Three-Valued Abstractions of Markov Chains: Completeness for a Sizeable Fragment of PCTL

Michael Huth, Nir Piterman, and Daniel Wagner

Department of Computing
Imperial College London
London, SW7 2AZ, United Kingdom
{m.huth, nir.piterman, dwagner}@doc.ic.ac.uk

Abstract. Three-valued Markov chains and their PCTL semantics abstract – via probabilistic simulations – labeled Markov chains and their usual PCTL semantics. This abstraction framework is *complete* for a PCTL formula if all labeled Markov chains that satisfy said formula have a finite-state abstraction that satisfies it in its abstract semantics. We show that not all PCTL formulae are complete for this abstraction framework. But PCTL formulae whose path modalities occur in a suitable combination of negation polarity and threshold type are proved to be complete, were abstractions are bounded, 3-valued unfoldings of their concrete labeled Markov chains. This set of complete PCTL formulae subsumes widely used PCTL patterns.

1 Introduction

Markov chains are an important modeling formalism for systems that contain stochastic uncertainty and for which the assumption of the "Markov property" (that the transition probability at a state depends only on that state and not on the execution history of the system) is feasible. Markov chains are used in a wide range of applications, we mention biological sequence analysis, statistical software testing, and formal verification of communication protocols or probabilistic algorithms as examples.

In formal verification, we want to validate a system model (and so hopefully the system, too) by proving that it satisfies critical properties. In the context of Markov chains as models, probabilistic computation tree logic [1] has emerged as the defacto standard for expressing such properties. The semantics of that probabilistic logic over Markov chains also renders algorithms for automatically deciding the truth of formulae over *finite-state* Markov chains, leading to the now mature and established methodology of *probabilistic model checking* [2].

But the initial models of systems often have infinite state. For example, a system state may implicitly encode the value of a continuous-time clock. Since we ultimately want to validate critical properties on systems and not on models, this begs the question of whether truth of some property on an infinite-state system or model can, in principal, be witnessed as truth of that same property on a suitable finite-state model. Suitability here means that the obtained model

abstracts certain features of the system but still contains sufficient state and behavior of the system it intends to model.

We therefore study the feasibility of this approach in a formal setting, where systems are identified with infinite-state Markov chains and abstractions are finite-state Markov chains with 3-valued atomic observables such that abstraction is based on probabilistic simulation [3, 4]. In this setting, we show negative and positive existence results for finite-state witnesses of truth that depend on the interplay between path modalities (e.g. "true at all reachable states") and threshold types (e.g. "true with probability at least .999"). As we will demonstrate, these results suggest that – from a *practical* perspective – finite-state abstractions for probabilistic computation tree logic and Markov chains more often than not exist. But there may not be an algorithm for computing them.

Related work. In [5], Markov chains and their PCTL semantics are soundly abstracted into 3-valued models, and a model checking algorithm is given for their 3-valued abstract semantics of PCTL. This gives a foundation for counter-example guided abstraction refinement where abstractions have intervals (not real numbers) as probability transitions.

In [6], game-theoretic foundations for truth of PCTL formulae ϕ over Markov chains M are developed. A Hintikka game for ϕ and M, with Büchi type acceptance conditions for infinite plays, is designed so that a "Verifier" player has a winning strategy if M satisfies ϕ. Dually, a "Refuter" player has a winning strategy if M doesn't satisfy ϕ. In loc. cit. it is also observed that a winning strategy could be chosen so that it forces always finite plays for certain path modalities. This insight provides the seed for the results reported here. But proving these results doesn't require any appeal to the games and results of loc. cit.

In [7], stochastic 2-player games are used as abstractions of Markov decision processes (MDPs) and a game simulation is developed and shown to be sound for PCTL. Interestingly, they also show incompleteness in the sense of our paper (for finite games) for the PCTL formula $[\mathsf{tt}\,\mathsf{U}\,q]_{>0}$, which *is* expressible in our complete fragment. This contradiction is only apparent since the incompleteness of that formula results solely from the non-determinism in MDPs whereas our work considers Markov chains, which are deterministic.

Outline of paper. In Section 2, we provide the background – notably our abstraction framework – needed for our technical development. The key concept of "completeness" for our abstraction framework and our incompleteness results are presented in Section 3. Completeness results for a fragment of PCTL are presented in Section 4. In Section 5, we put negative and positive results into context and conclude the paper. Selected proofs are provided in an appendix.

2 Background

We define the concrete models of systems considered here.

Definition 1 (Markov chains). *A* 3-valued, labeled Markov chain M *over a countable set* $\mathtt{AP}$ *of atomic propositions is a tuple* $(S, \mathbf{P}, L)$*, where*

1. S *is a countable set of* states*,*
2. $\mathbf{P}$ *is a* stochastic matrix $\mathbf{P}\colon S \times S \longrightarrow [0,1]$ *such that the countable sum* $\sum_{s' \in S} \mathbf{P}(s,s')$ *exists and equals* 1 *for all* $s \in S$*,*
3. *and* L *is a labeling function* $L\colon S \times \mathtt{AP} \longrightarrow \{\mathtt{tt}, \mathtt{?}, \mathtt{ff}\}$*.*

M *is* finitely branching *if* $\{s' \mid \mathbf{P}(s,s') > 0\}$ *is finite for all* $s \in S$*. We write* (M, s_0) *to denote that* M *has a designated initial state* s_0*.*

Throughout we refer to 3-valued, labeled Markov chains as *models*. Such models can be seen as (possibly infinite) labeled graphs where the outgoing transitions of state s to states s' are decorated with the positive transition probabilities $\mathbf{P}(s,s')$ of the corresponding distribution $\mathbf{P}(s,\cdot)$, and vertices $s \in S$ are labeled with atomic propositions as follows: label $\mathtt{q}$? marks the states s with $L(s,\mathtt{q}) = \mathtt{?}$, label $\mathtt{q}$ at s indicates $L(s,\mathtt{q}) = \mathtt{tt}$, and absence of any $\mathtt{q}$ or $\mathtt{q}$? label at state s implicitly marks $L(s,\mathtt{q}) = \mathtt{ff}$. When all labels for M have value $\mathtt{tt}$ or $\mathtt{ff}$, we call model M a *Markov chain*, *concrete* or *2-valued*. Thus $\mathtt{?}$ abstracts both $\mathtt{tt}$ and $\mathtt{ff}$ in the familiar information ordering [9].

A widely used notion of probabilistic (bi-)simulation was defined by Larsen and Skou in [3] for probabilistic processes with actions. We define *probabilistic simulation* for our 3-valued models, based on probabilistic simulation for *probabilistic specification systems* with propositional labels in [4].

