Hintikka Games for PCTL on Labeled Markov Chains

Harald Fecher
Informatik, Universität Freiburg
fecher@informatik.uni-freiburg.de

Michael Huth Nir Piterman Daniel Wagner
Dept. of Computing, Imperial College London
{mrh, nir.piterman, dwagner}@doc.ic.ac.uk

Abstract

We present Hintikka games for formulae of the probabilistic temporal logic PCTL and countable labeled Markov chains as models, giving an operational account of the denotational semantics of PCTL on such models. Winning strategies have a decent degree of compositionality in the parse tree of a PCTL formula and express the precise evidence for truth or falsity of a PCTL formula.

We also prove the existence of monotone winning strategies that are almost finitely representable. Thus this work serves as a foundation for witness and counterexample generation in probabilistic model checking through games.

This work is also of independent interest as it displays a subtle interplay between Büchi acceptance conditions on infinite plays, the strictness or non-strictness of probability thresholds in Strong and Weak Until PCTL formulae in "GreaterThan" normal form, and a finite-state approximation lemma for Strong Until formulae with strict thresholds.

1 Introduction

Countable labeled Markov chains [8, 3] are an important class of stochastic processes for the modeling of probabilistic systems. PCTL [6] is a probabilistic temporal logic whose formulae ϕ can express practically relevant specifications, e.g. "with probability at least $1 - 1/100$, a device will be elected leader" may be a requirement within a telecommunications standard such as [1], and can be written as $[\mathsf{tt}\,\mathsf{U}\,\mathsf{someLeaderElected}]_{\geq 1-1/100}$ in PCTL. A denotational semantics $[\![\phi]\!]_M$ over labeled Markov chains M then renders truth or falsity of ϕ, where $[\![\phi]\!]_M$ 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 [10]). 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. In this paper we mean to provide a mathematical formalism that caters for just that: 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, 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 [7]. The semantics of first-order logic over models is defined as 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 each model M, a Hintikka game $\mathsf{G}(M, \phi)$ involves two players, Verifier (who wants to prove that M satisfies ϕ) and Refuter (who wants to prove that M does not satisfy ϕ). For example, in game $\mathsf{G}(M, \phi_1 \wedge \phi_2)$ Refuter has initial control and chooses a move to the continuation game $\mathsf{G}(M, \phi_1)$ or a move to the continuation game $\mathsf{G}(M, \phi_2)$. So Refuter is a "universal" player. Dually, in game $\mathsf{G}(M, \exists x\, \phi)$, Verifier is in initial control, chooses an element a of the structure in M, binds x to a, and moves to the continuation game $\mathsf{G}(M[x \mapsto a], \phi)$ for the model that is M but with x interpreted as a. So Verifier is an "existential" player. Sequences of such moves always generate finite plays since the continuation games involve proper subformulae. Eventually, a game of form $\mathsf{G}(M, R(t_1, \ldots, t_2))$ is reached for $n$ary relation symbol R and terms t_i. Verifier wins that game if the interpretation of the tuple of terms in M is contained in the interpretation of relation R in M;

otherwise Refuter wins.

Strategies for both players are objects that allow them to make necessary choices for determining continuation games. For example, Verifier needs to make choices at disjunctions and existential quantifiers. 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.

In this paper we 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.

Our games retain the idea of Verifier and Refuter as being existential and universal players (respectively), and of both having 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.

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 $s_0 s_1 \ldots$ 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. Path formulae are interpreted as predicates over paths of M. The semantics is defined as usual: a path $\pi = s_0 s_1 \ldots$ satisfies

- $\mathsf{X}\,\phi$ iff $s_1 \in [\![\phi]\!]_M$
- $\phi\,\mathsf{U}^{\leq k}\psi$ iff there is a $l \in \mathbb{N}$ such that $l \leq k$, $s_l \in [\![\psi]\!]_M$ and for all $0 \leq j < l$ we have $s_j \in [\![\phi]\!]_M$
- $\phi\,\mathsf{W}^{\leq k}\psi$ iff for all $l \in \mathbb{N}$ such that $0 \leq l \leq k$ we have either $s_l \in [\![\phi]\!]_M$ or there is $0 \leq j \leq l$ with $s_j \in [\![\psi]\!]_M$

