \in R_{s, \varphi}, \text{val}(s', \varphi') \neq \perp\} \end{aligned}$$

Intuitively, $R_{s, \varphi}$ is the set of configurations reachable from (s, φ) using one transition of a state in φ . Thus, s' are the successors of s and φ' are subformulas of $\delta(q_i, L(s))$. Set $\text{Val}_{s, \varphi}$ includes 0, 1, and values of configurations in $R_{s, \varphi}$. In game $G_{M,((t))}$, a play proceeding from (s, φ) reaches either a configuration whose value is in $\text{Val}_{s, \varphi}$ or a configuration (s, ψ) for $\psi \in ((t))$. Sets $V_0^{s, \varphi}$, $V_1^{s, \varphi}$, and $E^{s, \varphi}$ are defined in Fig. 4, the definition of $\mathcal{F}_{s, \varphi}^*$ is deferred for now.

The intuition behind this weak game is as follows: Configuration (s, φ) means that the transition of each q_i holds with probability $\bowtie_i p_i$ where the sets X_i measured by these probabilities are pairwise disjoint. In order to check that, given configuration (s, φ) , Player 0 chooses a function $f \in \mathcal{F}_{s, \varphi}^*$ that associates with location $s' \in \text{succ}(s)$ and state q_i the value Player 0 can achieve playing from $(s', \delta(q_i, L(s)))$. The play continues with Player 1 choosing a successor s' of s and a state q_i , and the play then reaches configuration $(s', \delta(q_i, L(s)), f(i, s'))$. From such value-annotated configurations, Player 0 and Player 1 choose successors according to the usual resolution of \vee and \wedge :

- In a configuration for which the value v was already determined, either $f(i, s') \bowtie_i v$, i.e. Player 0 achieved the promised value and wins immediately; or Player 0 failed to achieve the promised value and loses immediately.
- Otherwise, the play ends up in another configuration of the form (s', φ') for $\varphi' \in \llbracket Q \rrbracket^*$ and the play continues and ignores the value $f(i, s')$ (as obviously $f(i, s') \leq 1$). If the play continues ad infinitum, the winner is determined according to acceptance condition $\tilde{\alpha}$.

We now define the function space $\mathcal{F}_{s, \varphi}^*$ that captures terms built from the separation operator $*$. For $n \in \mathbb{N}$, let $[n] = \{1, \dots, n\}$. Throughout, let $X \rightarrow Y$ be the set of total functions from set X to set Y . Let $\mathcal{F}_{s, \varphi}$ be $[n] \times \text{succ}(s) \rightarrow \text{Val}_{s, \varphi}$, the set of functions from pairs consisting of ‘sub-stars’ of φ and successors of s to values in $\text{Val}_{s, \varphi}$. Also, any $f \in \mathcal{F}_{s, \varphi}$ is *disjoint* if there are $\{a_{i, s'} \in [0, 1] \mid i \in [n] \text{ and } s' \in \text{succ}(s)\}$ such that (i) $\sum_{s' \in \text{succ}(s)} a_{i, s'} f(i, s') P(s, s') \bowtie_i p_i$ for all $i \in [n]$ and (ii) $\sum_{i \in [n]} a_{i, s'} = 1$ for all $s' \in \text{succ}(s)$. Intuitively, a function $f \in \mathcal{F}_{s, \varphi}$ associates with q_1, \dots, q_n and s' the value that Player 0 can achieve from configuration $(s', \delta(q_i, L(s)))$. Values in $\text{Val}_{s, \varphi}$ suffice, as no others are directly reachable. We call f “disjoint”, as all the requirements from the different q_i ’s can be achieved using a partition (realized by the existence of the above $a_{i, s}$) of the probability of all successors. Let $\mathcal{F}_{s, \varphi}^*$ be the set of disjoint functions. By Theorem 1, V partitions into winning regions W_0 and W_1 of configurations for Player 0 and Player 1, respectively. We set $\text{val}(c) = 1$ for $c \in W_0$ and $\text{val}(c) = 0$ for $c \in W_1$.

Case 3: Finally, let $((t))$ be an SCC such that none of its transitions is in E_u and none has $*$ markings. For formulas $\varphi \in ((t)) \cap \llbracket Q \rrbracket^{\bowtie}$ of form $\bowtie(\llbracket q_1 \rrbracket \bowtie_1 p_1, \dots, \llbracket q_n \rrbracket \bowtie_n p_n)$ we reuse the definitions of $R_{s, \varphi}$, $\text{Val}_{s, \varphi}$, and $\mathcal{F}_{s, \varphi}$. Weak game $G_{M, ((t))}$ is defined as in Case 2. Sets $V_0^{s, \varphi}$, $V_1^{s, \varphi}$, and $E^{s, \varphi}$ are defined as in Fig. 4, except that functions f don’t range over $\mathcal{F}_{s, \varphi}^*$ but now range over $\mathcal{F}_{s, \varphi}^{\bowtie}$, the set of intersecting functions and the dual of $\mathcal{F}_{s, \varphi}^*$ of Case 2: function $f \in \mathcal{F}_{s, \varphi}$ is *intersecting* if for all sets $\{a_{i, s'} \in [0, 1] \mid i \in [n] \text{ and } s' \in \text{succ}(s)\}$ either (i) there is $i \in [n]$ with $\sum_{s' \in \text{succ}(s)} a_{i, s'} f(i, s') P(s, s') \bowtie_i p_i$ or (ii) there is $s' \in \text{succ}(s)$ with $\sum_{i \in [n]} a_{i, s'} \neq 1$.

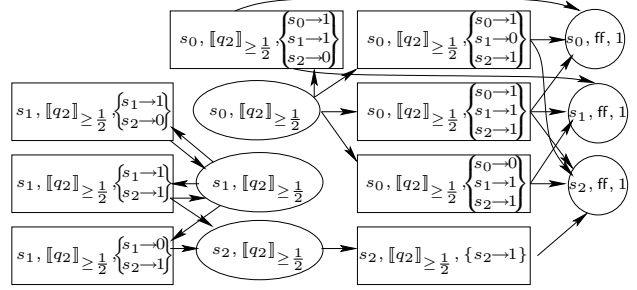


Figure 5. Case 3 of acceptance game

As in Case 2, we say that wins for Player 0 have value 1, and wins for Player 1 have value 0.

The intuition for this weak game is verbatim that of the weak game in Case 2, except that Player 0 chooses a function f that is in $\mathcal{F}_{s, \varphi}^{\bowtie}$ instead of in $\mathcal{F}_{s, \varphi}^*$.