Definition 2 (Probabilistic simulation). *Let* $M = (S, \mathbf{P}, L)$ *be a model over* $\mathtt{AP}$*. Relation* $H \subseteq S \times S$ *is a* probabilistic simulation *if whenever* $(t,s) \in H$ *then*

1. $L(t,\mathtt{q}) \leq L(s,\mathtt{q})$ *for all* $\mathtt{q} \in \mathtt{AP}$*.*
2. *For each* $s \in S$ *there is a* weight function $\rho_s\colon S \times S \longrightarrow [0,1]$ *such that*
   (a) $\sum_{s' \in S}(\mathbf{P}(s,s') \cdot \rho_s(s',t')) = \mathbf{P}(t,t')$ *for all* $t' \in S$*;*
   (b) $(t',s') \in H$ *whenever* $\rho_s(s',t') > 0$ *for some* $s \in S$*.*

We often write tHs *for* $(t,s) \in H$*, and say that* t simulates s*, written* $t \preceq s$*, if there is a probabilistic simulation* H *such that* tHs*. Model* A *simulates model* M*, written* $A \preceq M$*, if this is true of their respective initial states in the model* $A + M$ *that is the disjoint sum of the models* A *and* M*.*

Definition 3 (PCTL syntax). *The syntax of PCTL is as follows:*

$$\begin{array}{lll} \phi ::= & \mathtt{q} \mid \neg\phi \mid \phi \wedge \phi \mid \phi \vee \phi \mid [\alpha]_{\bowtie p} & \textit{(state formulae)} \\ \alpha ::= & \mathsf{X}\,\phi \mid \phi\,\mathsf{U}^{\leq k}\,\phi \mid \phi\,\mathsf{W}^{\leq k}\,\phi & \textit{(path formulae)} \end{array}$$

where $\mathtt{q} \in \mathtt{AP}$*,* $p \in [0,1]$*,* $\bowtie \in \{<, \leq, \geq, >\}$ *and* $k \in \mathbb{N} \cup \{\infty\}$*. Let PCTL be the set of* state *formulae* ϕ *generated in this manner. We write* $\mathtt{tt}$ *and* $\mathtt{ff}$ *for any PCTL formulae* $[\alpha]_{\geq 0}$ *and* $[\alpha]_{>1}$*, respectively.*

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

Fig. 1. Path-formula semantics on paths $\pi = s_0 s_1 \ldots$ in interpretation $m \in \{\mathsf{o}, \mathsf{p}\}$

Intuitively, $[\alpha]_{\bowtie p}$ specifies the property that the probability of all paths (infinite sequences of states $s_0 s_1 \ldots$ with positive transition probabilities $\mathbf{P}(s_i, s_{i+1})$) that begin at state s and satisfy path formula α is $\bowtie p$. The path modalities X, U, and W stand for Next, Strong Until, and Weak Until (respectively). The value $k = \infty$ is used to express unbounded Untils, whereas $k \in \mathbb{N}$ expresses a proper step bound on Untils. We write $\phi\,\mathsf{U}\,\phi$ as a shorthand for $\phi\,\mathsf{U}^{\leq\infty}\,\phi$, and $\phi\,\mathsf{W}\,\phi$ as shorthand for $\phi\,\mathsf{W}^{\leq\infty}\,\phi$. For example, $\mathsf{X}\,\mathtt{q}$ holds in paths whose second (next) state satisfies $\mathtt{q}$, whereas $\mathtt{q}\,\mathsf{U}\,\mathtt{r}$ holds in paths that have a finite prefix of states satisfying $\mathtt{q}$ followed by a state satisfying $\mathtt{r}$, and $\mathtt{q}\,\mathsf{W}\,\mathtt{r}$ holds in paths that either satisfy $\mathtt{q}\,\mathsf{U}\,\mathtt{r}$ or where all states satisfy $\mathtt{q}$.

We define semantics for PCTL formulae based on an optimistic and a pessimistic interpretation of labels [10, 11]. Optimistically, we interpret a proposition as true if it isn't false, i. e. $[\![\mathtt{q}]\!]^{\mathsf{o}}_M = \{s \in S \mid L(s, \mathtt{q}) \neq \mathtt{ff}\}$; pessimistically, $\mathtt{q}$ is true only if the labeling says so, i. e. $[\![\mathtt{q}]\!]^{\mathsf{p}}_M = \{s \in S \mid L(s, \mathtt{q}) = \mathtt{tt}\}$.

Definition 4 (PCTL semantics). *Let $m \in \{\mathsf{o}, \mathsf{p}\}$ be two modes of interpretation, $\neg\mathsf{o} = \mathsf{p}$, and $\neg\mathsf{p} = \mathsf{o}$. For ϕ in PCTL, we define $[\![\phi]\!]^m_M$:*

$$[\![\phi \wedge \psi]\!]^m_M = [\![\phi]\!]^m_M \cap [\![\psi]\!]^m_M \qquad [\![\phi \vee \psi]\!]^m_M = [\![\phi]\!]_M \cup [\![\psi]\!]^m_M$$
$$[\![\neg\phi]\!]^m_M = S \setminus [\![\phi]\!]^{\neg m}_M \qquad [\![[\alpha]_{\bowtie p}]\!]^m_M = \{s \in S \mid \mathsf{Prob}^m_M(s, \alpha) \bowtie p\}$$

where $\mathsf{Prob}^m_M(s, \alpha)$ is the probability of the measurable set $\mathsf{Path}^m(s, \alpha)$ of paths $\pi = s_0 s_1 \ldots$ in M that begin in $s_0 = s$ and satisfy $\pi \vDash^m \alpha$, defined in Figure 1.

We often write $M^s \vDash^m \phi$ for $s \in [\![\phi]\!]^m_M$ and use $M \vDash^m \phi$ as abbreviation of $M^{s_0} \vDash^m \phi$ for initial state s_0. For 2-valued Markov chains $\vDash^{\mathsf{o}}$ equals $\vDash^{\mathsf{p}}$ and coincides with the familiar and standard PCTL semantics $\vDash$ over Markov chains.