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 $s_0 s_1 \ldots s_k$, 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 $\phi\,\mathsf{U}^{\leq\infty}\psi$, $\phi\,\mathsf{W}\,\psi$ as shorthand for $\phi\,\mathsf{W}^{\leq\infty}\psi$, and record the familiar duality between (strong) Until and Weak Until:

$$\neg(\phi\,\mathsf{U}\,\psi) \equiv (\neg\psi)\,\mathsf{W}\,(\neg\phi \wedge \neg\psi) \tag{1}$$

PCTL formulae wrap path formulae with probability thresholds (turning predicates on paths into predicates on states), and may add a propositional logic layer on top of that, which may then be used to build up new Path formulae. The operators $\vee$ (disjunction) and $\rightarrow$ (implication) are derived as usual. Let $\mathtt{ff}$ be an abbreviation for any $[\alpha]_{>1}$, and $\mathtt{tt}$ denotes any $[\alpha]_{\geq 0}$. For labeled Markov chain $M = (S, P, L)$, the denotational semantics of PCTL formula ϕ is a subset $[\![\phi]\!]_M$ of S. We write $[\![\phi]\!]$ if M is clear from the context and define $[\![\phi]\!]$ by structural induction, as usual:

$$[\![q]\!] = L(q) \qquad [\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!]$$

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

where $\mathsf{Prob}_M(s,\alpha)$ is the probability of the measurable set $\mathsf{Path}(s,\alpha)$ of paths in M that begin in s and satisfy the

$\phi, \psi ::=$	*PCTL formulae*	$\alpha ::=$	*Path formulae*
q	Atom	$\mathsf{X}\,\phi$	Next
$\neg\phi$	Negation	$\phi\,\mathsf{U}^{\leq k}\,\psi$	Until
$\phi \wedge \psi$	Conjunction	$\phi\,\mathsf{W}^{\leq k}\,\psi$	Weak Until
$[\alpha]_{\bowtie p}$	Path Probability		

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

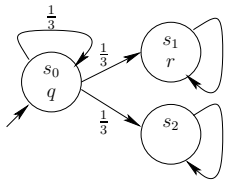

Figure 2. Labeled Markov chain M with $s_0 \in [\![q\,\mathsf{U}\,r]\!]_{\geq 1/2}[\![_M$, since $\mathsf{Prob}_M(s_0, q\,\mathsf{U}\,r) = 1/2$.

path formula α. The measure space of path sets is generated from cylinder path sets in the standard fashion [8]. Note that the semantics of PCTL and Path formulae is mutually recursive, reflecting the mutual recursion of their syntax. We say that PCTL formulae ϕ and ψ are semantically equivalent iff for all labeled Markov chains M we have $[\![\phi]\!]_M = [\![\psi]\!]_M$.

Example 1 *For the labeled Markov chain M in Figure 2 we have $[\![q\,\mathsf{U}\,r]\!]_{\geq 1/2}[\![_M = \{s_0, s_1\}$ and for the labeled Markov chain $M_2^{s_0}$ in Figure 3 we have $[\![q\,\mathsf{W}\,r]\!]_{\geq 5/9}[\![_M = \{s_0, s_0s_1, s_0s_1s_1, s_0s_0, s_0s_0s_1, s_0s_0s_0\}$.*

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

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 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 $k \leq j$, $P_{i,j}(s, t_{\text{sink}}) = 1 - \Sigma_{k=1}^{j} P_{i,j}(s, t_k)$ 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}.*

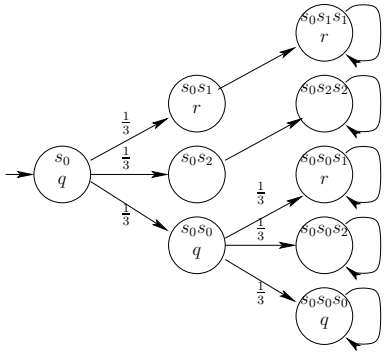

Figure 3. Unfolding $M_2^{s_0}$ of the labeled Markov chain of Figure 2 up to depth two.

