The Satisfiability Problem for Probabilistic CTL

Philipp Berger

November 17, 2014

Abstract

Probabilistic CTL formulas are used in model checking, which is a technique for formal verification with successful utilization in the analysis of systems from a diverse background like randomized algorithms, biological processes and communication protocols. In this paper we analyze a subclass of PCTL, the qualitative fragment, where all probabilistic bounds are either $= 0$, > 0, < 1 or $= 1$. We present a method for generating a satisfying model for a given formula in exponential time. We present a technique for representing infinite–state models in finite structures and therefore present an algorithm which finds a finite description of a satisfying model for every satisfiable qualitative PCTL formula.

1 Introduction

1.1 Results for Computation Tree Logic (CTL)

Computation Tree Logic (CTL) is a branching time logic which allows the specification of properties over discrete time systems. For this purpose, we distinguish state- and path formulas. State formulas allow assertions on the branching structure and atomic propositions exhibited by a single state, whereas path formulas allow assertions on future states. The Syntax of CTL state formulas is as follows:

$$\Phi := \quad \text{true} \quad | \quad a \quad | \quad \Phi_1 \wedge \Phi_2 \quad | \quad \neg\Phi \quad | \quad \exists\varphi \quad | \quad \forall\varphi$$

where $a \in AP$ is an atomic proposition and φ is a CTL path formula. CTL path formulas follow the syntax

$$\varphi := \quad \bigcirc\Phi \quad | \quad \Phi_1 \mathbf{U} \Phi_2,$$

where Φ, Φ_1 and Φ_2 are state formulas.

The satisfiability problem is defined as the question whether a given formula can be satisfied, e.g. if there exists a (possibly infinite) model on which the formula holds. If a logic possesses the small model property, than for every formula in the logic there exists a model that satisfies the formula and this model is finite. In Emerson and Halpern [2] it is shown that for CTL the satisfiability problem is EXPTIME–complete and that CTL possesses the small model property.

The algorithm proposed by Emerson and Halpern [2] focuses on creating a pseudo–model from the formula by applying a technique due to Fischer and Ladner involving the closure of a formula and its subsets. They argue that this technique is a generalized approach that is promising for other logics as well.

1.2 PCTL and the Qualitative Fragment

Brázdil et al. [1] build upon this approach and apply it to Probabilistic CTL (PCTL), a logic similar to CTL but extended by probabilistic operators. As in CTL the syntax of PCTL is divided into state formulas and path formulas. The syntax of PCTL state formulas is

$$\Phi := \quad \text{true} \quad | \quad a \quad | \quad \Phi_1 \wedge \Phi_2 \quad | \quad \neg\Phi \quad | \quad \mathbb{P}_J(\varphi),$$

where, again, $a \in AP$ is an atomic proposition and φ is a CTL path formula. J is of the form $\bowtie p$ with $\bowtie \in \{<, >, \leqslant, \geqslant\}$ and $p \in [0, 1]$.

PCTL path formulas follow the syntax

$$\varphi := \quad \bigcirc\Phi \quad | \quad \Phi_1 \mathbf{U} \Phi_2 \quad | \quad \Phi_1 \mathbf{U}^{\leqslant n} \Phi_2,$$

where Φ, Φ_1 and Φ_2 are state formulas and $n \in \mathbb{N}$ is an upper bound on the number of steps.

The qualitative fragment of PCTL is identical to PCTL, but the probability operator $\mathbb{P}_J(\varphi)$ is restricted to $J \in \{= 0, > 0, < 1, = 1\}$.

2 The Qualitative PCTL Satisfiability Problem

In this chapter we give an algorithm and accompanying descriptional structures for representing infinite–state models for the satisfiability problem on the qualitative fragment of PCTL.

We will prove that for every satisfiable qualitative PCTL formula there exists a finite representation of its possibly infinite model. Along with this proof we will introduce pseudo–models and marked graphs.