We point out that when n above is 1, i.e. in handling $\varphi = \llbracket q_1 \rrbracket \bowtie_1 p_1$, the definition of $*$ and \bowtie coincide. Indeed, there is then exactly one option for choosing $\{a_{1, s'} \mid s' \in \text{succ}(s)\}$: the value $a_{1, s'}$ has to be 1 for all $s' \in \text{succ}(s)$. This justifies dropping the $*$ or \bowtie when applied to one operand.

Trivial SCCs $((t))$, for which $((t)), E \cup E_b \cup E_u$ is cycle-free, may satisfy more than one of these three cases. This ambiguity is unproblematic as game values in $G_{M, ((t))}$ are then determined via propagation of pre-seeded game values.

Example 3 We verify that $M \in \mathcal{L}(A)$ for A from Example 1 and M from Fig. 3(b), where locations are labeled by propositions – e.g., $L(s_0) = \{a\}$. The weak game of SCC $((q_2))$, shown in Fig. 5, has only accepting configurations or dead ends. So Player 0 wins only $(s_1, \llbracket q_2 \rrbracket \geq 0.5)$ and $(s_1, \llbracket q_2 \rrbracket \geq 0.5, \{s_1 \rightarrow 1, s_2 \rightarrow 0\})$, and loses everywhere else.

The stochastic weak game $G_{M, ((q_1))}$ for the SCC $((q_1))$, shown in Fig. 6, depicts stochastic configurations with a diamond and configurations from other SCCs are put into hexagons (with the hexagon labeled $(s_1, \llbracket q_2 \rrbracket \geq 0.5)$ having value 1 and all others having value 0). As none of the configurations are accepting, Player 0 can only win by reaching optimal hexagons. Hexagon $(s_1, \llbracket q_2 \rrbracket \geq 0.5)$ has value 1 and is the optimal choice for Player 0 from configuration $(s_1, q_1 \vee \llbracket q_2 \rrbracket \geq 0.5)$. Player 0 configuration $(s_2, q_1 \vee \llbracket q_2 \rrbracket \geq 0.5)$ has value 0. So the value for Player 0 of diamond configuration (s_0, q_1) is 0.5. Initial configuration $(s_0, \llbracket q_1 \rrbracket \geq 0.5)$ makes up a trivial bounded SCC (e.g. Case 2), so its value equals 1 as $\frac{1}{3} \text{val}(s_0, q_1 \vee \llbracket q_2 \rrbracket \geq 0.5) + \frac{1}{3} \text{val}(s_1, q_1 \vee \llbracket q_2 \rrbracket \geq 0.5) + \frac{1}{3} \text{val}(s_2, q_1 \vee \llbracket q_2 \rrbracket \geq 0.5)$ is 0.5. Therefore, $M \in \mathcal{L}(A)$.

Theorem 2 Given a p -automaton $A = \langle 2^{\text{AP}}, \dots \rangle$, its language $\mathcal{L}(A)$ is well defined. If A and $M \in \text{MC}_{\text{AP}}$ are finite, $M \in \mathcal{L}(A)$ can be decided in EXPTIME.

For finite Markov chain M and p -automaton A with non-trivial, bounded SCCs, checking acceptance $M \in \mathcal{L}(A)$ is

$V_0^{s,\varphi} =$	$\{(s, \varphi)\} \cup \{(s', \varphi', v) \mid s' \in \text{succ}(s), \varphi' \in R_{s,\varphi}, v \in \text{Val}_{s,\varphi}, \text{val}(s', \varphi') \neq \perp, \text{ and } \text{val}(s', \varphi') < v\}$	\cup
	$\{(s', \varphi_1 \vee \varphi_2, v) \mid s' \in \text{succ}(s), \varphi_1 \vee \varphi_2 \in R_{s,\varphi}, v \in \text{Val}_{s,\varphi}, \text{ and } \text{val}(s', \varphi_1 \vee \varphi_2) = \perp\}$	
$V_1^{s,\varphi} =$	$\{(s, \varphi, f) \mid f \in \mathcal{F}_{s,\varphi}^*\} \cup \{(s', \varphi', v) \mid s' \in \text{succ}(s), \varphi' \in R_{s,\varphi}, v \in \text{Val}_{s,\varphi}, \text{val}(s', \varphi') \neq \perp, \text{ and } \text{val}(s', \varphi') \geq v\}$	\cup
	$\{(s', \varphi_1 \wedge \varphi_2, v) \mid s' \in \text{succ}(s), \varphi_1 \wedge \varphi_2 \in R_{s,\varphi}, v \in \text{Val}_{s,\varphi}, \text{ and } \text{val}(s', \varphi_1 \wedge \varphi_2) = \perp\}$	
$E^{s,\varphi} =$	$\{((s, \varphi), (s, \varphi, f)) \mid f \in \mathcal{F}_{s,\varphi}^*\} \cup \{((s', \varphi', v), (s', \varphi')) \mid s' \in \text{succ}(s), \varphi' \in \llbracket Q \rrbracket, v \in \text{Val}_{s,\varphi}, \text{ and } \text{val}(s', \varphi') = \perp\}$	\cup
	$\{((s, \varphi, f), (s', \delta(q_i, L(s)), f(i, s'))) \mid s' \in \text{succ}(s), i \in [n], \text{ and } f(i, s') > 0\}$	\cup
	$\{((s', \varphi_1 \vee \varphi_2, v), (s', \varphi_i, v)) \mid s' \in \text{succ}(s), \varphi_1 \vee \varphi_2 \in R_{s,\varphi}, i \in \{1, 2\}, v \in \text{Val}_{s,\varphi}, \text{ and } \text{val}(s', \varphi_1 \vee \varphi_2) = \perp\}$	\cup
	$\{((s', \varphi_1 \wedge \varphi_2, v), (s', \varphi_i, v)) \mid s' \in \text{succ}(s), \varphi_1 \wedge \varphi_2 \in R_{s,\varphi}, i \in \{1, 2\}, v \in \text{Val}_{s,\varphi}, \text{ and } \text{val}(s', \varphi_1 \wedge \varphi_2) = \perp\}$	

Figure 4. Components of the game $G_{M,((t))}$, where $((t))$ does not contain \forall transitions