Example 2 *The unfolding $M_2^{s_0}$ for the labeled Markov chain M of Figure 2 is depicted in Figure 3.*

Lemma 1 (Finite-State Approximation) *Let M be a labeled Markov chain (S, P, L), let $q, r \in \mathbb{AP}$, and $p \in [0,1]$. 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 $s \in [\![q\,\mathsf{U}\,r]\!]_{>p-1/n}[\![_{M_{k,l}^s}$.*

Example 3 *Consider the labeled Markov chain in Fig. 2. Probability $\mathsf{Prob}_M(s_0, q\,\mathsf{U}\,r) = 1/2$ is attained by paths of increasing length, as the value of the infinite sum $\Sigma_{j=1}^{\infty}(1/3)^j$. However, for every $n \in \mathbb{N}$ there exists some $i \in \mathbb{N}$ such that $\Sigma_{j=1}^{i}(1/3)^j > 1/2 - 1/n$ 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 $s_0 \in [\![q\,\mathsf{U}\,r]\!]_{>1/2-1/n}[\![_{M_2^{s_0}}$.*

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 $s_0 \in [\![[p \mathsf{U} q]_{>1/2-1/n}]\!]_{M_4^{s_0}}$. Lemma 1 promises that for every (countable) labeled Markov chain there is a similar approximation.

Lemma 1 has a dual version, required in the proof of Theorem 1: for labeled Markov chain $M = (S, P, L)$, $q, r \in \mathbb{AP}$, and $p \in [0, 1]$: $s \notin [\![[q \mathsf{W} r]_{>p}]\!]_M$ iff for all $n \in \mathbb{N}$ there are $k, l \in \mathbb{N}$ with $s \notin [\![[q \mathsf{W} r]_{\geq p+1/n}]\!]_{M^s_{k,l}}$.

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). 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 $s \in [\![\phi]\!]_M$; and won by player $\mathtt{R}$ iff $s \notin [\![\phi]\!]_M$. 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, below. Plays in game $\mathsf{G}_M(s, \phi)$ are finite or infinite sequences of elements in $\mathsf{Cf}_M(s, \phi)$ starting in the distinguished initial configuration $\langle s, \phi, \mathtt{V}\rangle$. A play is generated by game moves, specified in detail below.

Our game semantics treats Boolean connectives in the same manner as Hintikka games for first-order logic (here we take the point of view of Verifier): proving truth of formula ϕ at state s amounts to winning the game from configuration $\langle s, \phi, \mathtt{V}\rangle$. In order to prove a conjunction we allow Refuter to choose which branch of the conjunction to prove. In order to handle negation, the game continues in the same state but with the unnegated formula and a swap of the role of players, thus attempting to show that Refuter cannot win from the unnegated formula.

In games for branching-time logics such as CTL or the μ-calculus (see e.g. [13]), the universal quantification in $\forall \mathsf{X}$ is resolved by Refuter's choice of a successor state; and the existential quantification in $\exists \mathsf{X}$ is resolved by Verifier supplying one successor state, both as familiar from the case of quantifiers in first-order logic. For PCTL, however, things are more complicated. The next operator $[\mathsf{X}\phi]_{\bowtie p}$ includes a promised probability $\bowtie p$, "at least p" or "more than p". Verifier now resolves this "probabilistic quantification" by showing how to re-distribute the required probability between the successors of the state.

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 shown in Figure 2. But in order to prove this we have to appeal to the entire infinite sum $\Sigma_{i=1}^{\infty}(1/3)^i$. Any fairness constraint forcing a transition from s_0 into $\{s_1, s_2\}$ would cut 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 $\mathtt{V}$ 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 at most $p - \epsilon$. Thus, player $\mathtt{R}$ chooses the ϵ and player $\mathtt{V}$ 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.

