PCTL Model Checking of Markov Chains: Truth and Falsity as Winning Strategies in Games

Harald Fecher[a], Michael Huth[*,b], Nir Piterman[b], Daniel Wagner[b]

[a]Institut für Informatik, Universität Freiburg, Georges-Köhler-Allee 79, 79110 Freiburg, Germany
[b]Department of Computing, Imperial College London, South Kensington campus, London, SW7 2AZ, United Kingdom

Abstract

Probabilistic model checking is a technique for verifying whether a model such as a Markov chain satisfies a probabilistic, behavioral property – e.g. “with probability at least 0.999, a device will be elected leader.” Such properties are expressible in probabilistic temporal logics, e.g. PCTL, and efficient algorithms exist for checking whether these formulae are true or false on finite-state models.

Alas, these algorithms don’t supply diagnostic information for why a probabilistic property does or does not hold in a given model. We provide here complete and rigorous foundations for such diagnostics in the setting of countable labeled Markov chains and PCTL. For each model and PCTL formula, we define a game between a Verifier and a Refuter that is won by Verifier if the formula holds in the model, and won by Refuter if it doesn’t hold. Games are won by exactly one player, through monotone strategies that encode the diagnostic information for truth and falsity (respectively). These games are infinite with Büchi type acceptance conditions where simpler fairness conditions are shown not be to sufficient. Verifier can always force finite plays for certain PCTL formulae, suggesting the existence of finite-state abstractions of models that satisfy such formulae.

Key words: Markov chains, probabilistic computation tree logic, game theory, fairness conditions, probabilistic evidence

[*]Corresponding author
Email addresses: `fecher@informatik.uni-freiburg.de` (Harald Fecher), `M.Huth@doc.imperial.ac.uk` (Michael Huth), `Nir.Piterman@doc.imperial.ac.uk` (Nir Piterman), `dwagner@doc.imperial.ac.uk` (Daniel Wagner)

2000 MSC: 60J10 (Markov chains with discrete parameter), 03B44 (temporal logic), 91A05 (2-person games), 91A15 (stochastic games)

1. Introduction

Countable labeled Markov chains [12, 3] are an important class of stochastic processes for the modeling of probabilistic systems. Probabilistic Computation Tree Logic (PCTL) [7] is a probabilistic temporal logic whose formulae ϕ can express practically relevant specifications, e.g. "with probability at least 0.999, a device will be elected leader" – a requirement within a telecommunications standard such as [1] – can be written as $[\mathsf{tt}\,\mathsf{U}\,\mathsf{someLeaderElected}]_{\geq 0.999}$ in PCTL. A denotational semantics $\llbracket\phi\rrbracket_M$ over labeled Markov chains M then renders truth or falsity of ϕ, where $\llbracket\phi\rrbracket_M$ is the set of states in M at which ϕ is true.

Efficient algorithms exist that compute, over a finite-state labeled Markov chain, the set of states that satisfy a given PCTL formula (e.g. the ones implemented in the probabilistic model checker PRISM [14]). A specifier, however, may need more information than just knowledge of that set. The specifier may want to understand why a particular state or set of states of interest is in that set, and any such information may be seen as evidence or diagnostics of truth. Equally, the specifier may be interested in comprehending why a particular state is not in that set, and any such information would now be evidence for falsity.

We here mean to provide such a mathematical formalism: a precise operational account of truth and falsity of PCTL formulae, expressed in a manner that is explorable step by step by humans and machines alike. The formalism we suggest is that of Hintikka games [9], played between two players Verifier and Refuter, and their notion of strategy for these players. The appeal of these games is that truth amounts to the existence of a winning strategy for Verifier, whereas falsity is captured by the existence of a winning strategy for Refuter. Such Hintikka games for PCTL over labeled Markov chains are meant to establish firm foundations on which questions about the existence and computation of finitary evidence of truth and falsity of PCTL formulae can be phrased, studied, and evaluated.

We now sketch the idea behind Hintikka games for first-order logic. A Tarskian notion of truth, $\models$, is a formally defined predicate between models and formulae of first-order logic and "property ϕ is true in model M" is defined as "predicate $M \models \phi$ holds". For example, if M is the set of natural numbers $\{0, 1, \ldots\}$ and ϕ is $\neg\exists x\,((x * x < x + 1) \wedge (x > 1))$, then $M \models \phi$ holds since all natural numbers are either at most 1 or their square is bigger than their successor number.

For model M and formula ϕ, a Hintikka game $\mathtt{G}(M,\phi)$ involves two players, Verifier $\mathtt{V}$ (who wants to prove that M satisfies ϕ) and Refuter $\mathtt{R}$ (who wants to prove that M does not satisfy ϕ). Game $\mathtt{G}(M,\phi)$ has as configurations triples of form $\langle M[\vec{x} \mapsto \vec{a}], \psi, \mathtt{C}\rangle$ where $[\vec{x} \mapsto \vec{a}]$ binds a set of variables x_i to natural numbers a_i, $\mathtt{C}$ is either Refuter $\mathtt{R}$ or Verifier $\mathtt{V}$, ψ is either ϕ or a strict sub-formula of ϕ, and $\langle M, \phi, \mathtt{V}\rangle$ is the initial configuration – saying that $\mathtt{V}$ claims that ϕ is true in M, and generates a game tree whose paths are plays – finite sequences of configurations. Below, we write $!\mathtt{C}$ for the player other than $\mathtt{C}$, i.e. $!\mathtt{V} = \mathtt{R}$ and $!\mathtt{R} = \mathtt{V}$.

For sake of illustration, consider the game for our example. The formula ϕ is a negation, so the initial configuration has $\langle M, \psi_0, \mathtt{R}\rangle$ as sole next configuration for ψ_0 being $\exists x\, ((x * x < x + 1) \wedge (x > 1))$. (Thus, we just swap the player – if $\mathtt{C}$ claims $\neg\psi$ then $!\mathtt{C}$ claims ψ – and remove the negation symbol when processing a negation.) At configuration $\langle M, \psi_0, \mathtt{R}\rangle$, the formula is an existential one and so the player that claims its truth (here $\mathtt{R}$) can choose a natural number, say 5, and bind it to x, resulting in the next configuration $\langle M[x \mapsto 5], \psi_1, \mathtt{R}\rangle$, where ψ_1 is $(x * x < x + 1) \wedge (x > 1)$. (In particular, configuration $\langle M, \psi_0, \mathtt{R}\rangle$ has infinitely many next configurations, one for each natural number a bound to x.) That formula ψ_1 is a conjunction claimed to be true by $\mathtt{R}$, and so his opponent $\mathtt{V}$ can choose a conjunct.

If $\mathtt{V}$ chooses conjunct $x > 1$, the next configuration is $\langle M[x \mapsto 5], x > 1, \mathtt{R}\rangle$. Formula $x > 1$ is atomic and x is bound to 5 so we simply evaluate this to $5 > 1$, which is true. Refuter has won this play. But if V chooses $x * x < x + 1$, the next configuration is $\langle M[x \mapsto 5], x * x < x + 1, \mathtt{R}\rangle$ and now Verifier wins since $5 * 5 = 25 \not< 6 = 5 + 1$. Finally, if player $\mathtt{V}$ always chooses $x * x < x + 1$ whenever a is greater than 1, and chooses $x > 1$ whenever a is at most 1, he wins all plays in the game tree of $\mathtt{G}(M,\phi)$, affirming that ϕ is true in M.

To summarize, existential quantifiers in a configuration $\langle M, \exists x\, \psi, \mathtt{C}\rangle$ require binding its quantified variable x to an element a of the model, chosen by player $!\mathtt{C}$, with next configuration $\langle M[x \mapsto a], \psi, \mathtt{C}\rangle$. Negation in a configuration $\langle M, \neg\psi, \mathtt{C}\rangle$ determines a swap of players and the removal of the negation with next configuration $\langle M, \psi, !\mathtt{C}\rangle$. Conjunction in a configuration $\langle M, \psi_1 \wedge \psi_2, \mathtt{C}\rangle$ means that player $!\mathtt{C}$ chooses a conjunct ψ_i for the next configuration $\langle M, \psi_i, \mathtt{C}\rangle$. Atomic configurations $\langle M, \delta, \mathtt{C}\rangle$ are simply evaluated, using the binding information of the model: player $\mathtt{C}$ wins if δ is true in M, otherwise player $!\mathtt{C}$ wins.

Strategies for both players are objects that allow them to make necessary choices for determining continuation plays. For example, Verifier needs to make choices for existential quantifiers in configurations of form $\langle M, \exists x\, \psi, \mathtt{V}\rangle$, and for conjunctions in configurations of form $\langle M, \psi_1 \wedge \psi_2, \mathtt{R}\rangle$. A strategy σ is winning for a player if all plays played according to the choices offered by strategy σ are won by that player. Since all plays for first-order logic are finite, classical game theory guarantees that games $\mathsf{G}(M, \phi)$ are determined: exactly one of the two players has a winning strategy for that game.

It is well known that in ordinary set theory ZF the assumption of the Axiom of Choice is equivalent to that

> **(Correspondence)** "Verifier wins game $\mathsf{G}(M, \phi)$ if, and only if, predicate $M \models \phi$ holds".

holds. So one gets an operational and "small-step" account of truth in first-order logic from the Axiom of Choice.

We here also rely on the Axiom of Choice in proving **(Correspondence)** in our setting of PCTL and countable labeled Markov chains. This dependency appears to vanish for finite-state models and for PCTL formulae whose threshold types and controlling player satisfy simple consistency conditions developed in this paper. The latter is of interest since *any* PCTL formula can be rewritten with the help of small perturbations of thresholds that won't diminish their practical value to specifiers but that establishes, in some cases, said consistency conditions. For example, formulae $[\mathsf{tt}\,\mathsf{U}\,\mathsf{someLeaderElected}]_{\geq 0.999}$ and $[\mathsf{tt}\,\mathsf{U}\,\mathsf{someLeaderElected}]_{>p}$ with $p = 0.999 - 1^{-15}$ have different threshold types ($\geq$ versus $>$) but the latter formula may in practice be considered a valid substitute for the former one.

Our Hintikka games for PCTL retain the above idea: Verifier and Refuter are adversarial players, and both have to make choices of either sub-formulae or of structural elements – which for PCTL turn out to be *sub-distributions* that approximate transition distributions in labeled Markov chains.

Outline of paper. In Section 2, we review the familiar denotational semantics of PCTL for countable labeled Markov chains as models, and prove a finite-state approximation lemma for (Strong) Until formulae with non-strict thresholds under that semantics. In Section 3, the game semantics for PCTL over countable labeled Markov chains is being defined and these games are shown to be determined and to capture precisely the denotational semantics of PCTL. In Section 4, we discuss what structural properties one may assume in winning strategies for our games. A discussion of the relevance of our results to finding finite representations of winning strategies is contained in Section 5. In Section 6, we discuss related work, and we conclude in Section 7.

ϕ, ψ ::=	*PCTL formulae*
q	Atom
$\neg\phi$	Negation
$\phi \wedge \psi$	Conjunction
$[\alpha]_{\bowtie p}$	Path Probability

α ::=	*Path formulae*
$\mathsf{X}\,\phi$	Next
$\phi\,\mathsf{U}^{\leq k}\psi$	Until
$\phi\,\mathsf{W}^{\leq k}\psi$	Weak Until

Figure 1: Syntax of PCTL, where $\mathsf{q} \in \mathbb{AP}$, $k \in \mathbb{N} \cup \{\infty\}$, $p \in [0, 1]$, and $\bowtie\ \in \{<, \leq, >, \geq\}$

2. Preliminaries

(Countable) Labeled Markov chains M over a set of atomic propositions $\mathbb{AP}$ are triples (S, P, L), where S is a countable set of states, $P\colon S \times S \to [0, 1]$ is a countable stochastic matrix such that the countable sum of non-negative reals $\sum_{s' \in S} P(s, s')$ converges to 1 for all $s \in S$, and $L\colon \mathbb{AP} \to \mathbb{P}(S)$ is a labeling function where $L(q)$ is the set of states at which atomic proposition q is true. We say that M is finitely branching iff for all $s \in S$ the set $\{s' \in S \mid P(s, s') > 0\}$ is finite. A path π from state s in M is an infinite sequence of states $s_0 s_1 \dots$ with $s_0 = s$ and $P(s_i, s_{i+1}) > 0$ for all $i \geq 0$. For $Y \subseteq S$, we write $P(s, Y)$ as a shorthand for the (possibly infinite but well defined) sum $\sum_{s' \in Y} P(s, s')$.