The basic approach of this technique will be:

1. Deduce a model–like structure from the initial formula.
2. Iterate certain logical checks for consistency, prune parts that do not fit.
3. On convergence: check whether the initial formula is satisfied in a state.
4. If yes: return abstract model description, or
5. if no: return unsatisfiable.

A common way for solving the satisfiability problem on CTL formulas uses the Fischer–Ladner closure of the given formula as a basis for constructing a model. This closure of a CTL formula Φ contains all sub–formulas, negations, etc. of Φ. Of course

many formulas in this set contradict each other, and therefore the set is used to create a set of all eligible subsets of the closure. A set of formulas is eligible iff it is logically consistent with regard to all contained formulas.

We employ a similar construction for PCTL but with some important changes. A subtle difference between CTL and PCTL is the behavior on the existence of a single path with a specific trait. In CTL a single path is enough to refute a property like an until formula $\forall\Phi_1\mathbf{U}\Phi_2$. If there exists a single path on which the formula does not hold, than this path is a model for the formula $\exists\neg(\Phi_1\mathbf{U}\Phi_2)$ and is a counterexample for the universal formula. In PCTL logic one could say an equivalent formula is $\mathbb{P}_{=1}(\Phi_1\mathbf{U}\Phi_2)$. Because the probability of a single infinite path can be zero, this formula might still hold, even if there exists a path on which the property does not hold. To refute this formula, one of two cases must hold. Following the idea of the aforementioned path, a single path can be enough if it is a finite path. If there exists a finite path on which the until formula does not hold, then this path has a non–zero probability and therefore refutes the property. If there are only infinite paths and the model is finite then by standard probability theory this path must lead to a bottom strongly connected component (BSCC). As the probability of not visiting a state in this BSCC is zero, all states along the way leading to this BSCC and all states in the BSCC support $\Phi_1 \wedge \neg\Phi_2$.

Lemma 1. *Let $\Phi := \mathbb{P}_{=1}(\Phi_1 \boldsymbol{U}\Phi_2)$ be a qualitative PCTL formula, $\mathcal{M} = (S, \rightarrow, L)$ be a Markov chain with a state–set S, a transition function $\rightarrow$ and a labeling function $L : S \rightarrow 2^{AP}$. If there exists a state $s \in S$ in which Φ does not hold, e.g. $\mathcal{M}_s \not\models \mathbb{P}_{=1}(\Phi_1 \boldsymbol{U}\Phi_2)$, then either*

1. *there is a finite path $\pi = s_0 \dots s_n$ starting in s such that $s_0 = s$ and $\mathcal{M}_{s_i} \models \Phi_1 \wedge \neg\Phi_2$ for $i < n$ and $\mathcal{M}_{s_n} \models \neg\Phi_1 \wedge \neg\Phi_2$, or*
2. *$\Pr^{\mathcal{M}}(\{\pi | \forall s_i \in \pi : \mathcal{M}_{s_i} \models \Phi_1 \wedge \neg\Phi_2\}) > 0$. For a finite model $\mathcal{M}$, each infinite path eventually leads to a BSCC α. For all states s on the path and in α it holds that $\mathcal{M}_s \models \Phi_1 \wedge \neg\Phi_2$.*

2.1 Pseudo–Structures and -Models

We use a version of the Fischer–Ladner closure adapted to PCTL. Given a PCTL formula Φ, the closure of Φ, $Cl(\Phi)$, is the least set satisfying the following rules:

1. $\Phi_1 \in Cl(\Phi) \Rightarrow \neg\Phi_1 \in Cl(\Phi)$
2. $\Phi_1 \wedge \Phi_2 \in Cl(\Phi) \Rightarrow \Phi_1 \in Cl(\Phi)$ and $\Phi_2 \in Cl(\Phi)$
3. $\mathbb{P}_{\bowtie}(\bigcirc\Phi_1) \Rightarrow \Phi_1 \in Cl(\Phi)$
4. $\mathbb{P}_{\bowtie}(\Phi_1\mathbf{U}\Phi_2) \in Cl(\Phi) \Rightarrow \Phi_1 \in Cl(\Phi)$ and $\Phi_2 \in Cl(\Phi)$ and $\mathbb{P}_{\bowtie}(\bigcirc\mathbb{P}_{\bowtie}(\Phi_1\mathbf{U}\Phi_2)) \in Cl(\Phi)$
5. $\mathbb{P}_{=1}(\Phi_1\mathbf{U}\Phi_2) \in Cl(\Phi) \Rightarrow \mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2) \in Cl(\Phi)$

6. $\Phi \in Cl(\Phi)$

In a simple example we calculate the closure of the formula $\Phi := \mathbb{P}_{=1}(\neg a \mathbf{U} a)$.

1. $\mathbb{P}_{=1}(\neg a \mathbf{U} a)$ (Φ itself)
2. $\neg a$ (rule 4)
3. a (rule 4)
4. $\mathbb{P}_{=1}(\bigcirc \mathbb{P}_{=1}(\neg a \mathbf{U} a))$ (rule 4)
5. $\mathbb{P}_{>0}(\neg a \mathbf{U} a)$ (rule 5)
6. $\mathbb{P}_{>0}(\bigcirc \mathbb{P}_{>0}(\neg a \mathbf{U} a))$ (rule 4)

Of course, the negated version of entries $1, 4, 5$ and 6 are also present. The full set is:

$$\begin{aligned} Cl(\Phi) :=\{&a, \neg a, \Phi, \neg\Phi, \mathbb{P}_{=1}(\bigcirc \mathbb{P}_{=1}(\neg a \mathbf{U} a)), \neg \mathbb{P}_{=1}(\bigcirc \mathbb{P}_{=1}(\neg a \mathbf{U} a)), \\ &\mathbb{P}_{>0}(\neg a \mathbf{U} a), \neg \mathbb{P}_{>0}(\neg a \mathbf{U} a), \mathbb{P}_{>0}(\bigcirc \mathbb{P}_{>0}(\neg a \mathbf{U} a)), \\ &\neg \mathbb{P}_{>0}(\bigcirc \mathbb{P}_{>0}(\neg a \mathbf{U} a))\} \end{aligned}$$

As we want to create a state–set for a model that satisfies Φ, we need to embed the formula into the states. In each state, some formulas hold and some do not — the powerset of the closure is closely affiliated with this state–set. But we restrict the subsets to those which are logically consistent. Here we implicitly use the expansion law of the until formula known from model checking theory to capture the requirements for the formula to hold in implications for next–states, e.g. either the until formula is fulfilled in this state or it could hold as the current state satisfies the intermediate condition. We call a subset S of the closure of Φ consistent or eligible iff for every $\Psi \in Cl(\Phi)$ it holds that:

- $\Psi \in S \Leftrightarrow \neg\Psi \notin S$
- $\Psi_1 \wedge \Psi_2 \in S \Rightarrow \Psi_1 \in S \wedge \Psi_2 \in S$
- $\neg(\Psi_1 \wedge \Psi_2) \in S \Rightarrow \neg\Psi_1 \in S \vee \neg\Psi_2 \in S$
- $\mathbb{P}_{>0}(\Psi_1 \mathbf{U} \Psi_2) \in S \Rightarrow \Psi_2 \in S$ or ($\Psi_1 \in S$ and $\mathbb{P}_{>0}(\bigcirc \mathbb{P}_{>0}(\Psi_1 \mathbf{U} \Psi_2)) \in S$)
- $\neg(\mathbb{P}_{>0}(\Psi_1 \mathbf{U} \Psi_2)) \in S \Rightarrow (\neg\Psi_1 \in S$ and $\neg\Psi_2 \in S)$ or $(\neg\Psi_2 \in S \wedge \neg \mathbb{P}_{>0}(\bigcirc \mathbb{P}_{>0}(\Psi_1 \mathbf{U} \Psi_2)) \in S)$
- $\mathbb{P}_{=1}(\Psi_1 \mathbf{U} \Psi_2) \in S \Rightarrow \Psi_2 \in S$ or ($\Psi_1 \in S$ and $\mathbb{P}_{=1}(\bigcirc \mathbb{P}_{=1}(\Psi_1 \mathbf{U} \Psi_2)) \in S$)
- $\neg(\mathbb{P}_{=1}(\Psi_1 \mathbf{U} \Psi_2)) \in S \Rightarrow (\neg\Psi_1 \in S$ and $\neg\Psi_2 \in S)$ or $(\neg\Psi_2 \in S \wedge \neg \mathbb{P}_{=1}(\bigcirc \mathbb{P}_{=1}(\Psi_1 \mathbf{U} \Psi_2)) \in S)$

Based on these possible states, which are now abstractly defined by the formulas they satisfy, we construct a pseudo–structure $\mathcal{A} = (A, \rightarrow)$. The state–set A consists of all eligible subsets of the closure of Φ and $\rightarrow \subseteq A \times A$ is a total relation. The next step towards a suitable model is to add a probability distribution to the pseudo–structure. We use $\mathbf{P}$, a uniform probability distribution for each state. Combining $\mathcal{A}$ and $\mathbf{P}$ yields a Markov chain $\mathcal{M}$. We use the closure of Φ to implicitly define the atomic propositions on $\mathcal{M}$, where for each formula in the closure a new atomic proposition is inserted. This allows us to label each state with those formulas holding in the state, mapping the eligible subsets of $\mathcal{A}$ exactly onto the labeled states of $\mathcal{M}$.

Following the example from earlier, we construct all eligible subsets of $Cl(\Phi)$. In the table below, each numbered column represents a state and each row contains either a 0 (this formula is not contained in the state) or a 1 (this formula is contained in the state).

State	1	2	3	4	5	6
a	1	1	1	0	0	0
$\neg a$	0	0	0	1	1	1
$\mathbb{P}_{>0}(\neg a \mathbf{U} a)$	1	1	1	1	1	0
$\neg\mathbb{P}_{>0}(\neg a \mathbf{U} a)$	0	0	0	0	0	1
$\mathbb{P}_{=1}(\neg a \mathbf{U} a)$	1	1	1	1	0	0
$\neg\mathbb{P}_{=1}(\neg a \mathbf{U} a)$	0	0	0	0	1	1
$\mathbb{P}_{=1}(\bigcirc\mathbb{P}_{=1}(\neg a \mathbf{U} a))$	1	0	0	1	0	0
$\neg\mathbb{P}_{=1}(\bigcirc\mathbb{P}_{=1}(\neg a \mathbf{U} a))$	0	1	1	0	1	1
$\mathbb{P}_{>0}(\bigcirc\mathbb{P}_{>0}(\neg a \mathbf{U} a))$	1	1	0	1	1	0
$\neg\mathbb{P}_{>0}(\bigcirc\mathbb{P}_{>0}(\neg a \mathbf{U} a))$	0	0	1	0	0	1

This gives us a Markov chain with six states.

As discussed earlier, the validity or invalidity of a negated almost–sure until formula is non–trivial to see, but all other formulas can be evaluated on a state and its successors. To formalize the observation from Lemma 1, we define corresponding sub–pseudo–structures of $\mathcal{A}$. A witness for a formula $\Psi := \neg(\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)) \in Cl(\Phi)$ in $\mathcal{A}$ is a pseudo–structure $\mathcal{B} := (B, \hookrightarrow)$ with $\varnothing \neq B \subseteq A$ and $\hookrightarrow \subseteq \rightarrow$. In Lemma 1 we stated that in some cases we require a BSCC in $\mathcal{M}$ to show that Ψ holds in a state. The witness $\mathcal{B}$ models this BSCC, therefore we require, that

- $\mathcal{B}$ is strongly connected,
- for every $s \in B$ it holds that $\neg\Phi_2 \in s$ and
- for every $s \in B$ and every $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \in s$ it holds that $\mathcal{M}^{\mathcal{B}}_s \models \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$.