We next define the notion of the closure $\mathsf{cl}(\phi)$ of a PCTL formula ϕ, which is the union of two sets of PCTL formulae. The first set $\mathsf{cl}_1(\phi)$ is the actual set of sub-PCTL-formulae of ϕ, including ϕ itself. The second 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 $>$, and for some $p \in [0, 1]$ and $\bowtie' \in \{>, \geq\}$ we have $[\alpha]_{\bowtie' p} \in \mathsf{cl}_1(\phi)$,

(b) α is $\psi_1 \mathsf{W} \psi_2$, $\bowtie$ is $\geq$, and for some $p \in [0, 1]$ and $\bowtie' \in \{>, \geq\}$ we have $[\alpha]_{\bowtie' p} \in \mathsf{cl}_1(\phi)$,

(c) α is $\psi_1 \mathsf{U}^{\leq k'} \psi_2$ and for some $p \in [0, 1]$ and a finite $k > k'$ we have $[\psi_1 \mathsf{U}^{\geq k} \psi_2]_{\bowtie p} \in \mathsf{cl}_1(\phi)$,

(d) α is $\psi_1 \mathsf{W}^{\leq k'} \psi_2$ and for some $p \in [0, 1]$ and a finite $k > k'$ we have $[\psi_1 \mathsf{U}^{\geq k} \psi_2]_{\bowtie p} \in \mathsf{cl}_1(\phi)$

The second set $\mathsf{cl}_2(\phi)$ allows us to replace any probability thresholds p with other values $p' \in [0, 1]$ and finite time bounds with smaller ones, but to allow this in such a manner that it is consistent with the above intuition behind ϵ-moves:

- (strong) Until formulae with non-strict bounds may change to (strong) Until formulae with strict bounds
- Weak Until formulae with strict bounds may change to Weak Until formulae with non-strict bounds, and

- the finite time bounds in bounded untils should be allowed to decrease.

The difference between the strong and weak untils stems from their duality, the negation of a Weak Until formula is a (strong) Until formula and vice versa as seen in (1). Thus, a Weak Until formula with strict bound is the negation of a (strong) Until formula with non-strict bound. When Refuter is trying to disprove a Weak Until formula with strict bound, she is in fact trying to prove the dual (strong) Until formula with non-strict bound, and requires the same possible moves for the non-strict bound and strict bound versions.

Example 5 *Consider the following formula:*

$$\phi = [(r \wedge [\,\mathsf{X}\,[(p \wedge \neg r)\,\mathsf{W}\,(q \wedge \neg r)]_{\geq 1}]_{>0})\,\mathsf{W}\,\mathsf{ff}]_{>0} \quad (2)$$

Intuitively, ϕ says there is an infinite path labeled by r such that every state on this path has a successor for which $p\,\mathsf{W}\,q$ holds on (almost) all paths on which r does not hold during verification of $p\,\mathsf{W}\,q$. Let $\alpha = (p \wedge \neg r)\,\mathsf{W}\,(q \wedge \neg r)$, $\beta = \mathsf{X}\,[\alpha]_{>0}$, and $\gamma = (r \wedge [\beta]_{>0})\,\mathsf{W}\,\mathsf{ff}$. The closure of ϕ is:

$$\mathsf{cl}(\phi) = \left\{ \begin{array}{l} \phi, [\gamma]_{\geq b}, \mathsf{ff}, (r \wedge [\beta]_{>0}), \\ [\beta]_{>0}, [\alpha]_{\geq b}, (p \wedge \neg r), \\ p, \neg r, r, (q \wedge \neg r), q \end{array} \middle| \; b \in [0,1] \right\}$$

As γ appears in ϕ with a strict bound, it is in the closure of ϕ with its original bound as well as with all possible non-strict bounds. As α appears in ϕ with a non-strict bound, it appears in the closure of ϕ only with non-strict bounds.