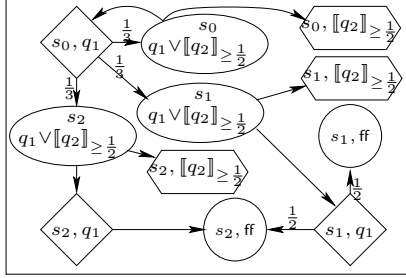


Figure 6. Case 1 of acceptance game

$\text{dual}(\forall(t_1, \dots, t_n)) =$	$*(\text{dual}(t_1), \dots, \text{dual}(t_n))$
$\text{dual}(*t_1, \dots, t_n) =$	$\forall(\text{dual}(t_1), \dots, \text{dual}(t_n))$
$\text{dual}(\varphi_1 \wedge \varphi_2) =$	$\text{dual}(\varphi_1) \vee \text{dual}(\varphi_2)$
$\text{dual}(\varphi_1 \vee \varphi_2) =$	$\text{dual}(\varphi_1) \wedge \text{dual}(\varphi_2)$
$\text{dual}(q) =$	\bar{q}
$\text{dual}(\bar{q}) =$	q
$\text{dual}(\llbracket q \rrbracket_{\geq p}) =$	$\llbracket \bar{q} \rrbracket_{\text{dual}(\geq p)}$
$\text{dual}(\geq p) =$	$> 1 - p$
$\text{dual}(> p) =$	$\geq 1 - p$

Figure 7. Definition of $\text{dual}(\varphi)$

exponential in the branching degree of M and in the branching degree of $*$ and \forall operators of A , but not in the number of states or locations. If A has only trivial bounded-SCCs, checking $M \in \mathcal{L}(A)$ reduces to solving a linear number of linear sized stochastic weak games.

5 Expressiveness of p-Automata

We show that the languages of p-automata are closed under Boolean operations and bisimulation, and emptiness and containment of languages are equi-solvable; that each Markov chain determines a p-automaton whose language is the bisimulation class of that Markov chain; and that each PCTL formula determines a p-automaton whose language consists of all Markov chains satisfying that formula.

Closure of Languages. It is routine to see that p-automata are closed under union and intersection. But they are also closed under complementation: Given a p-automaton $A = \langle \Sigma, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$, its dual $\text{dual}(A)$ is $\langle \Sigma, \bar{Q}, \bar{\delta}, \text{dual}(\varphi^{\text{in}}), Q \setminus \alpha \rangle$ with $\bar{Q} = \{\bar{q} \mid q \in Q\}$ and $\bar{\delta}(\bar{q}, \sigma) = \text{dual}(\delta(q, \sigma))$, where $\text{dual}(\varphi)$ is defined in Fig. 7. The structure of uniform weak p-automata ensures that

$\text{dual}(A)$ is also uniform weak. The languages of A and $\text{dual}(A)$ are complements.

Theorem 3 *Let A be a p-automaton with $\Sigma = 2^{\text{AP}}$. Then $\mathcal{L}(A) = \text{MC}_{\text{AP}} \setminus \mathcal{L}(\text{dual}(A))$.*

The key part of proving Theorem 3 is to show that, for all states q of A and all locations s of M , we have $\text{val}(s, q) = 1 - \text{val}(s, \bar{q})$ for the respective acceptance games.

Corollary 1 *Let $\Sigma = 2^{\text{AP}}$.*

- The set of languages accepted by p-automata with Σ is closed under Boolean operations.
- Language containment of p-automata with Σ reduces to language emptiness of such p-automata, and vice versa.

Languages of p-automata are closed under bisimulation.

Lemma 1 *For p-automaton $A = \langle 2^{\text{AP}}, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$ and $M_1, M_2 \in \text{MC}_{\text{AP}}$ with $M_1 \sim M_2$: $M_1 \in \mathcal{L}(A)$ iff $M_2 \in \mathcal{L}(A)$.*

To prove this, we use induction on the partial order on the SCCs in A to show that for all $t \in Q \cup \llbracket Q \rrbracket$ and for all locations s_1 in M_1 and locations s_2 in M_2 with $s_1 \sim s_2$ we have $\text{val}(s_1, t) = \text{val}(s_2, t)$.

Embedding of Markov Chains. A Markov chain $M = (S, P, L, s^{\text{in}}) \in \text{MC}_{\text{AP}}$ can be converted into a p-automaton $A_M = \langle 2^{\text{AP}}, Q, \delta, \varphi^{\text{in}}, \alpha \rangle$ whose language $\mathcal{L}(A_M)$ is the set of Markov chains bisimilar to M . The definition of A_M implicitly appeals in $*$ expressions to an enumeration of each set $\text{succ}(s')$:

$$\begin{aligned}
Q &= \{(s, s') \in S \times S \mid P(s, s') > 0\} \\
\delta((s, s'), L(s)) &= *(\llbracket (s', s'') \rrbracket_{\geq P(s', s'')} \mid s'' \in \text{succ}(s')) \\
\delta((s, s'), \sigma) &= \text{ff} \quad \text{if } \sigma \neq L(s) \\
\varphi^{\text{in}} &= *(\llbracket (s^{\text{in}}, s') \rrbracket_{\geq P(s^{\text{in}}, s')} \mid P(s^{\text{in}}, s') > 0) \\
\alpha &= Q
\end{aligned}$$

State (s, s') represents the transition from s to s' . Labels are compared for location s . Location s' is used to require that there are successors of probability at least $P(s, s')$. This p-automaton A_M has only bounded transitions and uses only the $*$ operator. In particular, it is uniform weak.

Theorem 4 *For any Markov chain $M \in \text{MC}_{\text{AP}}$, the language $\mathcal{L}(A_M)$ is the bisimulation equivalence class of M .*

By Lemma 1, one half of Theorem 4 follows from a proof that A_M accepts M . To show this, it suffices to demonstrate that Player 0 can infinitely often reach configurations of form $(s, *(\llbracket (s, s') \rrbracket_{\geq P(s, s')}))$ with $s' \in \text{succ}(s)$ for all locations s in M . For the other half, we use proof by contradiction: given M' with initial state t^{in} such that $M' \not\sim M$, we appeal to the usual partition-refinement algorithm to get a coarsest partition that witnesses $s^{\text{in}} \not\sim t^{\text{in}}$. That witnessing information can then be transformed into a winning strategy for Player 1 in the acceptance game for deciding $M' \in \mathcal{L}(A_M)$, and so $M' \notin \mathcal{L}(A_M)$ follows.

The construction of A_M for infinite Markov chains was the only reason why we allow p-automata with infinite state sets. Finite state sets suffice for embedding finite Markov chains. The construction of A_M was also our initial reason for introducing the $*$ and \forall operators. But we believe that the separation of concern expressed in these operators is useful in p-automata in general. The conjunctive operator $*$ used in the construction of A_M effectively hides an exponential blowup. If a Markov chain is deterministic (all successors of any location disagree on their labelings), we can eliminate the use of $*$ in A_M and still secure Theorem 4. But this embedding without $*$ does break Theorem 4 for non-deterministic Markov chain.