The syntax of PCTL is given in Fig. 1. Path formulae α are wrapping PCTL formulae into "LTL" operators for Next, (Strong) Until, and Weak Until familiar from linear-time temporal logic [18]. Until formulae $\phi\,\mathsf{U}^{\leq k}\psi$ are *Strong* Untils since paths that satisfy such a formula have to maintain temporary invariant ϕ until they reach a state satisfying ψ, and such a state has to be reached within finite transitions, and also within k transitions if $k \neq \infty$. Weak Until formulae $\phi\,\mathsf{W}^{\leq k}\psi$ are *Weak* Untils since reaching a state satisfying ψ is optional if ϕ is an invariant on the path $s_0 s_1 \dots s_k$, which is understood to be π when $k = \infty$. The value $k = \infty$ is being used to express unbounded Untils, whereas $k \in \mathbb{N}$ expresses a proper step bound on Untils. We write $\phi\,\mathsf{U}\,\psi$ as a shorthand for $\phi\,\mathsf{U}^{\leq\infty}\psi$, and $\phi\,\mathsf{W}\,\psi$ as shorthand for $\phi\,\mathsf{W}^{\leq\infty}\psi$.

Path formulae α are interpreted as predicates $\pi \models \alpha$ over paths π of M. PCTL formulae ϕ are interpreted as subsets $\llbracket\phi\rrbracket_M$ of S. The semantics of path and PCTL formulae is the standard one, given in Fig. 2. The measure space of path sets is generated from cylinder path sets in the standard fashion [12]. We thus write $\mathsf{Prob}_M(s, \alpha)$ for the probability of the measurable set $\mathsf{Path}(s, \alpha)$ of paths $\pi = s \dots$ with $\pi \models \alpha$. PCTL formulae wrap path formulae with probability

$$\begin{array}{lcl}
\pi & \models & \mathsf{X}\,\phi \text{ iff } \; s_1 \in \|\phi\|_M \\
\pi & \models & \phi\,\mathsf{U}^{\leq k}\psi \text{ iff } \exists\, l \in \mathbb{N}\colon l \leq k \;\&\; s_l \in \|\psi\|_M \;\&\; (\forall 0 \leq j < l\colon s_j \in \|\phi\|_M) \\
\pi & \models & \phi\,\mathsf{W}^{\leq k}\psi \text{ iff } \forall\, l \in \mathbb{N}\colon 0 \leq l \leq k \to (s_l \in \|\phi\|_M) \vee (\exists\, 0 \leq j \leq l\colon s_j \in \|\psi\|_M)
\end{array}$$

$$\|q\|_M = L(q) \qquad \|\phi \wedge \psi\|_M = \|\phi\|_M \cap \|\psi\|_M$$

$$\|\neg\phi\|_M = S \setminus \|\phi\|_M \qquad \|[\alpha]_{\bowtie p}\|_M = \{s \in S \mid \mathsf{Prob}_M(s, \alpha) \bowtie p\}$$

Figure 2: Semantics $\pi \models \alpha$ of path formulae for paths $\pi = s_0 s_1 \dots$, and semantics $\|\phi\|_M$ of PCTL formulae: $\mathsf{Prob}_M(s, \alpha)$ is probability of set $\mathsf{Path}(s, \alpha)$ of paths $\pi = s \dots$ in M with $\pi \models \alpha$

thresholds (turning predicates on paths into predicates on states), interpret atoms according to the labeling function L, and interpret negation and conjunction as complement and intersection of predicates (respectively). The operators $\phi \vee \psi$ (disjunction) and $\phi \to \psi$ (implication) are derived as $\neg(\neg\phi \wedge \neg\psi)$ and $\neg\phi \vee \psi$, respectively. Let $\mathtt{ff}$ be an abbreviation for any $[\alpha]_{>1}$, and $\mathtt{tt}$ denotes any $[\alpha]_{\geq 0}$.

Example 1 *For labeled Markov chain M in Fig. 3(a), $\|[q\,\mathsf{U}\,r]_{\geq 1/2}\|_M = \{s_0, s_1\}$. For the labeled Markov chain $M_2^{s_0}$ in Fig. 3 we have that $\|[q\,\mathsf{W}\,r]_{\geq 5/9}\|_M$ equals $\{s_0, s_0 s_1, s_0 s_1 s_1, s_0 s_0, s_0 s_0 s_1, s_0 s_0 s_0\}$.*

We say that PCTL formulae ϕ and ψ are semantically equivalent iff for all labeled Markov chains M we have $\|\phi\|_M = \|\psi\|_M$. Each PCTL formula ϕ is semantically equivalent to a PCTL formula in "GreaterThan" normal form obtained by replacing all occurrences of the form $[\alpha]_{<p}$ in ϕ with the PCTL formula $\neg[\alpha]_{\geq p}$, and by replacing any occurrences of the form $[\alpha]_{\leq p}$ in ϕ with the PCTL formula $\neg[\alpha]_{>p}$. For example, the "GreaterThan" normal form of the formula $[[\mathsf{X}\,[q\,\mathsf{U}\,r]_{<1/3}]_{\leq 1/2}\,\mathsf{U}\,r]_{>1/4}$ is $[\neg[\mathsf{X}\,\neg[q\,\mathsf{U}\,r]_{\geq 1/3}]_{>1/2}\,\mathsf{U}\,r]_{>1/4}$.

Assumption 1 (GreaterThan) *Without loss of generality, PCTL of Fig. 1 is restricted to $\bowtie\; \in \{\geq, >\}$.*

We now state and prove a finite-state approximation lemma for the validity of Until formulae with non-strict probability thresholds at states of labeled Markov

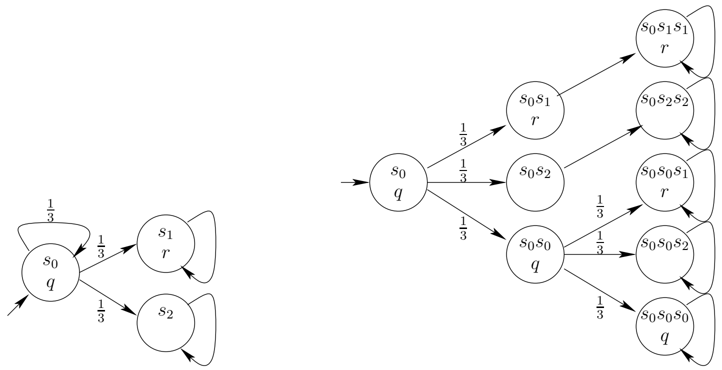

(a) Labeled Markov chain M (b) Unfolding $M_2^{s_0}$ of M

Figure 3: Labeled Markov chain M with $s_0 \in [\![[q \,\mathsf{U}\, r]_{\geq 1/2}]\!]_M$, since $\mathsf{Prob}_M(s_0, q \,\mathsf{U}\, r) = 1/2$

chains. This lemma will be crucial in proving that our game semantics of PCTL, developed in Section 3, captures exactly the denotational semantics in Fig. 2.

Definition 1 (Finite Unfoldings) *Let $M = (S, P, L)$ be a labeled Markov chain. For each $i \in \mathbb{N}$ and $s_0 \in S$ we define the labeled Markov chain $M_i^{s_0} = (S_i, P_i, L_i)$, a random tree with root s_0: unfold M from s_0 as a full tree of depth i, where edges have positive probability according to P. This may duplicate states but such duplicates will satisfy the same atomic propositions. States at level i have a self-loop with probability 1. The probability measures $P(s, \cdot)$ at levels $< i$ are those in M. For each $j \in \mathbb{N}$ we restrict $M_i^{s_0}$ to the finite-branching, and so finite-state, labeled Markov chain $M_{i,j}^{s_0} = (S_{i,j}, P_{i,j}, L_{i,j})$ with one additional state t_{sink} which satisfies $\mathtt{tt}$ but no other $q \in \mathbb{AP}$: For each $s \in S_i$, let $t_1, t_2, \ldots$ be an enumeration of $\{t_k \in S_i \mid P(s, t_k) > 0\}$ such that $P(s, t_k) \geq P(s, t_{k+1})$ for all $k \in \mathbb{N}$. Then $P_{i,j}$ is obtained from P_i by setting $P_{i,j}(s, t_k) = P_i(s, t_k)$ for $k \leq j$, $P_{i,j}(s, t_{\text{sink}}) = 1 - \sum_{k=1}^{j} P_{i,j}(s, t_k)$ and $P_{i,j}(t_{\text{sink}}, t_{\text{sink}}) = 1$; state set $S_{i,j}$ consists of those s reachable from s_0 via $P_{i,j}$, and $L_{i,j}$ is L_i restricted to set $S_{i,j}$ and extended to the new state t_{sink}.*

Example 2 *Fig. 3(b) shows unfolding $M_2^{s_0}$ for labeled Markov chain M of Fig. 3(a).*

Lemma 1 (Finite-State Approximation) *Let $q, r \in \mathbb{AP}$ and $p \in [0, 1]$ for la-*

beled Markov chain $M = (S, P, L)$*. Then* $s \in [\![[q \,\mathsf{U}\, r]_{\geq p}]\!]_M$ *iff for all* $n \in \mathbb{N}$ *there are* $k, l \in \mathbb{N}$ *with* $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_{k,l}}$*.*

PROOF OF LEMMA 1. Consider first the case that M is finitely branching. Recall that $\mathsf{Path}(s, q \,\mathsf{U}\, r)$ denotes the set of paths beginning in s that satisfy $q \,\mathsf{U}\, r$. Let $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ be $\mathsf{Path}(s, (q \,\mathsf{U}^{\leq i} r) \wedge \bigwedge_{0 \leq j < i} \neg(q \,\mathsf{U}^{\leq j} r))$, i.e., paths in which q holds until location i where r holds and r does not hold in locations smaller than i. We set $\mathsf{Path}_0(s, q \,\mathsf{U}\, r)$ to be $\mathsf{Path}(s, q \,\mathsf{U}^{\leq 0} r)$, i.e. the set $\{\pi = s_0 \cdots \mid s = s_0,\ s_0 \in L(r)\}$

For the "if" part, assume that for all $n \in \mathbb{N}$ there is $k \geq 0$ such that $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$. Then, $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_M$ follows by the monotonicity of the denotational semantics for "GreaterThan" thresholds. Thus, s is in the intersection $\bigcap_{n \in \mathbb{N}} [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_M$ which equals $[\![[q \,\mathsf{U}\, r]_{\geq p}]\!]_M$.

For the "only if" part, let $s \in [\![[q \,\mathsf{U}\, r]_{\geq p}]\!]_M$ and $n \in \mathbb{N}$. It suffices to find some $k \geq 0$ with $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$. As $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ is of form $\mathsf{Path}(s, \alpha)$ for a path formula α, that set of paths is measurable. For all $i \neq j$ note that sets $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ and $\mathsf{Path}_j(s, q \,\mathsf{U}\, r)$ are disjoint. Since $\mathsf{Path}(s, q \,\mathsf{U}\, r) = \bigcup_{i \geq 0} \mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ and as the latter is a disjoint union, we know that

$$\mathsf{Prob}_M(s, \mathsf{Path}(s, q \,\mathsf{U}\, r)) = \sum_{i \geq 0} \mathsf{Prob}_M(s, \mathsf{Path}_i(s, q \,\mathsf{U}\, r))$$

By definition of convergence for that infinite sum, for every $n \in \mathbb{N}$ there exists $k \geq 0$ such that

$$\sum_{i=0}^{k} \mathsf{Prob}_M(s, \mathsf{Path}_i(s, q \,\mathsf{U}\, r)) \geq \mathsf{Prob}_M(s, \mathsf{Path}(s, q \,\mathsf{U}\, r)) - 1/n$$

As $\sum_{i=1}^{k} \mathsf{Prob}_M(s, \mathsf{Path}_i(s, q \,\mathsf{U}\, r))$ equals $\mathsf{Prob}_{M^s_k}(s, q \,\mathsf{U}\, r)$ we obtain that s is in $[\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$ and we are done.