Similarly, for formula $\phi = [q\,\mathsf{U}\,r]_{\geq 1/2}$ we have $\mathsf{cl}(\phi) = \{\phi, q, r, [q\,\mathsf{U}\,r]_{>b} \mid b \in [0,1]\}$. As ϕ is a strong until with non-strict bound it is part of $\mathsf{cl}_1(\phi)$ and for every possible bound b its strict counterpart $[q\,\mathsf{U}\,r]_{>b}$ is in $\mathsf{cl}_2(\phi)$.

Subsequently, we write $!\mathtt{C}$ for the player other than $\mathtt{C}$, i.e. $!\mathtt{V} = \mathtt{R}$ and $!\mathtt{R} = \mathtt{V}$. The possible moves of game $\mathsf{G}_M(s_0, \phi)$ are defined through the moves of games $\mathsf{G}_M(s, \psi)$ by structural induction on $\psi \in \mathsf{cl}(\phi)$, simultaneously for all $s \in S$.

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

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

We may therefore assume that in subsequent moves configurations of the form $\langle s, [\alpha]_{\bowtie p}, \mathtt{C}\rangle$ never satisfy that $\bowtie\, p$ equals ≥ 0 or > 1.

M3. At configurations $\langle s, \mathsf{q}, \mathtt{C}\rangle$:

- player $\mathtt{C}$ 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 $\langle s, \psi, !\mathtt{C}\rangle$

So move M4 removes the negation from the formula but also swaps the role of players.

M5. At configuration $\langle s, \psi_1 \wedge \psi_2, \mathtt{C}\rangle$, player $!\mathtt{C}$ can choose as next configuration either $\langle s, \psi_1, \mathtt{C}\rangle$ or $\langle s, \psi_2, \mathtt{C}\rangle$

So player $!\mathtt{C}$ chooses a conjunct and the game continues with that conjunct instead of the conjunction.

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

- if $P(s, s') = 0$, player $!\mathtt{C}$ wins
- otherwise, $P(s, s') > 0$ and the next configuration is $\langle s', \psi, \mathtt{C}\rangle$

Move M6 is well defined. There is a non-empty set Y with $P(s, Y) \bowtie p$ as $p \in [0, 1]$, $P(s, \cdot)$ has mass one, and $\bowtie\, p$ is neither equal to > 1 nor to ≥ 0.

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

In move M7 such a choice is possible 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 over $n \in \mathbb{N}$.

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

In move M8 such a choice is possible 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 (1)), so it is like an *existential* quantification over $n \in \mathbb{N}$.

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 $\mathtt{C}$ is able to move to next configuration $\langle s, \psi_2, \mathtt{C}\rangle$
- if player $\mathtt{C}$ did not move, player $!\mathtt{C}$ is able to move to next configuration $\langle s, \psi_1, \mathtt{C}\rangle$
- if neither player moved above, the play must proceed as follows:

  Player $\mathtt{C}$ chooses a sub-distribution $d\colon S \to [0, 1]$ such that

$$\sum_{s'\in S} d(s') > 0 \quad \& \quad \sum_{s'\in S} d(s') \geq p \qquad (3)$$

$$\forall s' \in S\colon d(s') \leq P(s, s') \qquad (4)$$

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

In move M9, sub-distribution d 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 move M9 by (4).

M10. At configuration $\langle s, [\alpha]_{>p}, \mathtt{C}\rangle$ where α is either $\psi_1 \,\mathsf{U}^{\leq k}\, \psi_2$ 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 above; except when the last item of M9 applies, in which case the counter k in α is decreased to $k-1$ for that next configuration $\langle s', [\alpha]_{\bowtie d(s')\cdot P(s,s')^{-1}}, \mathtt{C}\rangle$

In move M10, a Bounded Until with bound 0 has to realize ψ_2 right away; and a Bounded Weak Until with bound zero has to realize at least one of ψ_1 or ψ_2 right away.

A finite play is won as explained in M1-M10 above. In most moves, the play either ends or moves to configurations where the formula is a *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.

It follows that every infinite play ends 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 C nor the sub-formulae ψ_1 and ψ_2 change.