The main idea of a witness is the following: Imagine a finite Markov chain with a state s. This state s does not satisfy $\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. As this is a finite Markov chain, only a finite path starting in s may fulfill this. To cover infinite models, we search for ways to enhance this. If the probabilities were dropped, an infinite path using a loop reachable from s would also be possible, but because of the probabilities this path has a probability mass

of zero and can therefore not be used. The witness captures this loop and will later on be used for augmentation of the finite model into an infinite model, where the probabilities around the states in the loop are adapted to not be of zero weight.

As the transition relation in a pseudo–structure is quite broad, the formulas on the underlying Markov chain do not necessarily hold. In a next step we therefore refine pseudo–structures by adding two rules about logical consistency. We call the resulting structure a pseudo–model. The first rule requires that for all formulas but formulas Ψ of the form $\neg(\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2))$, e.g. $\mathbb{P}_{=1}(\bigcirc\Phi)$, $\neg\mathbb{P}_{=1}(\bigcirc\Phi)$, $\mathbb{P}_{>0}(\bigcirc\Phi)$, $\neg\mathbb{P}_{>0}(\bigcirc\Phi)$, $\mathbb{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2)$, $\neg(\mathbb{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2))$, and $\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$, that are part of the eligible subset associated with the state s it holds that $\mathcal{M}_s \models \Psi$. This can be evaluated by looking at a state and its direct successors. The second rule takes care of the more involved check for the validity of $\neg(\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2))$. For a pseudo–structure to be a pseudo–model, we require that for $\neg(\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)) \in s, s \in A$ either

1. $\mathcal{M}_s \not\models \mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$ or
2. there exists a finite path $\pi = s_0 ... s_n$ and a witness $\mathcal{B}$ for $\neg(\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2))$, such that $s_0 \in s,\ s_n \in B$ and $\neg\Phi_2 \in s_i$, for $0 \leqslant i \leqslant n$.

We call a pseudo–model simple, iff the first rule always implies the second and no witness is necessary.

2.2 Marked Graphs

For some qualitative PCTL formulas, only infinite models exist. Using marked graphs we present a way of representing such infinite state models in a finite way. To achieve this, we exploit certain properties, i.e. similarities in the structure of the models. A marked graph $\mathcal{G} := (G, \hookrightarrow, L)$, where G is a finite node set, $\hookrightarrow \subseteq G \times G$ a total relation, $L \subseteq \hookrightarrow$ a set of marked transitions, with an induced Markov chain $\mathcal{M}_\mathcal{G} = (G^+, \rightarrow, \mathbf{P})$ is a graph, where the transition relation $\rightarrow$ is induced by $\hookrightarrow$ such that for $w \in G^*$, $x \in G$, there is a transition $wx \rightarrow wxy$ iff there exists some $y \in G$ with $x \hookrightarrow y$. A transition in $\mathcal{M}_\mathcal{G}$ is called marked iff the corresponding transition of $\mathcal{G}$ is marked. The transition probabilities of $\mathcal{M}_\mathcal{G}$ are determined as follows. If all outgoing transitions of a state w are unmarked or all outgoing transitions of w are marked, then the probability distribution in this state will be uniform. Otherwise, they depend on the depth of the state in the graph, i.e. $len(w)$. All marked transitions receive a uniform amount of the share $1 - \left(\frac{1}{4}\right)^{len(w)}$, the remaining part is distributed uniformly over all unmarked transitions. If there are states in a marked graph which have marked and unmarked transitions, the resulting Markov chain has an infinite state set, otherwise it is finite.