Embedding of PCTL Formulas. Each PCTL formula ϕ over $\mathbb{A}\mathbb{P}$ yields a p-automaton A_ϕ without $*$ markings, $\langle 2^{\mathbb{A}\mathbb{P}}, \text{cl}_t(\phi) \cup \mathbb{A}\mathbb{P}, \rho_X, \rho_\epsilon(\phi), F \rangle$, that accepts exactly the Markov chains satisfying ϕ . The construction resembles the translation from CTL to alternating tree automata:

- $\text{cl}_t(\phi)$ denotes the set of temporal subformulas of ϕ
- F consists of $\mathbb{A}\mathbb{P}$ and their negations, and all ψ of $\text{cl}_t(\phi)$ not of form $\psi_1 \text{U} \psi_2$
- functions ρ_X and ρ_ϵ are defined in Fig. 8.

Function ρ_X unfolds fix-points and replaces the threshold context $[\cdot]_{\triangleright p}$ with $\llbracket \cdot \rrbracket_{\triangleright p}$. That replacement is also done by function ρ_ϵ for the initial condition. The effect of these functions is similar to that achieved by using ϵ transitions to translate CTL formulas into two-way tree automata.

We now have that $\psi \in \text{cl}_t(\phi)$ for subformulas $[\psi]_{\triangleright p}$ of ϕ . Also, $[\psi_1 \text{U} \psi_2]_{\triangleright p}$ may be an element in $\text{cl}_t(\phi)$ whereas $\llbracket \psi_1 \text{U} \psi_2 \rrbracket_{\triangleright p}$ can only be an element of $\llbracket \text{cl}_t(\phi) \rrbracket_{>}$, it wraps $\psi_1 \text{U} \psi_2$ in the probabilistic quantification $\llbracket \cdot \rrbracket_{\triangleright p}$ of A_ϕ .

Theorem 5 *For any $M \in \text{MC}_{\mathbb{A}\mathbb{P}}$ and PCTL formula ϕ over $\mathbb{A}\mathbb{P}$, $M \models \phi$ iff $M \in \mathcal{L}(A_\phi)$. Deciding $M \in \mathcal{L}(A_\phi)$ is polynomial in the size of M and linear in the size of ϕ .*

We prove Theorem 5 by structural induction on PCTL formulas (i.e., state formulas), showing that for all locations s in M and all PCTL subformulas φ' of PCTL formula φ we have $s \in \llbracket \varphi' \rrbracket$ iff $\text{val}(s, \rho_\epsilon(\varphi')) = 1$ for configuration $(s, \rho_\epsilon(\varphi'))$ in the acceptance game of $M \in \mathcal{L}(A_\phi)$. As for the complexity, the membership game for $M \in \mathcal{L}(A_\phi)$ collapses to solving a sequence (linear in the size of ϕ) of

weak stochastic games with solely probabilistic configurations. Such games are solvable in polynomial time.

Example 4 *For $\varphi = [a \text{U} [X b]_{>0.5}]_{\geq 0.3}$ we have $A_\varphi = \langle 2^{\{a, b\}}, \text{cl}_t(\varphi) \cup \{a, b\}, \rho_X, \rho_\epsilon(\varphi), F \rangle$, where $\text{cl}_t(\varphi) = \{a \text{U} [X b]_{>0.5}, X b\}$, $\rho_\epsilon(\varphi) = (a \wedge \llbracket a \text{U} [X b]_{>0.5} \rrbracket_{\geq 0.3}) \vee \llbracket X b \rrbracket_{>0.5}$, set F equals $\{X b, a, b\}$, $\rho_X(X b) = b$, and $\rho_X(a \text{U} [X b]_{>0.5}) = (a \wedge a \text{U} [X b]_{>0.5}) \vee \llbracket X b \rrbracket_{>0.5}$.*

Corollary 1 and Theorem 5 imply that any algorithm for solving language emptiness or containment of p-automata would prove that satisfiability of PCTL is decidable [14, 2].

In comparing automata and temporal logic, automata usually can count but temporal logics cannot. Thus, just as alternating tree automata are more expressive than CTL and CTL*, p-automata are more expressive than PCTL.

Additionally, p-automata can encode recursive, probabilistic properties that we believe to be not expressible in PCTL. For example, $A_R = \langle 2^{\{a\}}, \{q_2\}, \delta, \llbracket q_2 \rrbracket_{>0}, \{q_2\} \rangle$ with $\delta(q_2, \{a\}) = \llbracket q_2 \rrbracket_{\geq 0.5}$ and $\delta(q_2, \{\}) = \text{ff}$, asserts the recursive, probabilistic property that a location is labeled a , and that the probability of its successors with the same property is at least 0.5. A naive attempt of expressing this in PCTL would be $\eta = a \wedge [(\neg a \vee [X a]_{\geq 0.5}) W \neg a]_{\geq 1}$. Then $\mathcal{L}(A_\eta) \subset \mathcal{L}(A_R)$ but this inclusion is strict.

6 Simulation of p-Automata

We now define simulation of p-automata that under-approximates language containment: if p-automaton B simulates p-automaton A (denoted $A \leq B$), then $\mathcal{L}(A)$ is contained in $\mathcal{L}(B)$, under qualifications detailed in the formal theorem below. This simulation is defined as a combination of fair simulation [13], simulation for alternating word automata [9], probabilistic bisimulation [20], and the games defined in Section 3. The simulation takes into account the structure of the automata, their acceptance condition, and local probabilistic constraints. We show that whether B simulates A can be decided in EXPTIME and that simulation under-approximates language containment.

We define simulation through a series of games G_{\leq} on the product of states and transitions of A and B : state u of B simulates state r of A iff Player 0 wins from configuration (r, u) in the corresponding game. More general configurations (α, β) are such that α is part of a transition of A and β is part of a transition of B . The classification of α and β as unbounded, bounded with $*$, bounded with \forall , or simple classifies (α, β) as one of 9 types. Here, we restrict our attention to the case that A and B do not use the \forall operator. Furthermore, a state that is part of a bounded SCC in B cannot simulate a state that is part of an unbounded SCC in A . These restrictions lead to the consideration of four cases, and are sufficient for handling simulation of automata that result from embedding PCTL formulas or Markov chains.