Definition 2 (Acceptance conditions) *Player* 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* 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* $\mathtt{C} = \mathtt{V}$*.*

These are Büchi type acceptance conditions, and so our games are known to be determined [11]. We use the notion of strategy for player 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 6 *Consider the game* $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{\geq 1/2})$*, where* M *is as in Fig. 2, and let* $\alpha = q \,\mathsf{U}\, r$*. The initial configuration is* $\langle s_0, [\alpha]_{\geq 1/2}, \mathtt{V}\rangle$*. In the first move player* R *chooses an* $n \in \mathbb{N}$ *and the next configuration is* $\langle s_0, [\alpha]_{>1/2-1/n}, \mathtt{V}\rangle$*. 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* 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* 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* 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* R *to move to* s_1 *and lose. This describes a winning strategy for player* V *in game* $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{\geq 1/2})$*.*

Example 7 *Although the choice of* d *in Example 6 may seem arbitrary, it meshes well with the use of Lemma 1. Consider again the game* $\mathsf{G}_M(s_0, [\alpha]_{\geq 1/2})$ *from Example 6. Suppose that in the first move player* R *chooses* $9 \in \mathbb{N}$*, and the next configuration is* $\langle s_0, [\alpha]_{>7/18}, \mathtt{V}\rangle$*. Since for the* $M_2^{s_0}$ *in Figure 3,* $\mathsf{Prob}_{M_2^{s_0}}(s_0, \alpha) = \frac{4}{9} > \frac{7}{18}$*, player* V *can use* $M_2^{s_0}$ *to guide her choices. In* $M_2^{s_0}$ *we have* $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_1, \alpha) = 1$ *and* $\mathsf{Prob}_{M_2^{s_0}}(s_0 s_0, \alpha) = \frac{1}{3}$*. Player* 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* 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* $s_0 s_1 \in [\![r]\!]_{M_2^{s_0}}$ *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* 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*.*

Definition 3 *1. A strategy* σ *for player* C *in game* $\mathsf{G}_M(s,\phi)$ *is* winning from a configuration Γ *in that game iff player* C *wins all plays beginning in configuration* Γ *when player* C *plays according to his strategy* σ *– regardless of how player* !C *plays.*

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

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 1 *Let $M = (S, P, L)$ be a labeled Markov chain over $\mathbb{AP}$, $s \in S$, and ϕ a PCTL formula. Then we have:*

1. $s \in [\![\phi]\!]_M$ *iff player* V *wins game* $\mathsf{G}_M(s, \phi)$
2. $s \notin [\![\phi]\!]_M$ *iff player* R *wins game* $\mathsf{G}_M(s, \phi)$.

In particular, game $\mathsf{G}_M(s, \phi)$ is determined.

Sketch of Proof: Given PCTL formula ϕ, both items 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 items in Theorem 1 for such ψ in their "only if" versions, which consists of six cases. We illustrate 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 $\geq$:

We show for all three cases above that **(#1)** $s \in [\![\phi]\!]_M$ implies player V wins game $\mathsf{G}_M(s, \phi)$ and **(#2)** $s \notin [\![\phi]\!]_M$ implies player 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 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 V from configuration $\langle s, \psi_1, \mathtt{V}\rangle$, so if player R chooses to go to this configuration, player V wins. If player R does not move to ψ_1, then M9 demands that player V chooses a sub-distribution $d : S \to [0, 1]$ satisfying (3)-(4). 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 $p' = \Sigma_{t\in T} P(s, t) \cdot \mathsf{Prob}_M(t, \alpha)$.

- 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 $p' = \Sigma_{t\in S} \mathsf{Prob}_M(t, \alpha) P(s, t)$ 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 $\Sigma_{t\in T} d(t) \geq p' - \delta \geq p$.

- 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 V.

Now regardless of the choice of player R, the next configuration is $\langle t, [\alpha]_{\bowtie p'}, \mathtt{V}\rangle$ such that $t \in [\![[\alpha]_{\bowtie p'}]\!]_M$. So player 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 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 V wins (cf. Example 7).