As M is finitely branching, there exists an upper bound l on the branching degree for all states in M^s_k. It follows that $\mathsf{Prob}_{M^s_k}(s, q \,\mathsf{U}\, r) = \mathsf{Prob}_{M^s_{k,l}}(s, q \,\mathsf{U}\, r)$.

For infinite branching M the proof is similar. We have to be more careful in noticing that every path set $\mathsf{Path}_i(s, q \,\mathsf{U}\, r)$ is still measurable and have to be careful in the way in which we sum up the probability of the set $\mathsf{Path}(s, q \,\mathsf{U}\, r)$. But this works out since all infinite sums have absolute convergence, establishing that for some k we have $s \in [\![[q \,\mathsf{U}\, r]_{>p-1/n}]\!]_{M^s_k}$. The existence of $M^s_{k,l}$ as required follows from convergence of $\mathsf{Prob}_{M^s_{k,l}}(s, q \,\mathsf{U}\, r)$ to $\mathsf{Prob}_{M^s_k}(s, q \,\mathsf{U}\, r)$. □

Example 3 *Consider the labeled Markov chain in Fig. 3(a):* $\mathsf{Prob}_M(s_0, q \,\mathsf{U}\, r) = 1/2$ *is attained by paths of increasing length, as the value of the infinite sum* $\sum_{j=1}^{\infty}(1/3)^j$*. But for all* $n \in \mathbb{N}$ *there is* $i \in \mathbb{N}$ *such that* $\sum_{j=1}^{i}(1/3)^j > 1/2 - 1/n$ *and where that finite sum is attainable in a finite unfolding of* M*. For example, for* $M_2^{s_0}$ *in Fig. 3 the probability of* $q \,\mathsf{U}\, r$ *at* s_0 *is* $\frac{4}{9}$ *so for every* $n < 18$ *we have* $s_0 \in \llbracket [q \,\mathsf{U}\, r]_{>1/2-1/n} \rrbracket_{M_2^{s_0}}$*. In* $M_4^{s_0}$ *the probability of* $q \,\mathsf{U}\, r$ *at* s_0 *is* $\frac{13}{27}$ *and so for every* $n < 54$ *we have* $s_0 \in \llbracket [p \,\mathsf{U}\, q]_{>1/2-1/n} \rrbracket_{M_4^{s_0}}$*. Lemma 1 promises a similar approximation for every (countable) labeled Markov chain.*

Lemma 1 has a dual version, required in the proof of Theorem 2 below.

Corollary 1 *For labeled Markov chain* $M = (S, P, L)$*,* $q, r \in \mathbb{AP}$*, and* $p \in [0, 1]$*:* $s \notin \llbracket [q \,\mathsf{W}\, r]_{>p} \rrbracket_M$ *iff for all* $n \in \mathbb{N}$ *there are* $k, l \in \mathbb{N}$ *with* $s \notin \llbracket [q \,\mathsf{W}\, r]_{\geq p+1/n} \rrbracket_{M_{k,l}^s}$*.*

PROOF OF COROLLARY 1. $s \notin \llbracket [q \,\mathsf{W}\, r]_{>p} \rrbracket_M$ iff $s \in \llbracket [\neg r \,\mathsf{U}\, (\neg q \wedge \neg r)]_{\geq 1-p} \rrbracket_M$, as $q \,\mathsf{W}\, r \equiv \neg(\neg r \,\mathsf{U}\, (\neg q \wedge \neg r))$. By Lemma 1, for all $n \in \mathbb{N}$ there are $k, l \in \mathbb{N}$ with $s \in \llbracket [\neg r \,\mathsf{U}\, (\neg q \wedge \neg r)]_{\geq 1-p-1/n} \rrbracket_{M_{k,l}^s}$. Thus, $s \notin \llbracket [q \,\mathsf{W}\, r]_{>p+1/n} \rrbracket_{M_{k,l}^s}$. □

3. Game semantics

Let $M = (S, P, L)$ be a labeled Markov chain over set of atomic propositions $\mathbb{AP}$. For each state $s \in S$ and PCTL formula ϕ we define a 2-person Hintikka game $\mathsf{G}_M(s, \phi)$. As already mentioned, these games are played between two players $\mathtt{V}$ (the Verifier) and $\mathtt{R}$ (the Refuter). As before, we let $!\mathtt{V} = \mathtt{R}$ and $!\mathtt{R} = \mathtt{V}$.

After having defined these games and their winning conditions, we show that each game $\mathsf{G}_M(s, \phi)$ is won by player $\mathtt{V}$ iff $s \in \llbracket \phi \rrbracket_M$; and won by player $\mathtt{R}$ iff $s \notin \llbracket \phi \rrbracket_M$. In particular, each game $\mathsf{G}_M(s, \phi)$ is *determined*, exactly one of the players $\mathtt{V}$ and $\mathtt{R}$ wins that game. The game $\mathsf{G}_M(s, \phi)$ has as set of configurations

$$\mathsf{Cf}_M(s, \phi) = \{\langle s', \psi, \mathtt{C} \rangle \mid s' \in S,\ \psi \in \mathsf{cl}(\phi),\ \mathtt{C} \in \{\mathtt{R}, \mathtt{V}\}\}$$

where we define the set of PCTL formulae $\mathsf{cl}(\phi)$, the *closure of* ϕ, in Fig. 4. This set merely delineates the universe of PCTL formulae ψ such that all configurations $\langle t, \psi, \mathtt{C} \rangle$ reachable in game $\mathsf{G}_M(s, \phi)$ satisfy $\psi \in \mathsf{cl}(\phi)$. Set $\mathsf{cl}_1(\phi)$ is part of the closure as familiar from first-order logic. Set $\mathsf{cl}_2(\phi)$ is specific to PCTL and will be discussed implicitly in game moves for clause *Path Probability* of PCTL. The intuition behind a configuration $\langle t, \psi, \mathtt{C} \rangle$ is that player $\mathtt{C}$ claims (or has the burden of proof) that ψ holds in state t.

Set $\mathsf{cl}_1(\phi)$ is the actual set of sub-PCTL-formulae of ϕ, including ϕ itself. Set $\mathsf{cl}_2(\phi)$ consists of all formulae $[\alpha]_{\bowtie p'}$ such that either

(a) α is $\psi_1 \,\mathsf{U}\, \psi_2$, $\bowtie$ is $>$, $p \in [0,1]$, and $\bowtie' \ \in \{>, \geq\}$ with $[\alpha]_{\bowtie' p} \in \mathsf{cl}_1(\phi)$,
(b) α is $\psi_1 \,\mathsf{W}\, \psi_2$, $\bowtie$ is $\geq$, $p \in [0,1]$, and $\bowtie' \ \in \{>, \geq\}$ with $[\alpha]_{\bowtie' p} \in \mathsf{cl}_1(\phi)$,
(c) α is $\psi_1 \,\mathsf{U}^{\leq k'} \psi_2$, $p \in [0,1]$, and $\infty > k > k'$ with $[\psi_1 \,\mathsf{U}^{\geq k} \psi_2]_{\bowtie p} \in \mathsf{cl}_1(\phi)$,
(d) α is $\psi_1 \,\mathsf{W}^{\leq k'} \psi_2$, $p \in [0,1]$, and $\infty > k > k'$ with $[\psi_1 \,\mathsf{U}^{\geq k} \psi_2]_{\bowtie p} \in \mathsf{cl}_1(\phi)$

Figure 4: Closure $\mathsf{cl}(\phi) = \mathsf{cl}_1(\phi) \cup \mathsf{cl}_2(\phi)$ of ϕ satsifying an invariant: For all configurations $\langle t, \psi, \mathtt{C}\rangle$ reachable in game $\mathsf{G}_M(s, \phi)$, formula ψ is in $\mathsf{cl}(\phi)$

Definition 2 1. *The moves of game $\mathsf{G}_M(s_0, \phi)$ are defined by structural induction on $\psi \in \mathsf{cl}(\phi)$, simultaneously for all $s \in S$, in Fig. 5.*

2. *A play in $\mathsf{G}_M(s_0, \phi)$ is an element of $\mathsf{Cf}_M(s, \phi)^+ \cup \mathsf{Cf}_M(s, \phi)^\omega$ beginning in $\langle s_0, \phi, \mathtt{V}\rangle$, where next configurations are determined as in Fig. 5.*

The intuition behind the moves is as follows. In move M1, any formula $[\alpha]_{>1}$ is made semantically equivalent to $\mathtt{ff}$ whereas move M2 encodes that any formula $[\alpha]_{\geq 0}$ is semantically equivalent to $\mathtt{tt}$.

Assumption 2 *By nature of the moves M1 and M2, moves to configurations of form $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ never satisfy that $\bowtie p$ equals ≥ 0 or > 1.*

In move M3, the winner of configurations $\langle s, \mathsf{q}, \mathtt{C}\rangle$ is determined according to whether atom q is true at state s. The moves M4 and M5 are basically those familiar from first-order logic for negation and conjunction (respectively).

In order to handle more complex operators we have to devise more complex moves. In games for branching-time logics such as CTL or the μ-calculus (see e.g. [21]), the universal quantification in $\forall \mathsf{X}\ \psi$ ("at all next states, ψ holds") is resolved by Refuter's choice of a successor state; and the existential quantification in $\exists \mathsf{X}\ \psi$ ("at some next state, ψ holds") is resolved by Verifier supplying one successor state, both as familiar from the case of quantifiers in first-order logic. For the next operator in PCTL, however, things are more complicated as reflected in move M6. The next operator $[\,\mathsf{X}\, \phi]_{\bowtie p}$ includes a promised probability $\bowtie p$, "at least p" or "more than p". At configuration $\langle s, [\,\mathsf{X}\, \psi]_{\bowtie p}, \mathtt{C}\rangle$ of move M6, player $\mathtt{C}$ chooses a subset Y of $\{s' \in S \mid P(s, s') > 0\}$ satisfying $P(s, Y) \bowtie p$. (If she

M1 At configurations $\langle s, [\alpha]_{>1}, \mathtt{C}\rangle$, player !C wins

M2 At configurations $\langle s, [\alpha]_{\geq 0}, \mathtt{C}\rangle$, player C wins

M3 At configurations $\langle s, \mathsf{q}, \mathtt{C}\rangle$: player C wins if $s \in L(\mathsf{q})$; player $!C$ wins if $s \notin L(\mathsf{q})$

M4 At configuration $\langle s, \neg\psi, \mathtt{C}\rangle$, the next configuration is $\langle s, \psi, !\mathtt{C}\rangle$

M5 At configuration $\langle s, \psi_1 \wedge \psi_2, \mathtt{C}\rangle$, player !C chooses $i \in \{1, 2\}$, next configuration is $\langle s, \psi_i, \mathtt{C}\rangle$

M6 At configuration $\langle s, [\mathsf{X}\, \psi]_{\bowtie p}, \mathtt{C}\rangle$, player C chooses a subset $Y \subseteq \{s' \in S \mid P(s, s') > 0\}$ satisfying $P(s, Y) \bowtie p$; then player !C chooses some $s' \in Y$, next configuration is $\langle s', \psi, \mathtt{C}\rangle$

M7 At configuration $\langle s, [\psi_1 \,\mathsf{U}\, \psi_2]_{\geq p}, \mathtt{C}\rangle$, player !C chooses some $n \in \mathbb{N}$ such that $p - 1/n \geq 0$ with resulting next configuration $\langle s, [\psi_1 \,\mathsf{U}\, \psi_2]_{>p-1/n}, \mathtt{C}\rangle$

M8 Dually, at configuration $\langle s, [\psi_1 \,\mathsf{W}\, \psi_2]_{>p}, \mathtt{C}\rangle$, now player C chooses $n \in \mathbb{N}$ such that $p + 1/n \leq 1$ with resulting next configuration $\langle s, [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p+1/n}, \mathtt{C}\rangle$

M9 At configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ where either α is $\psi_1 \,\mathsf{U}\, \psi_2$ and $\bowtie$ is $>$; or α is $\psi_1 \,\mathsf{W}\, \psi_2$ and $\bowtie$ is $\geq$

- player C is able to move to next configuration $\langle s, \psi_2, \mathtt{C}\rangle$
- if player C didn't move, player !C can move to next configuration $\langle s, \psi_1, \mathtt{C}\rangle$
- if neither player moved above, the play must proceed as follows: Player C chooses a sub-distribution $d\colon S \to [0, 1]$ such that

$$\sum_{s' \in S} d(s') > 0 \quad \& \quad \sum_{s' \in S} d(s') \geq p \quad \& \quad \forall s' \in S\colon d(s') \leq P(s, s') \tag{1}$$

Next, player !C chooses some state $s' \in S$ with $d(s') > 0$ and the next configuration is $\langle s', [\alpha]_{\bowtie d(s') \cdot P(s,s')^{-1}}, \mathtt{C}\rangle$.