The interpretation m is sound in that verifications of ϕ by $\vDash^{\mathsf{p}}$ on A ($A \vDash^{\mathsf{p}} \phi$) and refutations of ϕ by $\vDash^{\mathsf{o}}$ on A ($A \nvDash^{\mathsf{o}} \phi$) are verifications, respectively refutations, in any concrete M with $A \preceq M$. This soundness requires that PCTL formulae are presented in a particular normalform in which negations occur only on atomic propositions and where probability thresholds are either $\geq$ or $>$:

Definition 5 (Greater-than negation normal form). *The following subset of PCTL constitutes the* Greater-than negation normal form *(GTNNF):*

$$\phi ::= \mathtt{q} \mid \neg\mathtt{q} \mid \phi \wedge \phi \mid \phi \vee \phi \mid [\alpha]_{\bowtie p}$$
$$\alpha ::= \mathsf{X}\,\phi \mid \phi\,\mathsf{U}^{\leq k}\,\phi \mid \phi\,\mathsf{W}^{\leq k}\,\phi$$

where $\mathtt{q} \in \mathtt{AP}$, $p \in [0, 1]$, $\bowtie \in \{\geq, >\}$ and $k \in \mathbb{N} \cup \{\infty\}$.

Every formula ϕ of PCTL that is not in GTNNF can be transformed to a formula in GTNNF, equivalent in the two-valued semantics $\models$ over Markov chains, by (1) replacing each sub-formula of the form $[\alpha]_{<p}$ and $[\alpha]_{\leq p}$ by $\neg[\alpha]_{\geq 1-p}$ and $\neg[\alpha]_{>1-p}$ respectively, and then (2) pushing negations inwards. The second step, i. e. pushing negations inwards, is possible without breaking the syntactical restrictions of PCTL, only because the used definition includes both the Weak and the Strong Until. With an intermediate step into PCTL* [13] one gets:

$$\begin{aligned}
\neg[\mathsf{X}\,\phi]_{>p} &\equiv [\neg\,\mathsf{X}\,\phi]_{\geq 1-p} \equiv [\mathsf{X}\,\neg\phi]_{\geq 1-p} \\
\neg[\phi\,\mathsf{U}^{\leq k}\,\psi]_{>p} &\equiv [\neg(\phi\,\mathsf{U}^{\leq k}\,\psi)]_{\geq 1-p} \equiv [(\neg\psi)\,\mathsf{W}^{\leq k}(\neg\phi \wedge \neg\psi)]_{\geq 1-p} \\
\neg[\phi\,\mathsf{W}^{\leq k}\,\psi]_{>p} &\equiv [\neg(\phi\,\mathsf{W}^{\leq k}\,\psi)]_{\geq 1-p} \equiv [(\neg\psi)\,\mathsf{U}^{\leq k}(\neg\phi \wedge \neg\psi)]_{\geq 1-p}
\end{aligned}$$

Swapping the roles of $\geq$ and $>$ in the above equivalences yields the dualities for the remaining combinations of temporal operators and threshold types. The negations $\neg\phi$ and $\neg\psi$ above are then processed in the same manner, recursively.

We can now secure the desired soundness result:

Lemma 1. *Let M and A be models and $A \preceq M$. Then for all formulae ϕ in GTNNF we have the implications $A \models^{\mathsf{p}} \phi \Rightarrow M \models^{\mathsf{p}} \phi$ and $M \models^{\mathsf{o}} \phi \Rightarrow A \models^{\mathsf{o}} \phi$.*

This lemma is proved by structural induction on ϕ, using standard fixed-point and duality arguments for Weak and Strong Until formulae. As our paper focuses on completeness not on soundness, we don't feature this proof here.

3 Completeness for PCTL formulae

The notion of completeness we now define is relative to our class of models, their abstract PCTL semantics, and its abstraction via probabilistic simulation. We refer to this triad as "our abstraction framework" subsequently.

Definition 6 (Finitary completeness). *Our abstraction framework is* complete *for a PCTL formula ϕ iff for all Markov chains M that satisfy ϕ there is a finite-state model A such that $A \preceq M$ and $A \models^{\mathsf{p}} \phi$. Our abstraction framework is complete for a set of PCTL formulae Γ if it is complete for each $\phi \in \Gamma$.*

Completeness for ϕ thus means that all Markov chains that satisfy ϕ ($M \models \phi$) have a finite-state abstraction that also satisfies ϕ in the $\models^{\mathsf{p}}$ semantics. We chose $\models^{\mathsf{p}}$ for this definition since it, unlike $\models^{\mathsf{o}}$, is sound for verifications.

Example 1. The infinite-state Markov chain M depicted in Figure 2(a) satisfies $\varphi = [\mathsf{q}\,\mathsf{U}\,\mathsf{r}]_{>0.7}$. It is simulated by the finite-state model $M^{s_0}_{3,3}$ in Figure 2(b) and $M^{s_0}_{3,3} \models^{\mathsf{p}} \varphi$. In Section 4, we will see that φ is complete.

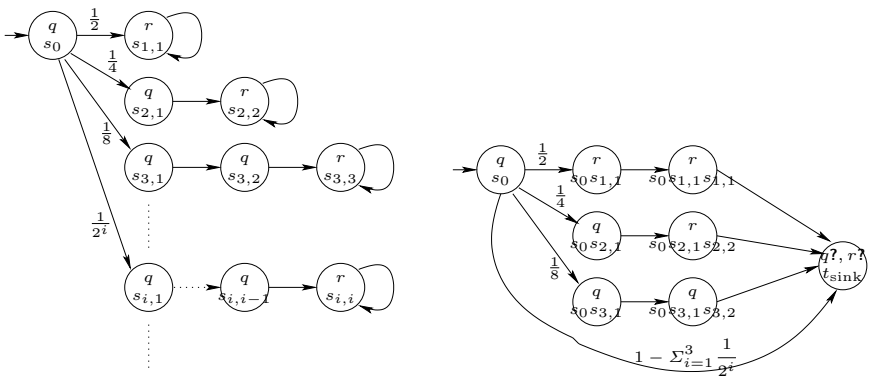

(a) M (b) $M^{s_0}_{3,3}$

Fig. 2. A model M satisfying $[\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{\geq 1}$ and $[\mathtt{q}\,\mathsf{W}\,\mathtt{r}]_{\geq 1}$, and its unfolding $M^{s_0}_{3,3}$

Incompleteness of PCTL. We show that full PCTL is incomplete by giving several counterexamples which consist of a concrete Markov chain M and a PCTL formula φ such that no finite-state model A can exist, which simulates M and for which $A \models^{\mathsf{p}} \varphi$. These examples are strongly inspired by Dams and Namjoshi's work on completeness for Kripke structures and the modal mu-calculus [12].