For sake of simplicity, p-automata $A = \langle \Sigma, Q, \delta, \varphi_a^{\text{in}}, F \rangle$ and $B = \langle \Sigma, U, \delta, \psi_b^{\text{in}}, F \rangle$ satisfy $Q \cap U = \{\}$ and we use δ

$\rho_x(\mathbf{a}, \sigma) = \mathbf{tt}$	if $\mathbf{a} \in \sigma$	$\rho_x(\mathbf{a}, \sigma) = \mathbf{ff}$	if $\mathbf{a} \notin \sigma$	$\rho_\epsilon(\mathbf{a}) = \mathbf{a}$	$\rho_\epsilon(\neg \mathbf{a}) = \neg \mathbf{a}$
$\rho_x(\neg \mathbf{a}, \sigma) = \mathbf{tt}$	if $\mathbf{a} \notin \sigma$	$\rho_x(\neg \mathbf{a}, \sigma) = \mathbf{ff}$	if $\mathbf{a} \in \sigma$	$\rho_\epsilon(\varphi_1 \vee \varphi_2) = \rho_\epsilon(\varphi_1) \vee \rho_\epsilon(\varphi_2)$	
$\rho_\epsilon(\varphi_1 \wedge \varphi_2) = \rho_\epsilon(\varphi_1) \wedge \rho_\epsilon(\varphi_2)$		$\rho_\epsilon(\llbracket X \varphi_1 \rrbracket_{\bowtie p}) = \llbracket X \varphi_1 \rrbracket_{\bowtie p}$		$\rho_x(X \varphi_1, \sigma) = \rho_\epsilon(\varphi_1)$	
$\rho_\epsilon(\llbracket \varphi_1 \mathbf{U} \varphi_2 \rrbracket_{\bowtie p}) = (\rho_\epsilon(\varphi_1) \wedge \llbracket \varphi_1 \mathbf{U} \varphi_2 \rrbracket_{\bowtie p}) \vee \rho_\epsilon(\varphi_2)$				$\rho_x(\varphi_1 \mathbf{U} \varphi_2, \sigma) = (\rho_\epsilon(\varphi_1) \wedge \varphi_1 \mathbf{U} \varphi_2) \vee \rho_\epsilon(\varphi_2)$	
$\rho_\epsilon(\llbracket \varphi_1 \mathbf{W} \varphi_2 \rrbracket_{\bowtie p}) = (\rho_\epsilon(\varphi_1) \wedge \llbracket \varphi_1 \mathbf{W} \varphi_2 \rrbracket_{\bowtie p}) \vee \rho_\epsilon(\varphi_2)$				$\rho_x(\varphi_1 \mathbf{W} \varphi_2, \sigma) = (\rho_\epsilon(\varphi_1) \wedge \varphi_1 \mathbf{W} \varphi_2) \vee \rho_\epsilon(\varphi_2)$	

Figure 8. Transition function ρ_x and auxiliary function ρ_ϵ of A_φ

for the transition function of both automata and F for both acceptance conditions. We determine whether B simulates A by a sequence of weak and stochastic weak games. The strict versions of the partial orders on equivalence classes of A and B are well-founded and so their lexicographical ordering is a well-founded ordering \prec on the sets of configurations of the game. Namely, $((\varphi), (\psi)) \prec ((\tilde{\varphi}), (\tilde{\psi}))$ if either $((\varphi)) \prec_A ((\tilde{\varphi}))$, or $((\varphi)) = ((\tilde{\varphi}))$ and $((\psi)) \prec_B ((\tilde{\psi}))$. Consider a pair of equivalence classes $((\varphi), (\psi))$, where φ is in A and ψ is in B . As before, all pairs larger than $((\varphi), (\psi))$ with respect to \prec have already been handled: for every φ' and ψ' with $((\varphi), (\psi)) \prec ((\varphi'), (\psi'))$ value $\text{val}(\varphi', \psi') \neq \perp$ is pre-seeded.

Case 1: Let $((\varphi))$ and $((\psi))$ be SCCs where $((\varphi))$ has no transitions in E_b , and $((\psi))$ no transitions in E_u and no \checkmark markings. We set $\text{val}(\varphi, \psi) = 0$; bounded-with-* states cannot simulate unbounded states.

Case 2: Let $((\varphi))$ and $((\psi))$ be SCCs such that both $((\varphi))$ and $((\psi))$ have no transitions in E_b . Then $G_{\leq}((\varphi), (\psi))$ is a stochastic weak game with V_0, V_1 , and E defined in Fig. 9, $V = \{(\tilde{\varphi}, \tilde{\psi}) \mid \tilde{\varphi} \preceq_A \varphi \text{ and } \tilde{\psi} \preceq_B \psi\}$, and $V_p = \{\}$. As pre-seeded values $\text{val}(\tilde{\varphi}, \tilde{\psi})$ for configurations $(\tilde{\varphi}, \tilde{\psi})$ with $((\varphi), (\psi)) \prec ((\tilde{\varphi}), (\tilde{\psi}))$ may be in the open interval $(0, 1)$, we treat $G_{\leq}((\varphi), (\psi))$ as a stochastic weak game.

Intuitively, Player 1 resolves disjunctions on the left and conjunctions on the right and does this before Player 0 needs to move. Player 0 resolves conjunctions on the left and disjunctions on the right when Player 1 cannot move. From configurations of the form (q', u') , where q' is a state of A and u' is a state of B , Player 1 chooses a letter $\sigma \in \Sigma$ and applies the transitions of q' and u' reading σ .

Finally, an infinite play in $G_{\leq}((q), (u))$ is winning for Player 0 if $((\varphi)) \cap Q \subseteq F$ implies $((\psi)) \cap U \subseteq F$. By Theorem 1 every configuration c has a value for Player 0. We set $\text{val}(c)$ to that value.

Case 3: Let $((\varphi))$ and $((\psi))$ be SCCs such that both have neither transitions in E_u nor \checkmark markings. Then $G_{\leq}((\varphi), (\psi))$ is a weak game. Let

$$\begin{aligned} \tilde{\varphi} &= *(\llbracket q_1 \rrbracket_{\bowtie_1 p_1}, \dots, \llbracket q_n \rrbracket_{\bowtie_n p_n}) \\ \tilde{\psi} &= *(\llbracket u_1 \rrbracket_{\bowtie'_1 p'_1}, \dots, \llbracket u_m \rrbracket_{\bowtie'_m p'_m}) \\ \mathcal{F}_{\tilde{\varphi}, \tilde{\psi}} &= [n] \times [m] \rightarrow [0, 1] \end{aligned}$$

Also, $f \in \mathcal{F}_{\tilde{\varphi}, \tilde{\psi}}$ is *disjoint* if there is $\{a_{i,j} \in [0, 1] \mid i \in [n] \text{ and } j \in [m]\}$ with (i) $\sum_{j \in [m]} a_{i,j} = 1$ for all $i \in [n]$ and (ii) $\sum_{i \in [n]} a_{i,j} \cdot p_i \cdot f(i, j) > p'_j$ for all $j \in [m]$, or