M10 At configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ where α is $\psi_1 \,\mathsf{U}^{\leq k} \psi_2$ or $\psi_1 \,\mathsf{W}^{\leq k} \psi_2$ with $k \in \mathbb{N}$:

- if $k = 0$ and α is $\psi_1 \,\mathsf{U}^{\leq k} \psi_2$, the next configuration is $\langle s, \psi_2, \mathtt{C}\rangle$
- if $k = 0$ and α is $\psi_1 \,\mathsf{W}^{\leq k} \psi_2$, player C chooses as next configuration either $\langle s, \psi_1, \mathtt{C}\rangle$ or $\langle s, \psi_2, \mathtt{C}\rangle$
- if $k > 0$, the moves are defined as in M9, except in the last item, where now k in α is decreased to $k - 1$ for that next configuration $\langle s', [\alpha]_{\bowtie d(s') \cdot P(s,s')^{-1}}, \mathtt{C}\rangle$

Figure 5: Moves of game $\mathsf{G}_M(s, \phi)$

can't find such a set, she loses the play.) The implicit claim of $\mathtt{C}$ is that ψ holds at all states of Y, and so this move acts like a universal quantification but not over all elements but over enough to make the probabilities add up. The other player $!\mathtt{C}$ then chooses an element s' from Y and the next configuration is $\langle s', \psi, \mathtt{C}\rangle$. (By Assumption 2, player $\mathtt{C}$ does not have the option of choosing $Y = \{\}$.)

This leaves us with explaining and motivating the moves for α being the Weak or Strong Until. In *qualitative* games, Until operators are resolved by using the logical equivalence $q \,\mathsf{U}\, r \;\equiv\; r \vee (q \wedge \,\mathsf{X}\,(q \,\mathsf{U}\, r))$ – and similarly for Weak Until operators. The only problem in adopting this for PCTL is in the possibility of deferring promises forever. For games in qualitative settings this is typically handled by fairness, but for PCTL fairness is not strong enough:

Example 4 *PCTL formula $[q \,\mathsf{U}\, r]_{\geq 1/2}$ holds at state s_0 in the labeled Markov chain of Fig. 3(a). But we have to appeal to the entire infinite sum $\sum_{i=1}^{\infty}(1/3)^i$ for proving this. Any fairness constraint forcing a transition from s_0 into $\{s_1, s_2\}$ cuts that infinite sum down to a finite one, failing to prove that formula for state s_0.*

However, allowing to defer the satisfaction of the Strong Until indefinitely is unsound. The PCTL formula $[q \,\mathsf{U}\, r]_{>.5}$ does not hold at s_0 but allowing Verifier to delay promises forever may be unsound, e.g., Verifier could supply the promise $1/3$ immediately, promising more than $1/6$ in the future, and then – by deferring the promise indefinitely – Verifier could win game $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{>0.5})$.

To address this problem we add a special ϵ-move as well as acceptance conditions for infinite plays. If the probability is at least p, player $\mathtt{C}$ (who claims this) should be able to prove that it is greater than $p - \epsilon$ for every $\epsilon > 0$. On the other hand, if the probability is strictly less than p then there exists an ϵ for which it is less than $p - \epsilon$; and player $!\mathtt{C}$ does not lose ground by giving up an ϵ. Thus, player $!\mathtt{C}$ chooses the ϵ and player $\mathtt{C}$ proves in finite time (appealing to Lemma 1) that she can get as close as needed to the bound. The same intuition (but dual) works for *Weak* Until, when the Weak Until formula in question does *not* hold. This follows from the semantic equivalence $\neg(\phi \,\mathsf{U}\, \psi) \equiv (\neg\psi)\,\mathsf{W}\,(\neg\phi \wedge \neg\psi)$ of path formulae.

In move M7, player $!\mathtt{C}$ makes such an $\epsilon = 1/n$ move and the next configuration is the original one except that the threshold changes from $\geq\, p$ to $>\, p - 1/n$. Player $!\mathtt{C}$ can indeed choose such an n since p cannot be 0. The intuition is that $[p, 1] = \bigcap_{n\in\mathbb{N}}(p - 1/n, 1]$ so this behaves like a *universal* quantification.

In move M8, player C can choose such an n since $p < 1$. The intuition is that a Weak Until with a $>$ threshold is the dual of a Strong Until with a $\geq$ threshold (based on $\neg(\phi \,\mathsf{U}\, \psi) \;\equiv\; (\neg\psi)\,\mathsf{W}\,(\neg\phi \wedge \neg\psi)$), so it is like an *existential* quantification. The next configuration is the original one except the threshold changes from $> p$ to $\geq p + 1/n$.

Move M9 is the most complex one. At configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$, player $\mathtt{C}$ can claim that ψ_2 is true. If she does not do this, player $!\mathtt{C}$ can claim that ψ_1 is not true. If none of these happen, player $\mathtt{C}$ has to chose a structural element of the model, a sub-distribution d of $P(s, \cdot)$ that has positive mass, approximates the probability distribution $P(s, \cdot)$, and specifies the re-distribution of promise $\bowtie p$ into promised probabilities at successor states. Since $d(s') > 0$, we also have $0 < d(s') \cdot P(s, s')^{-1} \leq 1$ in the next configuration $\langle s', [\alpha]_{\bowtie d(s') \cdot P(s,s')^{-1}}, \mathtt{C}\rangle$ of move M9 by (1). The promised probability at a next configuration with state s' is given by $d(s')$, divided by the actual transition probability $P(s, s')$.

Move M10 behaves like move M9 excect that a Bounded Until with bound 0 has to realize ψ_2 right away; a Bounded Weak Until with bound zero has to realize at least one of ψ_1 or ψ_2 right away; and the k is decreased to $k - 1$ in α if the next configuration does not have a proper sub-formula of α to consider.

In most moves, plays either end or move to configurations with *proper* sub-formula in the closure. In a configuration with Strong Until with non-strict bound or Weak Until with strict bound, the next configuration changes from non-strict to strict bound or vice versa. In a configuration with Strong Until with strict bound or Weak Until with non-strict bound, the next configuration has the same path formula and threshold type, or has a proper sub-formula.

Thus, all infinite plays end with an infinite suffix of configurations that are

A1. all of the form $\langle s_i, [\psi_1 \mathsf{W}\, \psi_2]_{\geq p_i}, \mathtt{C}\rangle$ or **A2.** all of the form $\langle s_i, [\psi_1 \mathsf{U}\, \psi_2]_{> p_i}, \mathtt{C}\rangle$

Configurations of these suffixes are either labeled by Strong Until with strict bound or Weak Until with non-strict bound, where the states and the exact probability bound may still change, but where neither the player $\mathtt{C}$ nor the sub-formulae ψ_1 and ψ_2 change.

Definition 3 (Acceptance conditions) *1. Player* $\mathtt{V}$ *wins all infinite plays with an infinite suffix either of type A1 above with* $\mathtt{C} = \mathtt{V}$*, or of type A2 above with* $\mathtt{C} = \mathtt{R}$*. Player* $\mathtt{R}$ *wins all other infinite plays: those with an infinite suffix either of type A1 when* $\mathtt{C} = \mathtt{R}$*, or of type A2 when* $\mathtt{C} = \mathtt{V}$*.*

2. Finite plays are won as stipulated in Fig. 5. In particular, if a player has to make a choice and cannot do so, the other player wins that play.

These are Büchi type acceptance conditions, and so our games are known to be determined [16]. We use the notion of strategy for player $\mathtt{C}$ informally. But such strategies contain, for each configuration of a game, at most one set of choices as required by the applicable move from M1-M10.

Example 5 *We describe a winning strategy for player $\mathtt{V}$ in game $\mathsf{G}_M(s_0, [\alpha]_{\geq 1/2})$ for M as in Fig. 3(a) and $\alpha = q \cup r$. The initial configuration is $\langle s_0, [\alpha]_{\geq 1/2}, \mathtt{V}\rangle$. In the first move, player $\mathtt{R}$ chooses $n \in \mathbb{N}$ with next configuration $\langle s_0, [\alpha]_{>1/2-1/n}, \mathtt{V}\rangle$. Then, as long as the play $\Gamma_0\Gamma_1 \ldots$ remains in configurations Γ_i of the form $\langle s_0, [\alpha]_{>p_i}, \mathtt{V}\rangle$, player $\mathtt{V}$ is going to choose the sub-distribution d with constant values $d(s_2) = 0$ and $d(s_1) = \frac{1}{3} - \frac{1}{2n}$, and dynamic value $d(s_0) = p_i - d(s_1)$. A simple calculation shows that as long as player $\mathtt{R}$ chooses s_0 as the next state (clearly, if she chooses s_1 she is going to lose as $s_1 \in L(r)$) the promised probability $> p_i$ is going to decrease according to the following sequence: $p_0 = \frac{1}{2} - \frac{1}{n}$, $p_1 = \frac{1}{2} - \frac{3}{2n}$, $p_2 = \frac{1}{2} - \frac{6}{2n}$, $p_3 = \frac{1}{2} - \frac{15}{2n}$, and in general $p_i = \frac{1}{2} - \frac{3^i+3}{4n}$ for $i \in \mathbb{N}$. Whenever p_i decreases below $\frac{1}{3}$ (and there is some $i \in \mathbb{N}$ for which this happens), player $\mathtt{V}$ still chooses d with $d(s_2) = 0$ as above but now defines $d(s_1) = p_i$ and $d(s_0) = 0$, thereby forcing player $\mathtt{R}$ to move to s_1 and lose.*

Example 6 *Although the choice of d in Example 5 may seem arbitrary, it meshes well with the use of Lemma 1. Consider again the game from Example 5. Suppose player $\mathtt{R}$ chooses $9 \in \mathbb{N}$ in the first move, with next configuration $\langle s_0, [\alpha]_{>7/18}, \mathtt{V}\rangle$. Since for the $M_2^{s_0}$ in Fig. 3, $\mathsf{Prob}_{M_2^{s_0}}(s_0, \alpha) = \frac{4}{9} > \frac{7}{18}$, player $\mathtt{V}$ can use $M_2^{s_0}$ to guide her choices. In $M_2^{s_0}$, $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_1, \alpha) = 1$ and $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_0, \alpha) = \frac{1}{3}$. Player $\mathtt{V}$ uses the gap of $\frac{1}{18}$ and re-distributes it between the successors of s_0. She can choose, for example, $d(s_1) = \frac{1}{3} - \frac{1}{54}$ and $d(s_0) = \frac{1}{9} - \frac{1}{54}$. The next possible configurations are then $\langle s_1, [\alpha]_{>17/18}, \mathtt{V}\rangle$ and $\langle s_0, [\alpha]_{>5/18}, \mathtt{V}\rangle$. Player $\mathtt{V}$ identifies the resulting states with those obtained in $M_2^{s_0}$, here $s_0 s_1$ and $s_0 s_0$ (respectively). As $s_0 s_1 \in \llbracket r \rrbracket_{M_2^{s_0}}$ the first is clearly a winning configuration. From $\langle s_0, [\alpha]_{>5/18}, \mathtt{V}\rangle$ and the corresponding location $s_0 s_0$ in $M_2^{s_0}$, player $\mathtt{V}$ notices that $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_0 s_1, \alpha) = 1$ and chooses $d(s_1) = 5/18$. The next configuration is $\langle s_1, [\alpha]_{>15/18}, \mathtt{V}\rangle$ (with corresponding $s_0 s_0 s_1$ in $M_2^{s_0}$) and won by supplying r.*

We define winning strategies and use them to define which player wins a game.