Lemma 2. *Not all formulae of form $[\phi\,\mathsf{U}\,\psi]_{\geq p}$ and $[\phi\,\mathsf{W}\,\psi]_{\geq p}$ are complete.*

Proof. We consider $[\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{\geq 1}$ and $[\mathtt{q}\,\mathsf{W}\,\mathtt{r}]_{\geq 1}$. Let M be the Markov chain illustrated in Figure 2(a): The initial state s_0 is labeled $\mathtt{q}$ and is infinitely branching with $\mathbf{P}(s_0, s_{i,1}) > 0$ for all $i \geq 1$; its i-th successor $s_{i,1}$ has probability $1/2^i$, all other transition probabilities are 1; the i-th path $s_0 s_{i,1} \dots s_{i,i}$ consists of $i-1$ states labeled $\mathtt{q}$ and ends in an absorbing state $s_{i,i}$ labeled $\mathtt{r}$. The Markov chain M obviously satisfies any $\varphi \in \{[\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{\geq 1}, [\mathtt{q}\,\mathsf{W}\,\mathtt{r}]_{\geq 1}\}$.

Now assume there is a finite-state model A with $n > 0$ states and initial state a_0, such that $A \models^{\mathsf{p}} \varphi$ and $A \preceq M$. Since A is finite-state there exists a state a_1 in A (a successor of a_0) which simulates infinitely many successors $s_{i_j,1}$ ($j > 0$) of s_0 in M. Of these states $s_{i_j,1}$ there must be a state $s_{n_0,1}$ which is starting point of a path $s_{n_0,1} \dots s_{n_0,n_0}$ with $n_0 > n+1$ states labeled $\mathtt{q}$ before reaching its absorbing $\mathtt{r}$ state. By the definition of simulation this path must be matched by a path $a_1 \dots a_{n_0}$ in A such that $a_j \preceq s_{n_0,j}$ for all $1 \leq j \leq n_0$. Since A is of finite size n there must be a state $a_{j'}$ re-occurring along this path, and thus there is a loop in A. As the states $s_{n_0,1} \dots s_{n_0,n_0-1}$ are labeled $\mathtt{q}$, all states of the path $a_1 \dots a_{n_0}$, and on the loop in this path, must be labeled $\mathtt{q}$ or $\mathtt{q}?$. Similarly, as the states $s_{n_0,1} \dots s_{n_0,n_0-1}$ are not labeled $\mathtt{r}$, for all states a_j of the loop we get $L(a_i, \mathtt{r}) = \mathsf{ff}$ or $L(a_i, \mathtt{r}) = \mathtt{?}$. Now, since $A \models^{\mathsf{p}} \varphi$ by assumption, the states a_j must actually be labeled with $\mathtt{q}$. Otherwise, let $\alpha \in \{\mathtt{q}\,\mathsf{U}\,\mathtt{r}, \mathtt{q}\,\mathsf{W}\,\mathtt{r}\}$. If one state a_{i_0} in the loop were labeled $\mathtt{q}?$, and so $A^{a_{i_0}} \not\models^{\mathsf{p}} \mathtt{q}$, there would be a finite prefix $a_0 a_1 \dots a_{i_0}$, and thus a measurable cylinder path set with positive

probability for which no path pessimistically satisfies α. Thus $\mathsf{Prob}^{\mathsf{p}}_M(a_0, \alpha) < 1$, contradicting $A \models^{\mathsf{p}} \varphi$.

But now we have an overall contradiction: no model that contains a loop of states labeled $\mathtt{q}$ can simulate M because this would imply that M contains an infinite path of states labeled $\mathtt{q}$, which the given M clearly does not. Hence there cannot be a finite-state model A such that $A \vDash^{\mathsf{p}} \varphi$ and $A \preceq M$. □

We can use the same concrete Markov chain M from Figure 2(a) and a similar proof structure to show the incompleteness of $[\mathsf{X}\,\phi]_{\geq p}$.

Lemma 3. *Not all formulae of form $[\mathsf{X}\ \phi]_{\geq p}$ are complete.*

Proof. We consider $\varphi = [\mathsf{X}\,[\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{>0}]_{\geq 1}$ and the Markov chain M from Figure 2(a) which satisfies φ. Again, assume there is a finite-state model A with n states and initial state a_0, such that $A \vDash^{\mathsf{p}} \varphi$ and $A \preceq M$.

Since A is finite-state there exists a state a_1 in A (a successor of a_0) which simulates infinitely many successors $s_{i,1}$ of s_0 in M. Since $A \vDash \varphi$ the state a_1 needs to satisfy $[\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{>0}$. Hence there must be a path $\pi = a_1 \ldots a_k$ where the states $a_1, \ldots, a_{k-1}$ are labeled $\mathtt{q}$ and a_k is labeled $\mathtt{r}$. If this path were loop-free, then every of the infinitely many states $s_{i,1}$ would be starting point of a path which reaches an $\mathtt{r}$ state after at most k steps. This is a contradiction to the definition of M. Thus π must contain a loop of states labeled $\mathtt{q}$. But this would force M to contain an infinite path $s_{i,1} \ldots$ where all states are labeled $\mathtt{q}$. Again we have a contradiction because M does not contain such a path. □

Sub-formula $[\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{>0}$ in and of itself does not imply incompleteness. In Section 4, we will actually show that formulae of this form are complete.

Incompleteness of formulae of form $[\mathsf{X}\,[\phi\,\mathsf{U}\,\psi]_{>p}]_{\geq p'}$ requires infinite branching, as in the Markov chain in Figure 2(a). For finitely branching Markov chains this form is complete, as then only a finite number of successor states needs to be considered, on each of which sub-formula $[\phi\,\mathsf{U}\,\psi]_{>p}$ can be finitely verified (as we show in the next section). Forms $[\phi\,\mathsf{U}\,\psi]_{\geq 1}$ and $[\phi\,\mathsf{W}\,\psi]_{\geq 1}$ are also incomplete for finitely branching models (for slightly different Until formulae). We summarize:

Corollary 1. *Full PCTL is incomplete.*

Our incompleteness proofs above work for any simulation notion $\preceq$ satisfying

1. $L(t, \mathtt{q}) \leq L(s, \mathtt{q})$ for all $\mathtt{q} \in \mathtt{AP}$
2. $\mathbf{P}(s, s') > 0$ implies $\mathbf{P}(t, t') > 0$ for some t' with $t' \preceq s'$
3. $\mathbf{P}(t, t') > 0$ implies $\mathbf{P}(s, s') > 0$ for some s' with $t' \preceq s'$

whenever $t \preceq s$. In their bi-directionality, these three conditions are reminiscent of Larsen and Skou's probabilistic $2/3$-bisimulation [3] and of Dams and Namjoshi's notion of (mixed) reverse simulation for labeled transition systems [12]: conditions (1) and (2) together constrain the abstract model in terms of the concrete model (and are necessary but not sufficient for sound abstraction à la Lemma 1); conditions (1) and (3) constrain the concrete model in terms of the abstract one (and are necessary for securing our incompleteness results).