$\sum_{i \in [n]} a_{i,j} \cdot p_i \cdot f(i, j) = p'_j$ and either \bowtie'_j is \geq or there is i' with $a_{i',j} > 0$ and $\bowtie_{i'}$ is $>$. Let $\mathcal{F}_{\tilde{\varphi}, \tilde{\psi}}^*$ be the set of disjoint functions. The configurations of $G_{\leq}((\varphi), (\psi))$ are

$$V = \{(\tilde{\varphi}, \tilde{\psi}, f) \mid \tilde{\varphi} \in ((\varphi)), \tilde{\psi} \in ((\psi)), \text{ and } f \in \mathcal{F}_{\tilde{\varphi}, \tilde{\psi}}^*\} \cup \{(\tilde{\varphi}, \tilde{\psi}), (\tilde{\varphi}, \tilde{\psi}, v) \mid \tilde{\varphi} \preceq_A \varphi, \tilde{\psi} \preceq_B \psi, \text{ and } v \in [0, 1]\}$$

and the definition of V_0, V_1 , and E are given in Fig. 10. Set V above is uncountable and infinitely branching, as branching includes a choice of a function $f: [n] \times [m] \rightarrow [0, 1]$. The techniques that were used in Section 3 can be used to make these games finite branching; and, if both A and B are finite, these games will be finite, too. For $(\gamma, \epsilon) \in \llbracket Q \rrbracket^* \times \llbracket U \rrbracket^*$ with

$$\begin{aligned} \gamma &= *(\llbracket q_1 \rrbracket_{\bowtie_1 p_1}, \dots, \llbracket q_n \rrbracket_{\bowtie_n p_n}) \\ \epsilon &= *(\llbracket u_1 \rrbracket_{\bowtie'_1 p'_1}, \dots, \llbracket u_m \rrbracket_{\bowtie'_m p'_m}) \end{aligned}$$

in order to show that ϵ simulates γ , Player 0 needs to show that the probability of ϵ (and its partition) can be supported by γ . Accordingly, from (γ, ϵ) Player 0 chooses $f: [n] \times [m] \rightarrow [0, 1]$ and moves to configuration (γ, ϵ, f) . Such a configuration relates to the claim that q_i is related to u_j with proportion $f(i, j)$ and that f can be partitioned (using the $\{a_{i,j}\}$ to support the different u_j 's). Then, Player 1 chooses i and j such that $f(i, j) > 0$ and an alphabet letter $\sigma \in \Sigma$, leading to a configuration of the form $(\delta(q_i, \sigma), \delta(u_j, \sigma), f(i, j))$. Conjunctions and disjunctions are resolved in the usual way until either reaching another configuration in $\llbracket Q \rrbracket^* \times \llbracket U \rrbracket^* \times [0, 1]$, in which case the value $f(i, j)$ is ignored (as $f(i, j) \leq 1$), or until the play reaches a configuration with a pre-seeded value v . Then, if $f(i) \leq v$ Player 0 has fulfilled her obligation and she wins. If $f(i) > v$, Player 0 failed and she loses. An infinite play in $G_{\leq}((\varphi), (\psi))$ is winning for Player 0 if $((\varphi)) \cap Q \subseteq F$ implies $((\psi)) \cap U \subseteq F$. By Theorem 1, every $c \in V$ has a value in $\{0, 1\}$ for Player 0. We set $\text{val}(c)$ to that value.

Case 4: Let $((\varphi))$ and $((\psi))$ be SCCs where $((\varphi))$ has no E_u transitions or \checkmark markings, and $((\psi))$ has no E_b transitions. Then $G_{\leq}((\varphi), (\psi))$ is a stochastic weak game defined in Fig. 11 where α, α_i and β, β_i range over formulas in transitions of A and B (resp.) while γ and u range over $\llbracket Q \rrbracket^*$ and U (resp.). For probabilities p_i that don't sum up to 1, we add a sink state (losing for Player 0) that fills that gap.

An infinite play in $G_{\leq}((\varphi), (\psi))$ is winning for Player 0 if $((\varphi)) \cap Q \subseteq F$ implies $((\psi)) \cap U \subseteq F$. By Theorem 1 every configuration c has a value for Player 0. We set $\text{val}(c)$ to that value.

V_0	$= \{c \in V \mid \exists \varphi_i, \psi_i: c = (\varphi_1 \wedge \varphi_2, \psi_1 \vee \psi_2)\} \cup \{c \in V \mid \exists q': c = (q', \psi_1 \vee \psi_2), \text{ or } \exists u': c = (\varphi_1 \wedge \varphi_2, u')\}$
V_1	$= \{c \in V \mid \exists q', u': c = (q', u')\} \cup \{c \in V \mid \exists \varphi_i, \psi: c = (\varphi_1 \vee \varphi_2, \psi), \text{ or } \exists \varphi, \psi_i: c = (\varphi, \psi_1 \wedge \psi_2)\}$
E	$= \{((q', u'), (\delta(q', \sigma), \delta(u', \sigma))) \in V \times V \mid \sigma \in \Sigma\} \cup$ $\{((\varphi_1 \vee \varphi_2, \psi), (\varphi_i, \psi_i)), ((\varphi, \psi_1 \wedge \psi_2), (\varphi, \psi_i)) \in V \times V \mid i \in \{1, 2\}\} \cup$ $\{((\varphi_1 \wedge \varphi_2, \psi_2 \vee \psi_2), (\varphi_i, \psi_j)) \in V \times V \mid i, j \in \{1, 2\}\} \cup$ $\{((\varphi_1 \wedge \varphi_2, u'), (\varphi_i, u')), ((q', \psi_1 \vee \psi_2), (q', \psi_i)) \in V \times V \mid i \in \{1, 2\}\}$

Figure 9. Game $G_{\leq}(((\varphi)), ((\psi)))$ for $((\varphi))$ and $((\psi))$ unbounded