2.3 Proof and Algorithm

The algorithm starts by computing the closure of the input formula Φ. The initial state set is the set of eligible subsets of said closure and the transition relation is initialized to be complete.

In this configuration the model is of course invalid, as for any non–trivial formula there will be logical inconsistencies. Imagine formulas of the form $\mathbf{P}_{=1}(\bigcirc\Phi_1)$, which may hold in a state s_1. Per construction that uses the closure of Φ, there will also be states in which Φ_1 does not hold, call such a state s_2. Since the transition relation is initially complete, there will be a transition $s_1 \to s_2$. Of course this transition implies that the formula does not hold in s_1.

This lays out the basic idea of the algorithm. Starting from the complete pseudo–model, we iteratively delete states and transitions that contain formulas which are not satisfied in the current model. As mentioned earlier, all formulas but the negated until formulas can be handled by examining the states and their successors. The until formulas of the form $\neg\,\mathbf{P}_{=1}(\Phi_1\mathbf{U}\Phi_2)$ will be handled differently, depending on the type of model required. For a finite model, we check whether there exists a finite path refuting the until formula. If no such path exists, the state will be removed. For a (possibly) infinite model, before removing the state we further check if a witness for the formula can be employed. Should there be no such witness, the state is finally removed. Whenever we encounter a state with no outgoing edges, we remove it. The algorithm terminates on convergence, i.e. when the state set and transition relation do not change between iterations. If the original formula Φ is still satisfied in a state, we found a model for Φ and can therefore conclude that Φ is indeed satisfiable. Otherwise, we return that Φ is unsatisfiable.

Theorem 2. *If the algorithm returns some $\mathcal{A} = (A, \to)$, then $\mathcal{A}$ is a (simple) pseudo–model for Φ.*

The proof for Theorem 2 is done by showing that the rules of the definition of a pseudo–model hold for the output of the algorithm, which boils down to showing that for each formula ξ in each state $s \in A$ it holds that $\mathcal{M}_s^{\mathcal{A}} \models \xi$. First up, we prove the first rule concerning next and some until formulas:

- For each formula $\xi = \mathbf{P}_{=1}(\bigcirc\Phi_1)$ or $\xi = \neg\,\mathbf{P}_{>0}(\bigcirc\Phi_1)$, $\mathcal{M}_s^{\mathcal{A}} \models \xi$ is ensured by construction. Note that a single transition is enough to invalidate these formulas — exactly those transitions are removed in the algorithms initial steps by looking at exactly these formulas and inspecting all successor states. All transitions which would now invalidate ξ have been removed.

- Formulas of the form $\xi = \mathbf{P}_{>0}(\bigcirc\Phi_1)$, $\xi = \neg\,\mathbf{P}_{=1}(\bigcirc\Phi_1)$ and $\xi = \mathbf{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$ are a bit different, since here a single path is not enough to refute them, but a single path can be enough for them to hold. In the algorithm we check in each iteration if for each ξ such a path still exists. If it has been removed or did never exist, the state containing ξ is also removed.

- Formulas of the form $\xi = \neg\,\mathbf{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$. Proof by contraposition - we assume that $\mathcal{M}_s^{\mathcal{A}} \models \mathbf{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$. As used in the previous part, if the formula holds there exists a finite path $\pi = s_0, ..., s_n$ such that $\Phi_2 \in s_n$ and for $0 \leqslant i \leqslant n : \Phi_1 \in s_i$. We now show by induction over the path length that if such a path exists, a state in the path violates the eligibility criterion.

Induction base ($n = 0$): For a path of length zero the state s must fulfill the until formula. Therefore $\Phi_2 \in s$. But since $\neg \mathbf{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2) = \xi \in s$ we know that by completeness $\neg\Phi_2 \in s$. Since $\Phi_2 \in s$ and $\neg\Phi_2 \in s$, the state s is not eligible.