4 Complete fragment of PCTL

We now present a complete fragment of PCTL: those PCTL formulae whose path modalities occur in a suitable combination of negation polarity and threshold type. The technical details of this definition, and its alternative characterization via a normal form will be formalized below. In fact, we will show that for this fragment the desired finite abstractions can be obtained by unfolding the infinite model up to a bounded height and width. We first formalize full unfoldings.

Definition 7 (Unfolding). *Let $M = (S, \mathbf{P}, L)$ be a model. The* full unfolding *of M at s_0 is the model $M^{s_0}_{\text{full}} = (S_{\text{full}}, \mathbf{P}', L')$ where S_{full} is the set of nonempty sequences π over S, transition probability $\mathbf{P}'(s_1 \dots s_n, s_1 \dots s_n s_{n+1})$ is $\mathbf{P}(s_n, s_{n+1})$, and $L'(\pi \cdot s) = L(s)$. We restrict the set S_{full} to the set of sequences reachable from s_0 with positive probability.*

If M is a concrete Markov chain, so is $M^{s_0}_{\text{full}}$. Also, M and $M^{s_0}_{\text{full}}$ simulate each other, and so are equivalent. We now formalize finite unfoldings.

Definition 8 (Finite Unfolding).

1. *For $i \in \mathbb{N}$ and $s_0 \in S$, the* finite unfolding $M^{s_0}_i = (S_i, \mathbf{P}_i, L_i)$ *is the model where S_i is the set of nonempty sequences over S of length at most i, plus a designated sink state t_{sink}. As above $\mathbf{P}_i(s_1 \dots s_n, s_1 \dots s_n s_{n+1}) = \mathbf{P}(s_n, s_{n+1})$, $\mathbf{P}_i(s_0 \dots s_{i-1}, t_{\text{sink}}) = 1$ for each sequence of length i, and $\mathbf{P}_i(t_{\text{sink}}, t_{\text{sink}}) = 1$. Again, $L_i(\pi \cdot s) = L(s)$, and $L(t_{\text{sink}}, \mathtt{q}) = \mathbf{?}$ for all $\mathtt{q} \in \mathtt{AP}$. We restrict S_i to sequences reachable from s_0 with positive probability.*
2. *For $j \in \mathbb{N}$, this model is further restricted to maximal branching degree j as follows. Let $M^{s_0}_{i,j} = (S_{i,j}, \mathbf{P}_{i,j}, L_{i,j})$, where the components of $M^{s_0}_{i,j}$ are as follows. For each $s \in S_i$, let $t_1, t_2, \dots$ be an enumeration of $\{t_k \in S_i \mid \mathbf{P}_i(s, t_k) > 0\}$ such that $\mathbf{P}_i(s, t_k) \geq \mathbf{P}_i(s, t_{k+1})$ for all $k \in \mathbb{N}$. We then define $\mathbf{P}_{i,j}$ by setting $\mathbf{P}_{i,j}(s, t_k) = \mathbf{P}_i(s, t_k)$ for $k \leq j$ and $\mathbf{P}_{i,j}(s, t_{\text{sink}}) = 1 - \Sigma^j_{k=1} \mathbf{P}_i(s, t_k)$. We set $L_{i,j} = L_i$ and again restrict $S_{i,j}$ to sequences reachable from s_0 with positive $\mathbf{P}_{i,j}$ transition probabilities.*

The unfolding $M^{s_0}_{3,3}$ for the labeled Markov chain M of Figure 2(a) is depicted in Figure 2(b). Finite unfoldings give rise to simulations:

Lemma 4. *For all models M with initial state s_0 and $i, j \in \mathbb{N}$, the finite unfolding $M^{s_0}_{i,j}$ simulates M.*

Now we show that Next and Strong Until with $> p$ bounds have 'finite' unfoldings of the model as witnesses.

Lemma 5. *Let M be a model, $\mathtt{q}, \mathtt{r} \in \mathtt{AP}$ be propositions, and $M \vDash^{\mathsf{P}} [\alpha]_{>p}$ for $\alpha \in \{\mathsf{X}\,\mathtt{q}, \mathtt{q}\,\mathsf{U}\,\mathtt{r}\}$. There are i_0, j_0 with $M^{s_0}_{i,j} \vDash^{\mathsf{P}} [\alpha]_{>p}$ for all $i \geq i_0$ and $j \geq j_0$.*

$$\phi_{\text{pos}} ::= \mathtt{q} \mid \neg\mathtt{q} \mid \phi_{\text{pos}} \wedge \phi_{\text{pos}} \mid \phi_{\text{pos}} \vee \phi_{\text{pos}} \mid \neg\phi_{\text{neg}} \mid [\alpha_{\text{pos}}]_{>p} \mid [\alpha_{\text{neg}}]_{<p}$$
$$\phi_{\text{neg}} ::= \mathtt{q} \mid \neg\mathtt{q} \mid \phi_{\text{neg}} \wedge \phi_{\text{neg}} \mid \phi_{\text{neg}} \vee \phi_{\text{neg}} \mid \neg\phi_{\text{pos}} \mid [\alpha_{\text{neg}}]_{\geq p} \mid [\alpha_{\text{pos}}]_{\leq p}$$
$$\alpha_{\text{pos}} ::= \mathsf{X}\,\phi_{\text{pos}} \mid \phi_{\text{pos}}\,\mathsf{U}^{\leq k}\,\phi_{\text{pos}} \qquad \alpha_{\text{neg}} ::= \mathsf{X}\,\phi_{\text{neg}} \mid \phi_{\text{neg}}\,\mathsf{W}^{\leq k}\,\phi_{\text{neg}}$$

Fig. 3. PCTL$_>$, our complete fragment of PCTL, defined as all ϕ_{pos} above where $\mathtt{q} \in \mathtt{AP}$, $k \in \mathbb{N} \cup \{\infty\}$ and $p \in [0,1]$

Proof. Let α be $\mathsf{X}\,\mathtt{q}$. By assumption $M \models^{\mathsf{p}} [\mathsf{X}\,\mathtt{q}]_{>p}$. If s_0 has finitely many successors, the claim is obviously true. Otherwise, let $t_1, t_2, \ldots$ be the successors of s_0 ordered so that $\mathbf{P}(s_0, t_l) \geq \mathbf{P}(s_0, t_{l+1})$ for every $l \geq 1$. Let $t_{m_1}, t_{m_2}, \ldots$ be the sub-sequence of those states t_i with $M^{t_i} \models^{\mathsf{p}} \mathtt{q}$. Then $M \models^{\mathsf{p}} [\mathsf{X}\,\mathtt{q}]_{>p}$ implies $\Sigma_{l=1}^{\infty} \mathbf{P}(s_0, t_{m_l}) > p$. Thus there is some l_0 with $\Sigma_{l=1}^{l_0} \mathbf{P}(s_0, t_{m_l}) > p$. Let $j_0 = m_{l_0}$. For every $i \geq 1$ and $j \geq j_0$ it is then easily seen that $M_{i,j}^{s_0} \models^{\mathsf{p}} [\mathsf{X}\,\mathtt{q}]_{>p}$.