V_0	$= \{(\alpha_1 \wedge \alpha_2, \beta_1 \vee \beta_2, v), (\alpha_1 \wedge \alpha_2, \epsilon, v), (\gamma, \beta_1 \vee \beta_2, v), (\gamma, \epsilon)\} \cup \{(\alpha, \beta, v) \mid \text{val}(\alpha, \beta) \neq \perp \text{ and } v > \text{val}(\alpha, \beta)\}$
V_1	$= \{(\gamma, \epsilon, f), (\alpha_1 \vee \alpha_2, \beta, v), (\alpha, \beta_1 \wedge \beta_2, v)\} \cup \{(\alpha, \beta, v) \mid \text{val}(\alpha, \beta) = \perp \text{ or } v \leq \text{val}(\alpha, \beta)\}$
E	$= \{((\alpha_1 \vee \alpha_2, \beta, v), (\alpha_i, \beta, v)), ((\alpha, \beta_1 \wedge \beta_2, v), (\alpha, \beta_i, v)) \mid i \in \{1, 2\}\} \cup \{((\gamma, \epsilon), (\gamma, \epsilon, f))\} \cup$ $\{((\alpha_1 \wedge \alpha_2, \epsilon, v), (\alpha_i, \epsilon, v)), ((\gamma, \beta_1 \vee \beta_2, v), (\gamma, \beta_i, v)) \mid i \in \{1, 2\}\} \cup$ $\{((\gamma, \epsilon, f), (\delta(q_i, \sigma), \delta(u_j, \sigma), f(i, j))) \mid f(i, j) > 0 \text{ and } \sigma \in \Sigma\} \cup$ $\{((\alpha_1 \wedge \alpha_2, \beta_2 \vee \beta_2, v), (\alpha_i, \beta_j, v)) \mid i, j \in \{1, 2\}\}$

Figure 10. Game $G_{\leq}(((\varphi)), ((\psi)))$ for $((\varphi))$ and $((\psi))$ bounded with $*$. Where α and β range over formulas in transitions of A and B , respectively, γ and ϵ range over formulas in $\llbracket Q \rrbracket^*$ and $\llbracket U \rrbracket^*$, respectively

$$\begin{aligned}
V &= \{(\tilde{\varphi}, \tilde{\psi}) \mid \tilde{\varphi} \preceq_A \varphi \text{ and } \tilde{\psi} \preceq_B \psi\} \cup \llbracket Q \rrbracket \times U \times \Sigma \\
V_0 &= \{(\alpha_1 \wedge \alpha_2, \beta_1 \vee \beta_2), (\alpha_1 \wedge \alpha_2, u), (\gamma, \beta_1 \vee \beta_2)\} \\
V_1 &= \{(\alpha_1 \vee \alpha_2, \beta), (\alpha, \beta_1 \wedge \beta_2), (\gamma, u)\} \\
V_p &= \llbracket Q \rrbracket^* \times U \times \Sigma \\
E &= \{((\alpha_1 \vee \alpha_2, \beta), (\alpha_i, \beta)) \mid i \in \{1, 2\}\} \cup \\
&\quad \{((\alpha, \beta_1 \wedge \beta_2), (\alpha, \beta_i)) \mid i \in \{1, 2\}\} \cup \\
&\quad \{((\alpha_1 \wedge \alpha_2, \beta_1 \vee \beta_2), (\alpha_i, \beta_j)) \mid i, j \in \{1, 2\}\} \cup \\
&\quad \{((\gamma, u), (\gamma, u, \sigma)), ((\gamma, u, \sigma), (\delta(q_i, \sigma), \delta(u, \sigma)))\} \\
\kappa((\gamma, u, \sigma))((\delta(q_i, \sigma), \delta(u, \sigma))) &= p_i
\end{aligned}$$

Figure 11. Weak stochastic game for Case 4.

Intuitively, a state u measures the probability of some regular set of paths, and a state $\llbracket q \rrbracket_{\times p}$ can restrict the immediate steps taken by a Markov chain as well as enforce some regular structure on paths. Thus, this stochastic weak game establishes the conditions under which a Markov chain accepted from $\llbracket q \rrbracket_{\times p}$ can be also accepted from u .

The case when $((\varphi))$ or $((\psi))$ is a trivial SCC is subsumed by at least one of the four preceding cases. As for the acceptance game, this ambiguity is unproblematic as the game values in $G_{\leq}(((\varphi)), ((\psi)))$ are then determined by the propagation of pre-seeded game values.

Definition 3 We say that B simulates A , denoted $A \leq B$, if the value of configuration $(\varphi_a^{\text{in}}, \psi_b^{\text{in}})$, computed in the previous sequence of games, is 1.

Theorem 6 Let A and B be p -automata over 2^{AP} that contain no occurrence of \heartsuit . If A and B are finite, then $A \leq B$ implies $\mathcal{L}(A) \subseteq \mathcal{L}(B)$ and $A \leq B$ can be decided in EXPTIME. If A is A_M for a $M \in \text{MC}_{\text{AP}}$, then $A \leq B$ iff $\mathcal{L}(A) \subseteq \mathcal{L}(B)$ for all $B \in \text{MC}_{\text{AP}}$.

We now get sound and complete verification of $M \models \phi$ through simulations, in the sense of Dams & Namjoshi [6].

Corollary 2 For every infinite Markov chain $M \in \text{MC}_{\text{AP}}$ and PCTL formula ϕ over AP we have $M \models \phi$ iff there is a finite p -automata A with $A_M \leq A$ and $A \leq A_\phi$.

To see this, any such A implies $\mathcal{L}(A_M) \subseteq \mathcal{L}(A)$ and $\mathcal{L}(A) \subseteq \mathcal{L}(A_\phi)$ by both parts of Theorem 6 – noting that neither A_M nor A_ϕ have any occurrence of \heartsuit . Thus, $M \models \phi$ holds by Theorems 4 and 5. Conversely, if there is no such A , then A_ϕ can also not be such an A . As $A_\phi \leq A_\phi$ this implies $A_M \not\leq A_\phi$ and so $\mathcal{L}(A_M) \not\subseteq \mathcal{L}(A_\phi)$. So there is some $M' \sim M$ with $M' \not\models \phi$. Since $M' \sim M$, we get $M \not\models \phi$ as well by Lemma 1.

This method for deciding $M \models \phi$ via simulations is thus complete in the sense of [6]. To our knowledge, this is the first such completeness result for PCTL and Markov chains.

We now discuss the proof of Theorem 6. The first claim of Theorem 6 is proved as follows. Assuming $M \in \mathcal{L}(A)$ and $A \leq B$ we consider configurations (s, φ) and (φ, ψ) in the corresponding games, respectively. This determines a configuration (s, ψ) in the acceptance game for $M \in \mathcal{L}(B)$. We show an invariant, that $\text{val}(s, \varphi) \cdot \text{val}(\varphi, \psi) \leq \text{val}(s, \psi)$ for all such “synchronized” configurations. In particular, we get $\text{val}(s^{\text{in}}, \varphi_a^{\text{in}}) \cdot \text{val}(\varphi_a^{\text{in}}, \psi_b^{\text{in}}) = 1 \cdot 1 \leq \text{val}(s^{\text{in}}, \psi_b^{\text{in}})$ which proves $M \in \mathcal{L}(B)$. Extending this result to infinite-state automata seems to require the treatment of infinite converging products of real numbers.