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 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 $s \notin [\![\psi_2]\!]_M$ and hence there is a winning strategy for player R from configuration $\langle s, \psi_2, \mathtt{V}\rangle$. Also, it is either the case that $s \notin [\![\psi_1]\!]_M$ or $s \notin [\![[\mathsf{X}\,\alpha]_{\bowtie p}]\!]_M$. In the first case player R has a winning strategy from configuration $\langle s, \psi_1, \mathtt{V}\rangle$ and chooses this configuration. In the second case, player V chooses a sub-distribution $d\colon S \to [0, 1]$ such that (3)-(4) hold.

We claim that there is some $s' \in S$ with $d(s') > 0$ and $\mathsf{Prob}_M(s', \alpha) \not\bowtie d(s') P(s, s')^{-1}$. 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 (3). But this renders

$$\sum_{s'|d(s')>0} \mathsf{Prob}_M(s',\alpha) \bowtie p$$

which directly contradicts $s \notin [\![[\mathsf{X}\,\alpha]_{\bowtie p}]\!]_M$.

Thus, player 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 $s' \notin [\![[\alpha]_{\bowtie p'}]\!]_M$. Notice that player 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 R for the different forms of path formula α:

Case (a): For Weak Until formulae, we appeal to the dual of Lemma 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, $s \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{>p'}]\!]_M$. Let $n \in \mathbb{N}$ be such that $p' < p + 1/n < p$. By the dual of Lemma 1 there are $k, l \geq 0$ with $s \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p'+1/n}]\!]_{M^s_{k,l}}$ and so the probability of $\psi_1 \,\mathsf{W}\, \psi_2$ in $M^s_{k,l}$ at s is less than p. Player 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 V. As $s \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{\geq p}]\!]_{M^s_{k,l}}$, there is some $s' \in S$ such that $s' \notin [\![[\psi_1 \,\mathsf{W}\, \psi_2]_{\geq d(s')P(s,t)^{-1}}]\!]_{M^s_{k,l}}$. So player R chooses this s'. By definition of $M^s_{k,l}$ there can be only finite sequence of configuration of the form $\langle s', [\alpha]_{\geq p}, \mathtt{V}\rangle$, and so player R wins. This is dual to the strategy depicted for V in Example 7.

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 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 R wins by induction. □

Game $\mathsf{G}_M(s,\phi)$ is defined such that its initial configuration $\langle s, \phi, \mathtt{V}\rangle$ is owned by player V. We can define a dual game with the same moves but with initial configuration $\langle s, \phi, \mathtt{R}\rangle$. Theorem 1 and its proof then remain to be valid if we swap the role of players in both.

Example 8 *Consider game $\mathsf{G}_M(s_0, [q \,\mathsf{U}\, r]_{>1/2})$, where M is as in Fig. 2, and let $\alpha = q \,\mathsf{U}\, r$. From configuration $\langle s_0, [\alpha]_{>1/2}, \mathtt{V}\rangle$, player V won't move to $\langle s_0, r, \mathtt{V}\rangle$ as she would then lose. For the same reason, player R won't move to $\langle s_0, q, \mathtt{V}\rangle$. So if both players play strategies that are 'optimal' for them, player V has to choose a sub-distribution d at the initial configuration.*

If $d(s_2) > 0$, player V loses as player R can then choose s_2. So $d(s_2) = 0$ for any 'optimal' strategy of player 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 (4), which would violate (3). Since player V plays an 'optimal' strategy, $d(s_1) \neq 1/3$, as otherwise player 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 (4) there is therefore $\epsilon > 0$ such that $d(s_1) = 1/3 - \epsilon$. In particular, player 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 R chooses s_0 and the next configuration is $\langle s_0, [\alpha]_{>3d(s_0)}, \mathtt{V}\rangle$. By (3), $3d(s_0)$ must be at least $1/2 + 3\epsilon$ and so player 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 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 V as described above; or the play is infinite and so lost for player V by the acceptance conditions A1 on infinite plays.