Now let α be $\mathtt{q}\,\mathsf{U}\,\mathtt{r}$. Consider first the case that M is finitely branching. It is simple to see that for all $i \geq 0$ we have $\mathsf{Prob}^{\mathsf{p}}_{M_i^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r}) \leq \mathsf{Prob}^{\mathsf{p}}_{M_{i+1}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r})$ and that $\lim_{i\to\infty} \mathsf{Prob}^{\mathsf{p}}_{M_i}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r}) = \mathsf{Prob}^{\mathsf{p}}_{M}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r})$. Hence, for some i_0 we have that $\mathsf{Prob}^{\mathsf{p}}_{M_{i_0}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r}) > p$ and for every $i \geq i_0$ we have $M_i^{s_0} \models^{\mathsf{p}} [\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{>p}$.

In the case that M has infinite branching the proof is similar. As before, there is some i_0 such that $M_{i_0}^{s_0} \models^{\mathsf{p}} [\mathtt{q}\,\mathsf{U}\,\mathtt{r}]_{>p}$. We notice that for every $j \in \mathbb{N}$ we have $\mathsf{Prob}^{\mathsf{p}}_{M_{i_0,j}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r}) \leq \mathsf{Prob}^{\mathsf{p}}_{M_{i_0,j+1}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r})$ and that $\lim_{j\to\infty} \mathsf{Prob}^{\mathsf{p}}_{M_{i_0,j}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r})$ equals $\mathsf{Prob}^{\mathsf{p}}_{M_{i_0}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r})$. Hence, for some j_0 we have $\mathsf{Prob}^{\mathsf{p}}_{M_{i_0,j_0}^{s_0}}(s_0, \mathtt{q}\,\mathsf{U}\,\mathtt{r}) > p$ and the lemma follows. □

Weak Until and Next with $\geq p$ bounds have finite counter-examples.

Corollary 2. *Let $M \not\models^{\mathsf{o}} [\alpha]_{\geq p}$ for $\alpha \in \{\mathsf{X}\,\mathtt{q}, \mathtt{q}\,\mathsf{W}\,\mathtt{r}\}$ and a model M. Then there exist i_0 and j_0 such that for all $i \geq i_0$ and $j \geq j_0$ we have $M_{i,j}^{s_0} \not\models^{\mathsf{o}} [\alpha]_{\geq p}$.*

Proof. For α being $\mathsf{X}\,\mathtt{q}$ this follows from $[\mathsf{X}\,\varphi]_{>p} \equiv \neg[\mathsf{X}\,\neg\varphi]_{\geq 1-p}$ over two-valued models and from the duality of the optimistic and pessimistic semantics in three-valued models. For α being $\mathtt{q}\,\mathsf{W}\,\mathtt{r}$, we similarly exploit that $[\varphi_1\,\mathsf{W}\,\varphi_2]_{\geq p}$ is equivalent to $\neg[\neg\varphi_2\,\mathsf{U}(\neg\varphi_1 \wedge \neg\varphi_2)]_{>1-p}$ over two-valued models. □

We state and prove our main result, the completeness of PCTL$_>$, which is defined in Figure 3. GTNNF normalforms of PCTL$_>$ allow only $[\mathsf{U}]_{>p}$ and $[\mathsf{X}]_{>p}$ type operators. That is, they disallow Weak Until and the comparison $\geq p$.

Although any finite-state abstraction would be sufficient for completeness we show a stronger result: the abstraction can be chosen as finite unfolding.

Theorem 1 (Completeness of PCTL$_>$). *Let M be a Markov chain with initial state s_0, ϕ a formula in PCTL$_>$, and $M \models \phi$. Then there exist i, j such that the finite unfolding $M_{i,j}^{s_0}$ of M pessimistically satisfies ϕ, i.e. $M_{i,j}^{s_0} \models^{\mathsf{p}} \phi$.*

Proof. We strengthen the claim with a dual claim for formulae in the negative part of $\mathrm{PCTL}_>$ and for the optimistic semantics: "For ϕ in the negative part ϕ_{neg} of $\mathrm{PCTL}_>$, if $M^s \not\models \phi$ then there exist i, j such that $M^s_{i,j} \not\models^{\mathsf{o}} \phi$." We show this extended claim by structural induction on ϕ, simultaneously for all states s.