Definition 4 1. *A strategy σ for player $\mathtt{C}$ in game $\mathsf{G}_M(s, \phi)$* is winning *from a* configuration *Γ in that game iff player $\mathtt{C}$ wins all plays beginning in configuration Γ when player $\mathtt{C}$ plays according to his strategy σ.*

2. *Player $\mathtt{C}$* wins *$\mathsf{G}_M(s, \phi)$ iff player $\mathtt{C}$ has a winning strategy from $\langle s, \phi, \mathtt{V}\rangle$.*

We can now formalize our main result that the denotational semantics of PCTL is captured exactly by the existence of winning strategies in games $\mathsf{G}_M(s,\phi)$.

Theorem 2 *Let $M = (S, P, L)$ be a labeled Markov chain over $\mathbb{AP}$, $s \in S$, and ϕ a PCTL formula. Then $s \in \|\phi\|_M$ iff player $\mathtt{V}$ wins game $\mathsf{G}_M(s,\phi)$; and $s \notin \|\phi\|_M$ iff player $\mathtt{R}$ wins game $\mathsf{G}_M(s,\phi)$. In particular, game $\mathsf{G}_M(s,\phi)$ is determined.*

PROOF OF THEOREM 2. Given PCTL formula ϕ, both "iff" claims are shown by structural induction on PCTL formulae ψ in the closure of ϕ, simultaneously on all states of M. As exactly one of $s \in \|\psi\|_M$ and $s \notin \|\psi\|_M$ holds, it suffices to show both "iff" claims in Theorem 2 for such ψ in their "only if" versions, which consists of six cases. We prove only the most interesting case here, when ϕ equals $[\alpha]_{\bowtie p}$ where either

(a) α is $\psi_1 \mathsf{U}\, \psi_2$ and $\bowtie$ is $>$

(b) α is $\psi_1 \mathsf{W}\, \psi_2$ and $\bowtie$ is $\geq$ or

(c) α is $\psi_1 \mathsf{U}^{\leq k} \psi_2$ or $\psi_1 \mathsf{W}^{\leq k} \psi_2$ with $k \in \mathbb{N}$ and $\bowtie$ is either $>$ or $\geq$:

(All other cases follow a routine argument.) We show for all three cases above that **(#1)** $s \in \|\phi\|_M$ implies player $\mathtt{V}$ wins game $\mathsf{G}_M(s,\phi)$ and **(#2)** $s \notin \|\phi\|_M$ implies player $\mathtt{R}$ wins game $\mathsf{G}_M(s,\phi)$.

(#1) First, let $s \in \|\phi\|_M$. The formula α is logically equivalent to $\psi_2 \vee (\psi_1 \wedge \mathsf{X}\,\alpha)$ and, in case that α is bounded, the bound decreases by 1. It follows that it is either the case that $s \in \|\psi_2\|_M$ or $s \in \|\psi_1 \wedge [\mathsf{X}\,\alpha]_{\bowtie p}\|_M$. In the first case, player $\mathtt{V}$ chooses to move to configuration $\langle s, \psi_2, \mathtt{V}\rangle$ and by induction she has a winning strategy from this configuration. In the second case, by induction there is a winning strategy for player $\mathtt{V}$ from configuration $\langle s, \psi_1, \mathtt{V}\rangle$, so if player $\mathtt{R}$ chooses to go to this configuration, player $\mathtt{V}$ wins. If player $\mathtt{R}$ does not move to ψ_1, then M9 demands that player $\mathtt{V}$ chooses a sub-distribution $d : S \to [0,1]$ satisfying (1). By assumption $s \in \|[\mathsf{X}\,\alpha]_{\bowtie p}\|_M$. Let T be the set of states t such that $\mathsf{Prob}_M(t,\alpha) > 0$ and $P(s,t) > 0$. We choose d such that $d(s') = 0$ for all $s' \in S \setminus T$.

So it suffices to specify d on set T. For that, let $p' = \sum_{t\in T} P(s,t)\cdot \mathsf{Prob}_M(t,\alpha)$. **Consider the case that $\bowtie$ is $>$.** By assumption $p' > p$. In the case that $p = 0$, we choose some state $t \in T$ such that $\mathsf{Prob}_M(t,\alpha) > 0$, we set $d(t) = \mathsf{Prob}_M(t,\alpha)\cdot P(s,t)$, and $d(t') = 0$ for all $t' \neq t$. In the case that $p > 0$, let δ be $p' - p$.

We are going to distribute this gap δ between all the states in T according to the distribution $P(s,\cdot)$. That is, for all $t \in T$

$$d(t) = \max(0, (\mathsf{Prob}_M(t,\alpha) - \delta) \cdot P(s,t))$$

In case that $\mathsf{Prob}_M(t,\alpha) \leq \delta$ we thus have $d(t) = 0$ (and so effectively remove t from set T above). As $p' = \sum_{t \in S} \mathsf{Prob}_M(t,\alpha) \cdot P(s,t)$ and $p > 0$ there must be at least one state t such that $\mathsf{Prob}_M(t,\alpha) \geq p'$ and hence $\mathsf{Prob}_M(t,\alpha) - \delta > 0$, implying $d(t) > 0$. It follows that $\sum_{t \in T} d(t) \geq p' - \delta \geq p$.

Consider the case that $\bowtie$ is $\geq$. By assumption $p' \geq p$. Let δ be $p' - p$. For all $t \in T$, let

$$d(t) = \max(0, \mathsf{Prob}_M(t,\alpha) - \delta \cdot P(s,t))$$

If $\mathsf{Prob}_M(t,\alpha) \leq \delta$, set $d(t) = 0$. This completes the specification of sub-distribution d chosen by player $\mathtt{V}$.

Now regardless of the choice of player $\mathtt{R}$, the next configuration is $\langle t, [\alpha]_{\bowtie p'}, \mathtt{V}\rangle$ such that $t \in \|[\alpha]_{\bowtie p'}\|_M$. So player $\mathtt{V}$ maintains the truth value of the configuration. Notice that also the distance from the promised bound p' and the real probability is being maintained.

Case (a): For (Strong) Until, we appeal to Lemma 1. We treat sub-formulae ψ_1 and ψ_2 as propositions (respectively, the q and r in that lemma) and annotate states of M by ψ_1 and ψ_2. Let $p' = \mathsf{Prob}_M(s, \psi_1 \mathsf{U} \psi_2)$. By assumption $p' > p$. In particular, $s \in \|[\psi_1 \mathsf{U} \psi_2]_{\geq p'}\|_M$. Let $n \in \mathbb{N}$ be such that $p' > p' - 1/n > p$. By Lemma 1 (applied to p' instead of p), there are $k, l \geq 0$ with $s \in \|[\psi_1 \mathsf{U} \psi_2]_{> p' - 1/n}\|_{M^s_{k,l}}$ and so the probability of $\psi_1 \mathsf{U} \psi_2$ in $M^s_{k,l}$ at s is greater than p. Player $\mathtt{V}$'s strategy is to consider this system $M^s_{k,l}$. She chooses sub-distributions $d\colon S \to [0,1]$ according to the probabilities $\mathsf{Prob}_{M^s_{k,l}}(t,\alpha)$ (instead of $\mathsf{Prob}_M(t,\alpha)$ but as explained above). By definition of $M^s_{k,l}$ there can be only finite sequences of configurations of the form $\langle s', [\alpha]_{>p}, \mathtt{V}\rangle$, and so player $\mathtt{V}$ wins (cf. Example 6).

Case (b): For Weak Until $\psi_1 \mathsf{W} \psi_2$, all infinite plays have a suffix of configurations of form $\langle s', [\psi_1 \mathsf{W} \psi_2]_{\geq p}, \mathtt{V}\rangle$ and are thus winning for player $\mathtt{V}$. Finite plays again reach configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$, where induction applies directly.

Case (c): For bounded operators, as the bound decreases, in a finite number of steps the play moves to configurations of the form $\langle s', \psi_i, \mathtt{V}\rangle$ for $i \in \{1,2\}$, where induction applies directly, and in the desired manner.

(#2) Let $s \notin \|\phi\|_M$. It follows that $\mathsf{Prob}_M(s,\alpha) \leq p$ in case that $\bowtie$ is $>$; and $\mathsf{Prob}_M(s,\alpha) < p$ in case that $\bowtie$ is $\geq$. As above, α is logically equivalent to $\psi_2 \vee (\psi_1 \wedge \mathsf{X}\,\alpha)$ and in case that α is bounded the bound decreases by 1.

It follows that $s \notin \llbracket \psi_2 \rrbracket_M$ and hence there is a winning strategy for player $\mathtt{R}$ from configuration $\langle s, \psi_2, \mathtt{V} \rangle$. Also, it is either the case that $s \notin \llbracket \psi_1 \rrbracket_M$ or $s \notin \llbracket [\,\mathsf{X}\, \alpha]_{\bowtie p} \rrbracket_M$. In the first case, player $\mathtt{R}$ has a winning strategy from configuration $\langle s, \psi_1, \mathtt{V} \rangle$ and chooses this configuration. In the second case, player $\mathtt{V}$ chooses a sub-distribution $d \colon S \to [0,1]$ such that (1) hold.

We claim that there is some $s' \in S$ with $d(s') > 0$ and $\mathsf{Prob}_M(s', \alpha) \not\bowtie d(s') \cdot P(s,s')^{-1}$. Proof by contradiction: otherwise, $\mathsf{Prob}_M(s', \alpha) \bowtie d(s')$ for all s' with $d(s') > 0$ implies that $\sum_{s' | d(s') > 0} \mathsf{Prob}_M(s', \alpha) \bowtie \sum_{s' \in S} d(s') \geq p$ by (1). But this renders $\sum_{s' | d(s') > 0} \mathsf{Prob}_M(s', \alpha) \bowtie p$ which directly contradicts $s \notin \llbracket [\,\mathsf{X}\, \alpha]_{\bowtie p} \rrbracket_M$. Thus, player $\mathtt{R}$ can choose such an s' and maintain the play in configurations of the form $\langle s', [\alpha]_{\bowtie p'}, \mathtt{V} \rangle$ such that $s' \notin \llbracket [\alpha]_{\bowtie p'} \rrbracket_M$. Notice that player $\mathtt{R}$ can choose a successor s' such that

$$p' - \mathsf{Prob}_M(s', \alpha) \geq p - \mathsf{Prob}_M(s, \alpha)$$

i.e., the gap between the promise and the actual probability does not decrease.

We now study the consequences of this capability of player $\mathtt{R}$ for the different forms of path formula α:

Case (a): For Weak Until formulae, we appeal to Corollary 1. As before, we treat ψ_1 and ψ_2 as propositions and annotate states of M by them. Let $p' = \mathsf{Prob}_M(s, \psi_1 \,\mathsf{W}\, \psi_2)$. By assumption $p' \leq p$. In particular, $s \notin \llbracket [\psi_1 \,\mathsf{W}\, \psi_2]_{>p'} \rrbracket_M$. Let $n \in \mathbb{N}$ be such that $p' < p + 1/n < p$. By Corollary 1, there are $k, l \geq 0$ with $s \notin \llbracket [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p' + 1/n} \rrbracket_{M^s_{k,l}}$ and so the probability of $\psi_1 \,\mathsf{W}\, \psi_2$ in $M^s_{k,l}$ at s is less than p. Player $\mathtt{R}$'s strategy is to consider this system $M^s_{k,l}$. Let $d \colon S \to [0,1]$ be the sub-distribution chosen by player $\mathtt{V}$. As $s \notin \llbracket [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p} \rrbracket_{M^s_{k,l}}$, there is some $s' \in S$ such that $s' \notin \llbracket [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq d(s') \cdot P(s,t)^{-1}} \rrbracket_{M^s_{k,l}}$. So player $\mathtt{R}$ chooses this s'. By definition of $M^s_{k,l}$, there can be only finite sequences of configurations of form $\langle s', [\alpha]_{\geq p}, \mathtt{V} \rangle$, and so player $\mathtt{R}$ wins. This is dual to the strategy depicted for $\mathtt{V}$ in Example 6.

Case (b): For (Strong) Until formulae, infinite plays of configurations of the form $\langle s', [\psi_1 \,\mathsf{U}\, \psi_2]_{\bowtie p}, \mathtt{V} \rangle$ are winning for player $\mathtt{R}$ by the winning conditions for infinite plays. Any finite play reduces to configurations of the form $\langle s', \psi_i, \mathtt{V} \rangle$ for $i \in \{1, 2\}$, where induction applies directly, and in the desired manner.