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

- *player R will never choose a configuration of form $\langle s_0, q, \mathtt{V}\rangle$, should such an opportunity arise*
- *whenever player V chooses sub-distribution d with $d(s_2) > 0$, player R will choose s_2*
- *otherwise, it must be the case that both $d(s_1)$ and $d(s_2)$ are positive; if $d(s_1) = 1/3$, player R chooses s_1*
- *if $d(s_1) \neq 1/3$, player 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, they 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 $d, d' : S \to [0, 1]$, 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 4 (Locally Monotone Strategies) *A strategy σ for player $\mathtt{C}$ in game $\mathsf{G}_M(s, \phi)$ is* locally 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 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 5 (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 2 *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.*

Sketch of Proof: 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 1 is memoryless, locally monotone, and cyclically monotone. That slightly modified strategy will clearly be memoryless by construction. The proof therefore describes this modified winning strategy and first proves its local monotonicity, by induction as in the proof of Theorem 1. Then it proves that it is cyclically monotone. The utilized techniques are similar to those used in the proof of Theorem 1, but cannot be demonstrated here due to lack of space. □

Example 9 *The winning strategy for Refuter in Example 8 is locally monotone as Refuter never encounters 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, in the next configuration $\langle s_0, [q \,\mathsf{U}\, r]_{>p'}, \mathtt{V}\rangle$ we have $p' \geq 1/2 + 3(p - 1/2) + 3\epsilon$. 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. For example, formula $[q \,\mathsf{U}\, r]_{>0.999} \wedge \neg[q \,\mathsf{W}\, r]_{\geq 0.9991}$ 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. From a practical point of view, it may be possible to change the strictness of the 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.

6 Related work

In [5] finite-state (discrete-time) labeled Markov chains and probabilistic CTL (PCTL) are considered in their stan-

Table 1. Sub-formulae that result in finite plays ($\surd$) or don't (x), 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	$\surd$	x ($\surd$)	x	x	$\surd$	x ($\surd$)
Refuter	x	x	x($\surd$)	$\surd$	x	x

dard semantics, and different forms of evidence 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 [5] 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. We hope that our Hintikka games provide a suitable foundation for understanding the trade-off between the precision of extant and future forms of evidence and the complexity of their supporting algorithms.

In [4] 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 [12] 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 [9] 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.

7 Conclusions

We captured the PCTL semantics over countably labeled Markov chains through Hintikka games with Büchi acceptance conditions. 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. This research was in part supported by DFG project (SFB/TR14 AVACS) and the UK EPSRC projects *Efficient Specification Pattern Library for Model Validation* (EP/D50595X/1) and *Complete and Efficient Checks for Branching-Time Abstractions* (EP/E028985/1).

References

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

[2] H. Aljazzar and S. Leue. Counterexamples for model checking of markov decision processes. Technical Report soft-08-01, University of Konstanz, 2007.

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

[4] D. Fischer, E. Grädel, and L. Kaiser. Model checking games for the quantitative μ-calculus. In *25th STACS*, Dagstuhl Seminar Proceedings, pages 301–312. IBFI, Schloss Dagstuhl, Germany, 2008.

[5] T. Han and J.-P. Katoen. Counterexamples in probabilistic model checking. In *13th TACAS*, LNCS 4424, pages 72–86. Springer-Verlag, 2007.

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

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

[8] J. G. Kemeny, J. L. Snell, and A. W. Knapp. *Denumerable Markov Chains*. Springer Verlag, 1976. Second Edition.

[9] D. Kozen. Semantics of probabilistic programs. *Journal of Computer and Systems Sciences*, 22:328–350, 1981.

[10] M. Kwiatkowska, G. Norman, and D. Parker. PRISM: probabilistic symbolic model checker. In *12th International Conference on Computer Performance Evaluation, Modelling Techniques and Tools*, LNCS 2324, pages 200–204. Springer-Verlag, 2002.

[11] D. A. Martin. Borel Determinacy. *Annals of Mathematics*, 102:363–371, 1975.

[12] C. Morgan and A. McIver. Results on the quantitative μ-calculus qmμ. *ACM Transactions on Computational Logic (TOCL)*, 8(1), 2007.

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