- Let ϕ be q. If $M^s \models \mathsf{q}$ then for every $i \geq 0$ and $j \geq 0$ we have $M^s_{i,j} \models^{\mathsf{p}} \mathsf{q}$. Dually, if $M^s \not\models \mathsf{q}$ then for every $i \geq 0$ and $j \geq 0$ we have $M^s_{i,j} \not\models^{\mathsf{o}} \mathsf{q}$.
- For the Boolean connectives $\phi_1 \wedge \phi_2$ and $\phi_1 \vee \phi_2$ and a state s, we take as bounds the maximum of the bounds i_k and j_k for sub-formulae ϕ_k obtained by induction for state s. These bounds work for the dual case as well.
- For a negation $\varphi = \neg\psi_{\mathrm{neg}}$ and a state s, if $M^s \models \neg\psi_{\mathrm{neg}}$, then $M^s \not\models \psi_{\mathrm{neg}}$. By induction, there are i and j with $M^s_{i,j} \not\models^{\mathsf{o}} \psi_{\mathrm{neg}}$. Thus $M^s_{i,j} \models^{\mathsf{p}} \neg\psi_{\mathrm{neg}}$.
  Dually, for a negation $\varphi = \neg\psi_{\mathrm{pos}}$ and a state s, if $M^s \not\models \neg\psi_{\mathrm{pos}}$, then $M^s \models \psi_{\mathrm{pos}}$. By induction, there are i and j with $M^s_{i,j} \models^{\mathsf{p}} \psi_{\mathrm{pos}}$, so $M^s_{i,j} \not\models^{\mathsf{o}} \neg\psi_{\mathrm{pos}}$.
- We now consider the path modalities X, U, and W.
  - For formula $\varphi = [\mathsf{X}\,\psi_{\mathrm{pos}}]_{>p}$ and a state s such that $M^s \models \varphi$, we treat ψ_{pos} as a proposition that labels the states of M. By (the proof of) Lemma 5, there is some j'_0 such that for every $i \geq 1$ and $j \geq j'_0$ we have $M^s_{i,j} \models^{\mathsf{p}} \varphi$. Let $t_1, \ldots, t_{j_0}$ be the first j'_0 successors of s. For t_k there exists i^k_0 and j^k_0 such that if $M^{t_k} \models \psi_{\mathrm{pos}}$ we have $M^{t_k}_{i^k_0, j^k_0} \models^{\mathsf{p}} \psi_{\mathrm{pos}}$. Let $i_0 = 1 + \max_k(i^k_0)$ and $j_0 = \max(j'_0, \max_k(j^k_0))$. It follows that $M^s_{i_0,j_0} \models^{\mathsf{p}} \varphi$.
  - Let $\varphi = [\mathsf{X}\,\psi_{\mathrm{neg}}]_{\geq p}$ with $M^s \not\models \varphi$. The proof is similar to the one in the previous item and uses Corollary 2.
  - For $\varphi = [\psi_1\,\mathsf{U}\,\psi_2]_{>p}$, with ψ_1 and ψ_2 in the positive fragment ϕ_{pos}, and a state s with $M^s \models [\psi_1\,\mathsf{U}\,\psi_2]_{>p}$, we initially treat ψ_1 and ψ_2 as propositions that label the states of M. By Lemma 5 there are i'_0 and j'_0 such that for every $i \geq i'_0$ and $j \geq j'_0$ we have $M^s_{i,j} \models^{\mathsf{p}} [\psi_1\,\mathsf{U}\,\psi_2]_{>p}$. Now we no longer treat the ψ_i as atoms: Let $t_1, \ldots, t_m$ be all the states appearing in $M_{i'_0,j'_0}$. For $\alpha \in \{1, 2\}$ and every t_k there exists $i^{k,\alpha}_0$ and $j^{k,\alpha}_0$ such that if $M^{t_k} \models \psi_\alpha$ we have $M^{t_k}_{i^{k,\alpha}_0, j^{k,1}_0} \models^{\mathsf{p}} \psi_\alpha$. Let $i_0 = i'_0 + \max_{k,\alpha}(i^{k,\alpha}_0)$ and $j_0 = \max(j'_0, \max_{k,\alpha}(j^{k,\alpha}_0))$ (see Figure 4). It follows that $M^s_{i_0,j_0} \models^{\mathsf{p}} \varphi$.
  - The proof for $\varphi = [\psi_1\,\mathsf{W}\,\psi_2]_{\geq p}$, with ψ_1 and ψ_2 in the fragment ϕ_{neg}, and a state s such that $M^s \not\models [\psi_1\,\mathsf{U}\,\psi_2]_{\geq p}$ is similar to the one in the previous item and uses Corollary 2.
  - Formula $[\alpha_{\mathrm{neg}}]_{<p}$ is equivalent to $\neg[\alpha_{\mathrm{neg}}]_{\geq 1-p}$ of form $\neg\varphi_{\mathrm{pos}}$. Formula $[\alpha_{\mathrm{pos}}]_{\leq p}$ is equivalent to $\neg[\alpha_{\mathrm{pos}}]_{>1-p}$ of form $\neg\varphi_{\mathrm{neg}}$. Thus this case follows by induction. For example, for state s, we have e.g. $M^s_{i_0,j_0} \models^{\mathsf{p}} [\alpha_{\mathrm{neg}}]_{<p}$ iff $M^s_{i_0,j_0} \models^{\mathsf{p}} \neg[\alpha_{\mathrm{neg}}]_{\geq 1-p}$ iff $M^s_{i_0,j_0} \not\models^{\mathsf{o}} [\alpha_{\mathrm{neg}}]_{\geq 1-p}$. □

We now show that the results in Section 3 imply that PCTL fragments that allow combinations of the operators we disallow cannot be complete. To that end, we first prove an additional incompleteness result.

Lemma 6. *Not all formulae of form* $[[\phi\,\mathsf{U}\,\psi]_{>p}\,\mathsf{W}\,\rho]_{>p'}$ *are complete.*

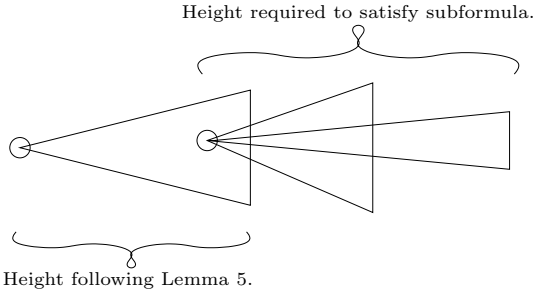

Fig. 4. Intuitively an unfolding for a sub-formula can be attached to every inner state of the unfolding of the formula. The resulting maximal height is still finite

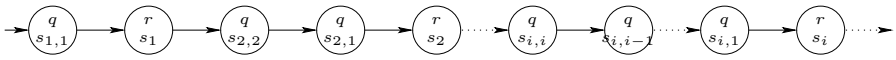

Fig. 5. Concrete Markov chain M that satisfies $[[\mathsf{q}\,\mathsf{U}\,\mathsf{r}]_{>0}\,\mathsf{W}\,\mathsf{ff}]_{>0}$

Proof. Let φ be $[[\mathsf{q}\,\mathsf{U}\,\mathsf{r}]_{>0}\,\mathsf{W}\,\mathsf{ff}]_{>0}$ and M be as in Figure 5. It is simple to see that $M \models \varphi$. Suppose there is a finite-state model A such that $A \preceq M$ and $A \models^{\mathsf{P}} \varphi$. Let a be the initial state of A such that $a \preceq s_0$. As $A \models^{\mathsf{P}} \varphi$, there is a bottom strongly connected component (SCC) in A such that every state in this SCC satisfies pessimistically $[\mathsf{q}\,\mathsf{U}\,\mathsf{r}]_{>0}$. By a pigeon-hole principle, we can find a state a' in this SCC that is labeled by r and simulates states s_i for infinitely many i. Consider a cycle from a' to itself. This cycle has some fixed length n. As for every $i > 0$ the distance from s_i to s_{i+1} is $i + 1$, this is a contradiction. □

We can now prove that static extensions of $\mathrm{PCTL}_>$ are incomplete.