Case (c): For bounded operators, as the bound decreases, in a finite number of steps the play moves to configurations of the form $\langle s', \psi_i, \mathtt{V} \rangle$ for $i \in \{1, 2\}$ and so player $\mathtt{R}$ wins by induction. □

In game $\mathsf{G}_M(s, \phi)$, player $\mathtt{V}$ owns initial configuration $\langle s, \phi, \mathtt{V} \rangle$. For a dual game, with the same moves but with initial configuration $\langle s, \phi, \mathtt{R}\rangle$, Theorem 2 and its proof then remain to be valid if we swap the role of players in both.

Example 7 *Consider game $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{>1/2})$, where M is as in Fig. 3(a), and let $\alpha = q \,\mathsf{U}\, r$. From configuration $\langle s_0, [\alpha]_{>1/2}, \mathtt{V}\rangle$, player $\mathtt{V}$ won't move to $\langle s_0, r, \mathtt{V}\rangle$ as she would then lose. For the same reason, player $\mathtt{R}$ won't move to $\langle s_0, q, \mathtt{V}\rangle$. So if both players play strategies that are "optimal" for them, player $\mathtt{V}$ has to choose a sub-distribution d at the initial configuration.*

If $d(s_2) > 0$, player $\mathtt{V}$ loses as player $\mathtt{R}$ can then choose s_2. So $d(s_2) = 0$ for any "optimal" strategy of player $\mathtt{V}$. But both $d(s_1)$ and $d(s_0)$ have to be positive since otherwise the mass of d can be at most $1/3$ by (1), which would violate (1). Since player $\mathtt{V}$ plays an "optimal" strategy, $d(s_1) \neq 1/3$, as otherwise player $\mathtt{R}$ could choose as next configuration $\langle s_1, [\alpha]_{>(1/3)\cdot(1/3)^{-1}}, \mathtt{V}\rangle$ and would then win by move M1. By (1), there is therefore $\epsilon > 0$ such that $d(s_1) = 1/3 - \epsilon$. In particular, player $\mathtt{R}$ won't choose s_1 as she would lose the next configuration $\langle s_1, [\alpha]_{>1-3\epsilon}, \mathtt{V}\rangle$ (since $s_1 \in L(r)$). So player $\mathtt{R}$ chooses s_0 and the next configuration is $\langle s_0, [\alpha]_{>3d(s_0)}, \mathtt{V}\rangle$. By (1), $3d(s_0)$ must be at least $1/2 + 3\epsilon$ and so player $\mathtt{V}$ promises more *in $> 3d(s_0)$ than she promised in the previous configuration.*

At configuration $\langle s_0, [\alpha]_{>3d(s_0)}, \mathtt{V}\rangle$, player $\mathtt{V}$ avoids losing only by choosing a sub-distribution d that maps s_0 to 0 and all other states to positive mass as before, and for the same reasons. Similarly, $d(s_1) < 1/3$ has to hold. So although a new function d with a new value of ϵ may be chosen, the next configuration is still of the same type $\langle s_0, [\alpha]_{>p'}, \mathtt{V}\rangle$ with $p' > 1/2$. Thus, either the play is finite and so lost for player $\mathtt{V}$ as described above; or the play is infinite and so lost for player $\mathtt{V}$ by the acceptance conditions A1 on infinite plays.

We conclude that player $\mathtt{R}$ wins that game. A winning strategy for her from the initial configuration only needs to be specified for move M9:

- *player $\mathtt{R}$ will never choose a configuration of form $\langle s_0, q, \mathtt{V}\rangle$, should such an opportunity arise*
- *whenever player $\mathtt{V}$ chooses sub-distribution d with $d(s_2) > 0$, player $\mathtt{R}$ will choose s_2*
- *otherwise, it must be the case that both $d(s_1)$ and $d(s_2)$ are positive; if $d(s_1) \geq 1/3$, player $\mathtt{R}$ chooses s_1*
- *if $d(s_1) < 1/3$, player $\mathtt{R}$ chooses s_0*

4. Winning strategies

We show that when a player can win game $\mathsf{G}_M(s, \phi)$ she can use winning strategies that are of a very specific type. In addition to being memoryless in the classical sense, these winning strategies choose very structured distributions when re-visiting a state in a configuration with a Strong or Weak Until operator.

As before, we use the notion of strategy informally. A strategy is *memoryless* if the choices of its player depend solely on the current configuration, not on the finite history of configurations that preceded the current one in a play. In our games, there can be configurations of type $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ for the same state s and the same path formula α (e.g., $\psi_1 \,\mathsf{U}\, \psi_2$) but with different bounds $\bowtie\; p$. We show that it is enough to consider winning strategies which induce bounds that change monotonically, as defined below. Subsequently, for sub-distributions $d, d' \colon S \to [0, 1]$, we write

- $d' \leq d$ iff for all $s \in S$ we have $d'(s) \leq d(s)$
- $d' < d$ iff $d' \leq d$ and $d'(s) < d(s)$ for some $s \in S$

For a *locally monotone* strategy the choice of sub-distribution d at configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ is monotone in $\bowtie p$, regardless of the history of a play.

Definition 5 (Locally Monotone Strategies) *Strategy σ for player $\mathtt{C}$ in $\mathsf{G}_M(s, \phi)$ is* locally monotone *iff for all configurations $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{\bowtie p'}, \mathtt{C}\rangle$ that occur in plays consistent with σ (but not necessarily in the same play), where d and d' are the sub-distributions chosen according to σ at these two configurations (respectively), then $p \geq p'$ implies $d \geq d'$ and $p > p'$ implies $d > d'$.*

A *cyclically monotone* strategy is monotone on cyclic paths within single plays: its player can force a decrease or increase of the thresholds depending on the path formula and on whether it is a $\mathtt{V}$ or $\mathtt{R}$ configuration.

Definition 6 (Cyclically Monotone Strategies) *A strategy σ for player $\mathtt{C}$ in game $\mathsf{G}_M(s, \phi)$ is* cyclically monotone *iff for any two configurations $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}'\rangle$ and $\langle s, [\alpha]_{\bowtie p'}, \mathtt{C}'\rangle$ that occur in this order on some play consistent with σ, then*

- *$\alpha = \psi_1 \,\mathsf{U}\, \psi_2$ and $\mathtt{C} = \mathtt{C}'$ imply $p' < p$,*
- *$\alpha = \psi_1 \,\mathsf{W}\, \psi_2$ and $\mathtt{C} = \mathtt{C}'$ imply $p' \leq p$,*
- *$\alpha = \psi_1 \,\mathsf{U}\, \psi_2$ and $!\mathtt{C} = \mathtt{C}'$ imply $p' \geq p$,*

- $\alpha = \psi_1 \,\mathsf{W}\, \psi_2$ *and* $!\mathtt{C} = \mathtt{C}'$ *imply* $p' > p$.

The existence of winning strategies implies the existence of winning strategies that are locally monotone and cyclically monotone.

Theorem 3 *For every game $\mathsf{G}_M(s, \phi)$, there exists a winning strategy for player $\mathtt{C}$ iff there exists a memoryless winning strategy for player $\mathtt{C}$ that is also locally monotone and cyclically monotone.*

PROOF OF THEOREM 3. Assuming that there exists some winning strategy for player $\mathtt{C}$ in game $\mathsf{G}_M(s, \phi)$, it suffices to show that a slight modification of the winning strategy synthesized in the proof of Theorem 2 is memoryless, locally monotone, and cyclically monotone. That slightly modified strategy will clearly be memoryless by construction. We now describe this modified winning strategy and first prove its local monotonicity, by induction as in the proof of Theorem 2. Then we prove that it is cyclically monotone.

Modified winning strategy and its local monotonicity. The only configurations where player $\mathtt{C}$ needs to make choices are of form $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}' \rangle$, $\langle s, \psi_1 \vee \psi_2, \mathtt{C} \rangle$, and $\langle s, \psi_1 \wedge \psi_2, !\mathtt{C} \rangle$.

With the latter two, we restrict $\mathtt{C}$'s strategy to choose ψ_1 whenever possible and, only when impossible, to choose ψ_2. This is similar to what one can do in Hintikka games for first-order logic. We show that the way configurations of the form $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}' \rangle$ are handled induces a memoryless and monotone strategy.

If $\alpha = \mathsf{X}\,\psi$, then the strategy defined in the proof of Theorem 2 chooses the set of successors according to the state s, and is clearly memoryless.

If $!\mathtt{C} = \mathtt{C}'$ and either $\alpha = \psi_1 \,\mathsf{U}\, \psi_2$ and $\bowtie\ =\ \geq$; or $\alpha = \psi_1 \,\mathsf{W}\, \psi_2$ and $\bowtie\ =\ >$, then player $\mathtt{C}$ has to choose a value $n \in \mathbb{N}$. By choosing the minimal possible n she ensures that the strategy is memoryless.

Consider two configurations $\langle s, [\alpha]_{\bowtie p_1}, \mathtt{C}' \rangle$ and $\langle s, [\alpha]_{\bowtie p_2}, \mathtt{C}' \rangle$. Whenever the play moves to configurations of the form $\langle s', \psi_i, \mathtt{C}' \rangle$ for $i \in \{1, 2\}$, the strategy is memoryless, locally monotone, and cyclically monotone by induction. We start with proving local monotonicity for moves that may choose sub-distributions.

1. For configurations where $\alpha = \psi_1 \,\mathsf{W}\, \psi_2$, $\alpha = \psi_1 \,\mathsf{W}^{\leq k} \psi_2$, or $\alpha = \psi_1 \,\mathsf{U}^{\leq k} \psi_2$, and $\mathtt{C} = \mathtt{C}'$ we claim that the strategy composed in the proof of Theorem 2 is locally monotone by induction. Intuitively, this can be seen by the strategy using the gap δ between the probability of the formula and the required threshold. The strategy partitions this gap between all successors, so if the same state is visited with different thresholds, the partition of the gap implies that the distribution does not increase.

Let $p' = \mathsf{Prob}_M(s, \alpha)$ and $\delta_i = p' - p_i$ for $i \in \{1, 2\}$. According to the proof of Theorem 2 in configuration $\langle s, [\alpha]_{\bowtie p_i}, \mathtt{C}\rangle$ player $\mathtt{C}$ chooses the distribution

$$d_i(t) = \max(0, (\mathsf{Prob}_M(t, \alpha) - \delta_i) \cdot P(s, t))$$

It follows that if $p_1 \geq p_2$, then for every $t \in S$ we have $d_1(t) \geq d_2(t)$. If follows that if $p_1 = p_2$, then $d_1 = d_2$. Consider the case that $p_1 > p_2$. Then $p_1 > 0$ and for some t we have $d_1(t) > 0$ and $d_1(t) = \mathsf{Prob}_M(t, \alpha) - \delta_1$. As $\delta_1 < \delta_2$ and $d_2(t) = \mathsf{Prob}_M(t, \alpha) - \delta_2$ it follows that $d_1(t) > d_2(t)$.

2. For the case where $\alpha = \psi_1 \mathsf{U}\, \psi_2$ and $\mathtt{C} = \mathtt{C}'$, the strategy as defined in the proof of Theorem 2 is not locally monotone. We modify it as follows: For every configuration $\langle s, [\psi_1 \mathsf{U}\, \psi_2]_{>p}, \mathtt{C}\rangle$ the sub-distribution d is chosen according to the minimal k such that some fraction of $\mathsf{Prob}_{M_k^s}(s, \alpha)$ is greater than p. The exact definition of this fraction is given below. Furthermore, we use the gap between $\mathsf{Prob}_{M_k^s}(s, \alpha)$ and $\mathsf{Prob}_{M_{k-1}^s}(s, \alpha)$ to ensure local (and later cyclic) monotonicity. The definition of the sub-distribution d and the proof itself are quite technical.