The second claim of Theorem 6 follows since the simulation game collapses to an acceptance game when the automaton A in $A \leq B$ is derived from a Markov chain.

7 Related and Future Work

Automata for coalgebras [24], for the functor whose coalgebras are Markov chains, have a corresponding logic that enjoys the finite model property. Since PCTL does not have that property, these automata cannot express PCTL – notably its path modalities. Probabilistic processes [19] use automata-theoretic techniques for refinement checking

only. Probabilistic automata [21] give only rise to probabilistic languages of non-probabilistic models. And probabilistic verification of specifications written in linear-time temporal logic (LTL) [23] uses automata-theoretic machinery but cannot reason about combinations of LTL operators and probability thresholds as found in PCTL. The stochastic games of [18] abstract Markov decision processes as a 2-person game where two sources of non-determinism, stemming from the MDP and the state space partition respectively, are controlled by different players. This separation allows for more precision of abstractions but is not complete in the sense of [6], as shown in [16]. In [8], a Hintikka game was defined for satisfaction, $M \models \phi$, between Markov chains and PCTL formulas. That game resembles our acceptance game for $M \in \mathcal{L}(A_\phi)$.

We are in the process of developing a more general notion of game such that acceptance of input for *non-uniform* p-automata can be decided by solving a *single* such game. These games generalize stochastic games and will be the subject of a future paper. Apart from the aforementioned research questions, we plan to do the following: (i) Prove matching lower bounds for acceptance by p-automata. (ii) Develop p-automata that embed Markov decision processes and stochastic games. (iii) Extend the framework to Markov chains with infinite branching.

8 Conclusions

We presented a novel kind of automata, p-automata, that read in an entire Markov chain and either accept or reject that input. We demonstrated how this acceptance can be decided through a series of stochastic weak games and weak games, at worst case exponential in the size of the automaton and in the size of the Markov chain.

We proved p-automata to be closed under Boolean operations, that language containment and emptiness are decidable, and that the language of a p-automaton is closed under bisimulation. Bisimulation equivalence classes of any Markov chain as well as the set of models of any PCTL formula were shown to be expressible as such languages. In particular, the complexity of the acceptance game matches that of probabilistic model checking for such formulas.

These results suggest that emptiness, universality, and containment of p-automata are all tightly related to the open problem of decidability of PCTL satisfiability. We then developed a (fair) simulation between p-automata that stem from Markov chains or PCTL formulas. We proved simulation to be decidable in EXPTIME and to under-approximate language containment. In particular, p-automata are a complete abstraction framework for PCTL: if an infinite Markov chain satisfies a PCTL formula, there is a finite p-automaton that abstracts this Markov chain and whose language is contained in that of the p-automaton for that PCTL formula.

Acknowledgments. This research was supported by grant UK EPSRC EP/E028985/1.

References

- [1] A. Aziz, V. Singhal, F. Balaxin, R.K. Brayton, and A.L. Sangiovanni-Vincentelli. It usually works: The temporal logic of stochastic systems. In *Proc. of CAV*, LNCS 939, 1995.
- [2] T. Brázdil, V. Forejt, J. Kretínský, and A. Kucera. The Satisfiability Problem for Probabilistic CTL. In *Proc. of LICS*, 2008.
- [3] F. Ciesinski and C. Baier. LiQuor: A tool for Qualitative and Quantitative linear time analysis of reactive systems. In *Proc. of QEST*, 2006.
- [4] A. Condon. The complexity of stochastic games. *Inform. and Comp.*, 96(2), 1992.
- [5] C. Courcoubetis and M. Yannakakis. The Complexity of Probabilistic Verification. *J. ACM*, 42(4), 1995.
- [6] D. Dams and K. Namjoshi. The existence of finite abstractions for branching time model checking. In *Proc. of LICS*, 2004.
- [7] D. Dams and K. Namjoshi. Automata as Abstractions. In *Proc. of VMCAI*, LNCS 3385, 2005.
- [8] H. Fecher, M. Huth, N. Piterman, and D. Wagner. Hintikka games for PCTL on labeled Markov chains. *Performance Evaluation*, 2009. To appear.
- [9] C. Fritz and T. Wilke. State space reductions for alternating Büchi automata: Quotienting by simulation equivalences. In *Proc. of FSTTCS*, 2002.
- [10] E. Grädel, W. Thomas, and Th. Wilke (editors). *Automata, Logics, and Infinite Games*. LNCS 2500, Springer, 2002.
- [11] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Comp.*, 6, 1994.
- [12] A. Hinton, M. Kwiatkowska, G. Norman, and D. Parker. PRISM: A tool for automatic verification of probabilistic systems. In *Proc. of TACAS*, LNCS 3920, 2006.
- [13] T. Henzinger, O. Kupferman, and S. Rajamani. Fair simulation. *Inform. and Comp.*, 173(1), 2002.
- [14] S. Hart and M. Sharir. Probabilistic propositional temporal logics. *Inform. and Cont.*, 70(2–3), 1986.
- [15] S. Hart, M. Sharir, and A. Pnueli. Termination of Probabilistic Concurrent Programs. In *Proc. of POPL*, 1982.
- [16] M. Kattenbelt and M. Huth. Verification and Refutation of Probabilistic Specifications via Games. In *Proc. of FSTTCS*, LIPIcs 4, 2009.
- [17] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains*. Springer Verlag, 1976.
- [18] M. Kwiatkowska, G. Norman, and D. Parker. Game-based abstraction for Markov decision processes. In *Proc. of QEST*, 2006.
- [19] K. G. Larsen and B. Jonsson. Specification and Refinement of Probabilistic Processes. In *Proc. of LICS*, 1991.
- [20] K. G. Larsen and A. Skou. Bisimulation through probabilistic testing. *Inform. and Comp.*, 94, 1991.
- [21] M. O. Rabin. Probabilistic automata. *Inform. and Cont.*, 6, 1963.
- [22] P. Ramadge and W. Wonham. The control of discrete event systems. *Trans. on Control Theory*, 77, 1989.
- [23] Moshe Y. Vardi. Automatic Verification of Probabilistic Concurrent Finite-State Programs. In *Proc. of FOCS*, 1985.
- [24] Y. Venema. Automata and Fixed Point Logic: a Coalgebraic Perspective. *Inform. and Comp.*, 204(4), 2006.