Theorem 2. *Consider a PCTL fragment κ that contains one of the following combinations of PCTL operators: (i) $[\phi\,\mathsf{W}\,\psi]_{\geq p}$, (ii) $[\phi\,\mathsf{U}\,\psi]_{\geq p}$, (iii) $[\mathsf{X}\,\phi]_{\geq p}$ and $[\phi\,\mathsf{U}\,\psi]_{>p}$, or (iv) $[\phi\,\mathsf{W}\,\psi]_{>p}$ and $[\phi\,\mathsf{U}\,\psi]_{>p}$. Then κ is incomplete.*

Proof. The first three items follow from Lemmas 2 and 3 in Section 3. The last item follows from Lemma 6 above. □

5 Discussion and Conclusions

From a practical perspective, our completeness results mean that finite-state, 3-valued Markov chains are complete as abstractions for all of PCTL as long as Strong Untils occur under positive polarity and Weak Untils under negative polarity: Given such a formula, we can determine all its occurrences of path modalities whose negation polarity and threshold type do not match. Then we can change all such threshold types and adjust their probability with a small perturbation in situ. For example, a Weak Until under negative polarity with

$> .99$ threshold could be made complete by making it a Weak Until with $\geq .99 + 10^{-12}$ threshold without compromising the original intent of that property.

Let us conclude. We investigated whether the truth of formulae in probabilistic computation tree logic over infinite-state Markov chains can, in principle, be witnessed by finite-state Markov chains that simulate such infinite-state models of formulae and allow for 3-valued interpretations of atomic propositions. Negative results were presented for certain combinations of path modalities and probability threshold type, e.g. for Weak Until with strict threshold type. Positive results were proved for a sizeable fragment of PCTL formulae whose path modalities all occur in a statically determined combination of negation polarity and threshold type. Finally, we showed that static extensions of that complete fragment of PCTL are incomplete.

Acknowledgments. Mark Kattenbelt discussed with us incompleteness of stochastic games for MDPs. This research was supported by UK EPSRC project *Complete and Efficient Checks for Branching-Time Abstractions (EP/E028985/1)* and the Computing Laboratory Oxford University, which hosted the first author's sabbatical leave.

References

1. Hansson, H., Jonsson, B.: A logic for reasoning about time and reliability. Formal Aspects of Computing **6** (1994) 512–535
2. Kwiatkowska, M.: Model checking for probability and time: From theory to practice. In: Proc. of LICS 2003, IEEE Computer Society.
3. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Information and Computation **94** (1991) 1–28
4. Jonsson, B., Larsen, K.: Specification and refinement of probabilistic processes. In: Proc. of LICS 1991, IEEE Computer Society.
5. Fecher, H., Leucker, M., Wolf, V.: Don't know in probabilistic systems. In: SPIN Workshop on Model Checking of Software. (2006)
6. Fecher, H., Huth, M., Piterman, N., Wagner, D.: Hintikka Games for PCTL on Labeled Markov Chains. In: Proc. of QEST 2008, IEEE Computer Society.
7. Kattenbelt, M., Huth, M.: Abstraction Framework for Markov Decision Processes and PCTL Via Games. Technical Report RR-09-01, Oxford University Computing Laboratory (2009)
8. Kemeny, J., Snell, J., Knapp, A.: Denumerable Markov Chains, Springer (1976).
9. Kleene, S.: Introduction to Metamathematics. Van Nostrand (1952)
10. Godefroid, P., Huth, M., Jagadeesan, R.: Abstraction-based model checking using modal transition systems. In: Proc. of CONCUR 2001, LNCS 2154, Springer.
11. Huth, M.: On finite-state approximants for probabilistic computation tree logic. Theoretical Computer Science **346** (2005) 113–134
12. Dams, D., Namjoshi, K.: The existence of finite abstractions for branching time model checking. In: Proc. of LICS 2004, IEEE Computer Society.
13. Bianco, A., de Alfaro, L.: Model checking of probabilistic and nondeterministic systems. In: Proceedings of the 15th Conference on Foundations of Software Technology and Theoretical Computer Science, Springer-Verlag (1995) 499–513

A Ancillary Material

Proof (of Lemma 4). We show that the relation

$$H = \{(s, \pi \cdot s) \mid s \in S, \pi \in S^*\} \cup \{(s, t_{\text{sink}}) \mid s \in S\}$$

is a simulation relation.

First, for all $\pi \in S^*$, $s \in S$ and $\mathtt{q} \in \mathtt{AP}$ we have $L(s, \mathtt{q}) = L(\pi \cdot s, \mathtt{q})$. Since $L(t_{\text{sink}}, \mathtt{q}) = \mathtt{?}$ for all $\mathtt{q} \in \mathtt{AP}$ we clearly have $L(t_{\text{sink}}, \mathtt{q}) \leq L(s, \mathtt{q})$ for all $s \in S$.

We now define the weight function $\rho_s \colon S \times S \longrightarrow [0, 1]$:

- $\rho_s(s', t_{\text{sink}}) = 1$ if there is no path π such that $\pi \cdot s'$ is in $M_{i,j}^{s_0}$ (which equivalently means $|s_0 \dots s'| > i$ for every such path in M or $s' = t_{k'}$ with $k' > j$ for the ordering t_k of the successor states of a state in $M_i^{s_0}$ as described in Definition 8).
- $\rho_s(s', t_{\text{sink}}) = 0$ for all other $s' \in S$.
- $\rho_s(s', \pi) = 1$ if $\pi = \pi' \cdot s \cdot s'$.
- $\rho_s(s', \pi) = 0$ for all other paths.

Then the condition $\sum_{s' \in S}(\mathbf{P}(s, s') \cdot \rho_s(s', \pi')) = \mathbf{P}(\pi, \pi')$ for all $\pi' \in M_{i,j}^{s_0}$ collapses to $\mathbf{P}(s, s') = \mathbf{P}(\pi, \pi')$ for $\pi = \pi'' \cdot s$ and $\pi' = \pi'' \cdot s \cdot s'$ and $\sum_{s' \in S} \mathbf{P}(s, s') = 1 = \mathbf{P}(\pi, t_{\text{sink}})$ for all other states. Both equations are obviously true by the construction of $M_{i,j}^{s_0}$.

Finally, we need to check the co-inductive condition for simulation: Whenever $\rho_s(s', t_{\text{sink}}) > 0$ we have $(s', t_{\text{sink}}) \in H$; whenever $\rho_s(s', \pi') > 0$, then $\pi' = \pi'' \cdot s \cdot s'$ and hence $(s', \pi') = (s', \pi'' \cdot s \cdot s') \in H$. □