Consider the configuration $\langle s, [\alpha]_{>p}, \mathtt{C}\rangle$. We assume, without loss of generality, that $s \notin \llbracket \psi_2 \rrbracket_M$. We measure the exact probability to satisfy α within i steps. For every $t \in S$ let

$$n_0^t = \mathsf{Prob}_{M_0^t}(t, \alpha) \qquad n_i^t = \mathsf{Prob}_{M_i^t}(t, \alpha) - \mathsf{Prob}_{M_{i-1}^t}(t, \alpha) \quad (i > 0)$$

Consider the following increasing sequence:

$$N_0^t = \tfrac{n_0^t}{2} \qquad N_i^t = N_{i-1}^t + \textstyle\sum_{j=0}^{i} \tfrac{1}{2^{i+1-j}} n_j^t \quad (i > 0)$$

That is, $N_1^t = \frac{3}{4}n_0^t + \frac{1}{2}n_1^t$, $N_2^t = \frac{7}{8}n_0^t + \frac{3}{4}n_1^t + \frac{1}{2}n_2^t$, $N_3^t = \frac{15}{16}n_0^t + \frac{7}{8}n_1^t + \frac{3}{4}n_2^t + \frac{1}{2}n_1^t$, and so on. Notice that $\lim_{i \to \infty} N_i^t = \mathsf{Prob}_{M_k^t}(t, \alpha)$. Let i_0 be minimal such that $\sum_{t \in S} N_{i_0}^t \cdot P(s, t) > p$. By abuse of notation for $i \geq 0$, we write $N_{i+1}^s = \sum_{t \in S} N_i^t \cdot P(s, t)$. That is, N_i^s is the sum of the different N_{i-1}^t normalized by their probabilities to get from s to t. To simplify notations, for $i < 0$ and for all t we set $N_i^t = N_{i+1}^s = 0$. The value $N_{i_0}^t \cdot P(s, t)$ is going to be the basis for defining $d(t)$. Notice that it must be the case that $N_{i_0}^s \leq p$ and that $N_{i_0}^t - N_{i_0-1}^t > 0$. In order to maintain local monotonicity we distribute the gap between the required threshold p and $N_{i_0}^s$ between all the states t where $N_{i_0+1}^t > 0$. We have to be extremely careful with the states s for which $N^s_{i_0} = p$. For these states, we take a constant fraction of $N^t_{i_0} - N^t_{i_0-1}$ and distribute it among the successors t. We then have to scale the distribution d for all states s for which this constant fraction surpasses the required bound.

We set $d(t)$ as follows:

$$d(t) = \left(N^t_{i_0-1} + \left(\frac{1}{4} + \frac{3}{4}\frac{p - N^s_{i_0}}{N^s_{i_0+1} - N^s_{i_0}}\right)\left(N^t_{i_0} - N^t_{i_0-1}\right)\right) \cdot P(s,t)$$

It is simple to see that $\sum_{t\in S} d(t) > p$. Indeed, $\sum_{t\in S} d(t)$ is the sum of the following three expressions:

$$\sum_{t\in S} N^t_{i_0-1} \cdot P(s,t) = N^s_{i_0}$$

$$\sum_{t\in S} \frac{N^t_{i_0} - N^t_{i_0-1}}{4} \cdot P(s,t) = \frac{N^s_{i_0+1} - N^s_{i_0}}{4}$$

$$\sum_{t\in S} \frac{3}{4}\frac{p - N^s_{i_0}}{N^s_{i_0+1} - N^s_{i_0}} \cdot (N^t_{i_0} - N^t_{i_0-1}) \cdot P(s,t) = \frac{3}{4}(p - N^s_{i_0})$$

As $N^s_{i_0+1} > p$ the result follows.

Furthermore, when going to some successor t of s the choice of i_0 for s implies that for the choice of the sub-distribution d for t some value $i'_0 < i_0$ is going to be used. Thus, the sequence of configurations of the form $\langle t', [\alpha]_{>p'}, \mathtt{C}\rangle$ is finite and player $\mathtt{C}$ is winning.

We show that this definition of the sub-distribution d implies local monotonicity. Consider two configurations $\langle s, [\alpha]_{>p_1}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{>p_2}, \mathtt{C}\rangle$. Let d_1 and d_2 be the sub-distributions chosen by σ in these configurations and let i^1_0 and i^2_0 be the values used to define d_1 and d_2, respectively. By definition, $d_j(t)$ is in the open interval $(N^t_{i^j_0-1}P(s,t), N^t_{i^j_0}P(s,t))$ for $j \in \{1,2\}$. By definition, if $p_1 = p_2$, then $i^1_0 = i^2_0$ and it follows that $d_1 = d_2$. Similarly, if $p_1 > p_2$, then $i^1_0 \geq i^2_0$. If $i^1_0 > i^2_0$, the strictness of $d_1 > d_2$ follows from the strictness of the sequence N^t_i. If $i^1_0 = i^2_0$, then $d_1 > d_2$ as $p_1 > p_2$.

Cyclic monotonicity of modified winning strategy. We turn now to consider cyclic monotonicity. Consider the configurations $\langle s, [\alpha]_{\bowtie p_1}, \mathtt{C}'\rangle$ and $\langle s, [\alpha]_{\bowtie p_2}, \mathtt{C}'\rangle$ that appear in a play consistent with σ according to this order.

1. Consider the case where $\alpha = \psi_1 \,\mathsf{W}\, \psi_2$, $\alpha = \psi_1 \,\mathsf{W}^{\leq k} \psi_2$; or $\alpha = \psi_1 \,\mathsf{U}^{\leq k} \psi_2$ and $\mathtt{C} = \mathtt{C}'$. The strategy defined in the proof of Theorem 2 is also cyclically monotone. Indeed, from configuration $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ where $\mathsf{Prob}_M(s,\alpha) - p = \delta$ we pass to configuration $\langle t, [\alpha]_{\bowtie p'}, \mathtt{C}\rangle$ and we know that $\mathsf{Prob}_M(t,\alpha) - p' = \delta$. Hence, if configurations $\langle s, [\alpha]_{\bowtie p_1}, \mathtt{C}\rangle$ and $\langle s, [\alpha]_{\bowtie p_2}, \mathtt{C}\rangle$ appear in the same play, we have $p_1 \geq p_2$.

2. Consider the case where $\alpha = \psi_1 \,\mathsf{U}\, \psi_2$ and $\mathtt{C} = \mathtt{C}'$ and the strategy is as defined above. Let i_0^1 be the bound used for choosing the sub-distribution d in configuration $\langle s, [\alpha]_{>p_1}, \mathtt{C}\rangle$. By construction, values smaller than i_0^1 are going to be used to define the sub-distributions in successor configurations. It follows that if configuration $\langle s, [\alpha]_{>p_2}, \mathtt{C}\rangle$ is visited, a value $i_0^2 < i_0^1$ is going to be used to define its sub-distribution. From the strictness of the sequence N_i^t (and N_i^s), and as $N^s_{i_0^j} \leq p_j < N^s_{i_0^j+1}$, it follows that $p_2 < p_1$.

3. Consider the case where $\alpha = \psi_1 \,\mathsf{U}\, \psi_2$, $\alpha = \psi_1 \,\mathsf{U}^{\leq k}$; or $\alpha = \psi_1 \,\mathsf{W}^{\leq k} \psi_2$ and $!\mathtt{C} = \mathtt{C}'$. Let $p' = \mathsf{Prob}_M(s', \alpha)$ and $\delta_i = p_i - p'$ for $i \in \{1, 2\}$. Let d be the distribution suggested by player $!\mathtt{C}$ in configuration $\langle s, [\alpha]_{\bowtie p_1}, !\mathtt{C}\rangle$. By definition of d, we have $\sum_{t \in S} d(t) \geq p_1$. By assumption, $\langle s, [\alpha]_{\bowtie p_2}, !\mathtt{C}\rangle$ is reachable from $\langle s, [\alpha]_{\bowtie p_1}, !\mathtt{C}\rangle$, so both players do not choose to go to configurations of the form $\langle t, \psi_i, !\mathtt{C}\rangle$ for $i \in \{1, 2\}$. If follows that

$$\mathsf{Prob}_M(s,\alpha) = \sum_{t \in S} P(s,t) \cdot \mathsf{Prob}_M(t, \alpha)$$

We know that $\sum_{t \in S} d(t) \geq p' + \delta_1$. Then, there must exist some $t \in S$ such that

$$d(t) \cdot P(s,t)^{-1} \geq \mathsf{Prob}_M(t,\alpha) + \delta_1$$

It follows that if player $\mathtt{C}$ chooses this state t, the gap between the actual probability and the threshold does not decrease. Thus $p_1 \leq p_2$.

4. Consider the case where $\alpha = \psi_1 \,\mathsf{W}\, \psi_2$ and $!\mathtt{C} = \mathtt{C}'$. Then the proof is similar to the previous item. By assumption, $\mathtt{C}$ wins from $\langle s, [\alpha]_{\geq p_1}, !\mathtt{C}\rangle$ and hence $s \notin \llbracket [\alpha]_{\geq p_1} \rrbracket_M$. Let $p' = \mathsf{Prob}_M(s, \alpha)$. As player $\mathtt{C}$ wins from $\langle s, [\alpha]_{\geq p_1}, !\mathtt{C}\rangle$, we conclude that $p' < p_1$. In particular, $s \notin \llbracket [\psi_1 \,\mathsf{W}\, \psi_2]_{>p'} \rrbracket_M$. Let $n \in \mathbb{N}$ be such that $p' < p + 1/n < p$. By Corollary 1, there are $k, l \geq 0$ with $s \notin \llbracket [\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p' + 1/n} \rrbracket_{M^s_{k,l}}$ and so the probability of $\psi_1 \,\mathsf{W}\, \psi_2$ in $M^s_{k,l}$ at s is less than p_1. Player $\mathtt{C}$ is going to use system $M^s_{k,l}$ to guide her decisions. As usual $\mathsf{Prob}_{M^s_{k,l}}(s,\alpha)$ is equal to the sum $\sum_{t \in S_{k,l}} P(s,t) \cdot \mathsf{Prob}_{M^s_{k,l}}(t,\alpha)$. Let

$$p'' = \mathsf{Prob}_{M^s_{k,l}}(s,\alpha)$$

As mentioned $p'' < p_1$. Let $\delta_1 = p_1 - p''$ and let d be the distribution suggested by player $!\mathtt{C}$ in configuration $\langle s, [\alpha]_{\geq p_1}, !\mathtt{C}\rangle$. By definition of d, we have $\sum_{t \in S} d(t) \geq p_1 = \delta_1 + p''$. Then, there must exist some $t \in S$ such that

$$d(t) \cdot P(s,t)^{-1} \geq \mathsf{Prob}_{M^s_{k,l}}(t,\alpha) + \delta_1$$

Thus, if player $\mathtt{C}$ chooses this state t, the gap between the actual probability in $M^s_{k,l}$ and the threshold doesn't decrease. In Lemma 4 below, we prove that the probability of α increases when revisiting the state in $M^s_{k,l}$. Hence, $p_2 > p_1$. □

Lemma 4 *Let M be a labeled Markov chain, q and r in $\mathbb{AP}$, α the path formula $q \,\mathsf{W}\, r$, and $M^s_{k,l}$ given for some state s of M and $k, l \in \mathbb{N}$. Let t and t' be different states in $M^s_{k,l}$ that both correspond to some state s' of M such that*

- *there is a path from t to t' in $M^s_{k,l}$, and*
- *q holds throughout the unique and finite path from the root of $M^s_{k,l}$ to t'.*

If we have $\mathsf{Prob}_{M^s_k}(t,\alpha) < 1$, then $\mathsf{Prob}_{M^s_k}(t',\alpha) > \mathsf{Prob}_{M^s_k}(t,\alpha)$ follows.

PROOF OF LEMMA 4. As $\mathsf{Prob}_{M^s_k}(t, q \,\mathsf{W}\, r) < 1$ it follows that there is some "leaf" t'' in $M^s_{k,l}$ that is reachable from t in $M^s_{k,l}$ such that the unique finite path from t to t'' in $M^s_{k,l}$ does not satisfy $q \,\mathsf{W}\, r$. As $M^s_{k,l}$ is an unwinding of M, it follows that the subtree reachable from t' in $M^s_{k,l}$ is contained in the subtree reachable from t in $M^s_{k,l}$. Clearly, $\mathsf{Prob}_{M^s_{k,l}}(t',\alpha) \geq \mathsf{Prob}_{M^s_{k,l}}(t,\alpha)$. Indeed, if a path satisfies $q \,\mathsf{W}\, r$ then every prefix of the path also satisfies $q \,\mathsf{W}\, r$. We use proof by contradiction to argue that there is a path from t that does not satisfy $q \,\mathsf{W}\, r$ and does not pass through t'. Assume such a path does not exist. Then every path beginning in t that does not satisfy $q \,\mathsf{W}\, r$ has to pass through t'. However, both t and t' correspond to state s' in M. It follows that the only option to falsify $q \,\mathsf{W}\, r$ in game $\mathsf{G}_M(s',\alpha)$ is by "going in a loop" from state s' to itself. But by assumption all states on the path between t and t' satisfy q, a contradiction. □

Example 8 *The winning strategy for player* $\mathtt{R}$ *in Example 7 is locally monotone as* $\mathtt{R}$ *never meets a pair of configurations that need to be checked for local monotonicity. That strategy is also cyclically monotone: From configuration* $\langle s_0, [q \,\mathsf{U}\, r]_{>p}, \mathtt{V}\rangle$*, the only possible cycles lead to configurations* $\langle s_0, [q \,\mathsf{U}\, r]_{>p'}, \mathtt{V}\rangle$*. As explained already, Verifier is restricted to* $d(s_2) = 0$ *and* $d(s_1) < 1/3$ *or she loses in the next step. Let* $p > 1/2$ *and* $\epsilon = 1/3 - d(s_1)$*. Then* $d(s_0) \geq 1/6 + (p - 1/2) + \epsilon$*. Thus,* $p' \geq 1/2 + 3(p - 1/2) + 3\epsilon$ *in the next configuration* $\langle s_0, [q \,\mathsf{U}\, r]_{>p'}, \mathtt{V}\rangle$*. As* $\epsilon > 0$ *and* $p - 1/2 > 0$ *we have* $p' > p$*. Finally, if* $p_1, p_2, \ldots$ *is the sequence of bounds obtained in this manner, then* $p_{i+2} - p_{i+1} > p_{i+1} - p_i$ *for all* $i \geq 1$*.*

5. Discussion

Table 1 summarizes which PCTL sub-formulae can always be coerced into finite plays if the winning player plays according to a winning strategy. For example, a Strong Until with strict bound is ensured to have a finite strategy and explore a finite portion of the game before going to sub-formulae, and similarly from a negated Weak Until with a non-strict bound. To determine whether a PCTL formula is won by means of such finite plays only, we can either convert it into "GreaterThan" normal form and check whether each such sub-formula has a negation polarity that corresponds to the desired player in that table, or we can convert it into negation normal form and interpret that table *as is* on the resulting sub-formulae. As already discussed, one can change the strictness of a threshold bound by slightly changing the required probabilities in the formula. Thus, an ϵ-correction may change a formula that does not allow finite plays to a formula that does allow finite plays. Note that the operator $\mathsf{X}_{\geq}$ does not lead to inifinite plays but may lead to using infinite sets of states.

For example, formula $\eta = [q \,\mathsf{U}\, r]_{>0.999} \wedge \neg[q \,\mathsf{W}\, r]_{\geq 0.9991}$ is such that player $\mathtt{V}$ can win be ensuring only finite plays, if she can win at all. Furthermore, if the Markov chain is infinite, the game explores only a finite portion of it. In future work, we will demonstrate that this leads to a completeness result for abstraction: abstractions are *finite-state* labeled Markov chains A where the labeling function L has type $L\colon \mathbb{AP} \times S \to \{0, 1, \bot\}$ (instead of $L\colon \mathbb{AP} \to \mathbb{P}(S)$), there is a notion of satisfaction between PCTL formulae and abstract models A, the abstraction relation $(A, a) \prec (M, s)$ for countable labeled Markov chains M is a variant of Larsen-Skou probabilistic simulation [15]; and "completeness" means if there is η for which Verifier can force finite plays in all M, then $s \in [\![\eta]\!]_M$ implies there is some abstraction A with state a with $(A, a) \prec (M, s)$ where (A, a) satisfies η.

Table 1: Sub-formulae that result in finite plays (✓) or don't (✗), for which winning player; ticks in parentheses indicate finite plays after an initial ϵ-correction of bounds

	$X_>$	$X_\geq$	$W_>$	$W_\geq$	$U_>$	$U_\geq$
Verifier	✓	✗ (✓)	✗	✗	✓	✗ (✓)
Refuter	✗	✗	✗ (✓)	✓	✗	✗

It is known that these 3-valued labeled Markov chains and probabilistic simulation cannot render such completeness for all of PCTL [20]. Future work will therefore also attempt to generalize these abstractions to a kind of tree automata such that we secure completeness for the entire logic PCTL.

6. Related work

In [6], finite-state (discrete-time) labeled Markov chains and probabilistic CTL (PCTL) are considered in their standard semantics, and different forms of evidence are being developed for documenting the falsity of a PCTL formula in a given state. One form computes those paths that contribute most to the falsity of a formula. Another form computes most probable sub-trees to gain more precise diagnostic evidence. Both forms, studied for Strong and Weak Until, are supported with shortest-path type algorithms for computing such evidence.

In [2], the line of work from [6] is being pushed into the world of Markov decision processes, with a focus on upwards-bounded probability thresholds in PCTL formulae – whereas we study the downwards-bounded case without loss of generality. The shortest-path algorithms in [2] are then combined with AND/OR trees in order to filter the computed set of paths to one with high explanatory value, and to compute the probability of that filtered path set.

In [22], bounded model-checking techniques are applied to the generation of counter-examples for probabilistic reachability properties. These techniques are combined with optimizations such as loop-detection to speed up that computation and to contain the size of these counter-example path sets.

In [8], the soundness of probabilistic counter-examples based on simulation preorders of [10, 19], represented as finite-state Markov chains, appeals to properties of the possibly infinitely many concretizations of that finite-state Markov chain. An alternative approach is that proposed in [11], where finite, stochastic, 2-person games G are used as abstractions of Markov decision processes M. These games have a satisfaction relation for PCTL that is sound with respect to abstraction. Therefore, the winning strategies that witness such satisfaction $G \models \phi$ are guaranteed to transfer into winning strategies that witness the satisfaction $M \models \phi$ for the model M that G abstracts. This is an incomplete abstraction method in the sense discussed in Section 5 of [11].

In [5], a quantitative μ-calculus with an explicit discount operator, and with models whose transitions are labeled with discount factors has non-negative real numbers as results of model checks. Quantitative parity games are developed and shown to correspond to model checks for formulae of the quantitative μ-calculus. However, winning strategies are no longer memoryless in general as they may have to "make up" for discount factors encountered en-route in a play – even in games with finite set of configurations.

In [17], a quantitative μ-calculus (qMμ) is defined over models that contain both non-deterministic and probabilistic choice but no discounting. A denotational semantics generalizing Kozen's familiar one [13] is given. For any finite-state model and formula of qMμ a probabilistic analogue of parity games is given, the determinacy of this game is shown. It is also proved that its game value equals that of the denotational semantics for the model and formula in question and that there exist memoryless winning strategies.

This paper is a journal version of the paper [4].

7. Conclusions

We captured the denotational PCTL semantics over countably labeled Markov chains through Hintikka games with Büchi acceptance conditions. This therefore renders an operational account of truth and falsity of PCTL model checks on such models in terms of winning strategies for the players Verifier and Refuter (respectively). Game moves depend on the strictness or non-strictness of probability thresholds for path formulae. Winning strategies may be assumed to be memoryless and monotone in their choice of structural elements (here sub-distributions). PCTL formulae in "GreaterThan" normal form that contain Until operators with a certain combination of threshold type and negation polarity – statically derived from Table 1 – have winning strategies that may be interpreted as a finitary witness of the falsity (respectively, truth) of the formula under consideration.

Acknowledgments. We acknowledge support by DFG project (SFB/TR14 AVACS), the UK EPSRC project *Efficient Specification Pattern Library for Model Valida-*

tion (EP/D50595X/1), the UK EPSRC project *Complete and Efficient Checks for Branching-Time Abstractions* (EP/E028985/1), and the Computing Laboratory at Oxford University (which hosted the second author during his sabbatical leave).

References

[1] IEEE standard for a high performance serial bus, August 1996, Std 1394-1995.

[2] H. Aljazzar, S. Leue, Counterexamples for model checking of Markov decision processes, Technical Report soft-08-01 (abstract), University of Konstanz, December 2007.

[3] J. Desharnais, A. Edalat, P. Panangaden, Bisimulation for Labelled Markov Processes, Information and Computation 179 (2002) 163–193.

[4] H. Fecher, M. Huth, N. Piterman, D. Wagner, Hintikka Games for PCTL on Labeled Markov Chains, in: Proc. 5th Intl. Conf. Quantitative Evaluation of SysTems (QEST), 14-17 September 2008, St Malo, France, pp. 169–178.

[5] D. Fischer, E. Grädel, L. Kaiser, Model checking games for the quantitative μ-calculus, in: Proc. 25th Annual Symposium Theoretical Aspects of Computer Science (STACS 2008), pp. 301–312. arXiv:0802.2871v1 [cs.LO], 2008.

[6] T. Han, J.-P. Katoen, Counterexamples in probabilistic model checking, in: Proc. 13th Intl. Conf. Tools and Algorithms for the Construction and Analysis of Systems, 24 March - 1 April 2007, Braga, Portugal, pp. 72–86.

[7] H. Hansson, B. Jonsson, A logic for reasoning about time and reliability, Formal Aspects of Computing 6 (1994) 512–535.

[8] H. Hermanns, B. Wachter, L. Zhang, Probabilistic CEGAR, in: Proc. 20th Intl. Conf. Computer Aided Verification (CAV'08), 7-14 July 2008, Princeton, USA, pp. 162–175.

[9] J. Hintikka, Logic, Language-Games and Information: Kantian Themes in the Philosophy of Logic, Clarendon Press, Oxford, 1973.

[10] B. Jonsson, K. G. Larsen, Specification and refinement of probabilistic processes, in: Proc. 6th Annual Symp. Logic in Computer Science (LICS '91), 15-18 July 1991, Amsterdam, The Netherlands, pp. 266–277.

[11] M. Kattenbelt, M. Huth, Abstraction Framework for Markov Decision Processes and PCTL Via Games, Technical Report RR-09-01, Oxford University Computing Laboratory, February 2009.

[12] J. G. Kemeny, J. L. Snell, A. W. Knapp, Denumerable Markov Chains (second edition), Springer Verlag, 1976.

[13] D. Kozen, Semantics of probabilistic programs, Journal of Computer and Systems Sciences 22 (1981) 328–350.

[14] M. Kwiatkowska, G. Norman, D. Parker, PRISM: probabilistic symbolic model checker, in: Proc. 12th Intl. Conf. Computer Performance Evaluation, Modelling Techniques and Tools (TOOLS 2002), 14-17 April 2002, London, United Kingdom, pp. 200–204.

[15] K. G. Larsen, A. Skou, Bisimulation Through Probabilistic Testing, in: Proc. 16th Annual ACM Symp. Principles of Programming Languages, January 1989, Austin, Texas, pp. 344–352.

[16] D. A. Martin, Borel Determinacy, Annals of Mathematics 102 (1975) 363–371.

[17] C. Morgan, A. McIver, Results on the quantitative μ-calculus qMμ, ACM Trans. Comp. Logic 8 (1) (2007) 1–43.

[18] A. Pnueli, A temporal logic of programs, Theor. Comp. Science 13 (1981) 45–60.

[19] R. Segala, N. A. Lynch, Probabilistic simulations for probabilistic processes, in: Proc. 5th Intl. Conf. Concurrency Theory, Uppsala, Sweden, 22-25 August 1994, pp. 481–496.

[20] D. Wagner, MPhil/PhD Transfer Report, Department of Computing, Imperial College London. April 2008.

[21] T. Wilke, Alternating tree automata, parity games, and modal μ-calculus, Bull. Soc. Math. Belg. 8 (2) (2001).

[22] R. Wimmer, B. Braitling, B. Becker, Counterexample Generation for Discrete-Time Markov Chains Using Bounded Model Checking, in: Proc. 10th Intl. Conf. Verification, Model Checking, and Abstract Interpretation (VMCAI'09), 18-20 January 2009, Savannah, USA, pp. 366–380.