Induction step ($n \to n+1$): There are two possibilities for $\neg \mathbf{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2) \in s = s_n$ — either both $\neg\Phi_1 \in s_n$ and $\Phi_2 \in s_n$ — or $\neg \mathbf{P}_{>0}(\bigcirc \mathbf{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2)) \in s_n$. In the first case, with $\neg\Phi_1 \in s_n$ and $\Phi_2 \in s_n$, we know that per definition of the finite path we have that $\Phi_1 \in s_n$, too. Then s_n is not eligible as Φ_1 is in the state both negated and unnegated. If $\neg \mathbf{P}_{>0}(\bigcirc \mathbf{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2)) \in s_n$ then we have that $\neg \mathbf{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2) \in s_{n+1}$. We can use the induction hypothesis.

- Formulas of the form $\xi = \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. Proof by contraposition — lets assume that $\mathcal{M}_s^{\mathcal{A}} \models \neg \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. Recall Lemma 1. There exists a finite path starting in s which either contains a state in neither Φ_1 nor Φ_2 hold or leads to a BSCC where $\neg\Phi_2$ holds. For the first case, the finite path $\pi = s_0, ..., s_n$ holds a positive probability. All states along the path satisfy $\neg\Phi_2$ and the last state satisfies $\neg\Phi_1$. Then this path contains a state which is not eligible. Proof by induction over n:

  Induction base ($n = 0$): For a path of length zero the state s_0 must fulfill the negated until formula. Therefore $\neg\Phi_1 \in s_0$. But since $\mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2) = \xi \in s_0$ we know that by completeness $\Phi_1 \in s_0$. Since $\Phi_1 \in s$ and $\neg\Phi_1 \in s$, the state s_0 is not eligible.

  Induction step: There are two possibilities for $\neg \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2) \in s = s_n$ — either both $\neg\Phi_1 \in s_n$ and $\neg\Phi_2 \in s_n$ — or $\neg \mathbf{P}_{=1}(\bigcirc \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)) \in s_n$. In the first case, with $\neg\Phi_1 \in s_n$ and $\neg\Phi_2 \in s_n$, we know that per definition of the finite path we have that $\Phi_1 \in s_n$, too. Then s_n is not eligible as Φ_1 is in the state both negated and unnegated.

  If $\neg \mathbf{P}_{=1}(\bigcirc \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)) \in s_n$ then we have that $\neg \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2) \in s_{n+1}$. We can use the induction hypothesis.

- Formulas of the form $\xi = \neg \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. Following the definition of the pseudo–model, either $\mathcal{M}_s^{\mathcal{A}} \models \neg \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$ or there is a suitable witness for ξ. Proof by contraposition — lets assume that $\mathcal{M}_s^{\mathcal{A}} \models \mathbf{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. Then there exists no such finite path π starting in s where at some point neither Φ_1 nor Φ_2.

  In the case where the algorithm aims to construct a simple pseudo–model, the non–existance of π would have lead to the deletion of s.

  If on the other hand a witness may be employed and there es no such witness for ξ in s, then s would have been deleted.

Theorem 3. *If Φ is (finite) satisfiable, then the algorithm returns a (simple) pseudo-model for Φ*

For this proof we first formulate another theorem:

Theorem 4. *Let Φ be a qualitative PCTL formula. If Φ is satisfiable, then there is a pseudo-model $\mathcal{A} = (A, \to)$ for Φ. Moreover, if Φ is finitely–satisfiable, then $\mathcal{A}$ is simple.*

We will prove this theorem by giving a construction. In accordance with the theorem we start with a satisfiable formula Φ. Since Φ is satisfiable, there exists a Markov chain $\mathcal{M} = (S, \rightarrow, L)$ containing a state $s \in S$ for which $\mathcal{M}_s \models \Phi$. We define the labeling function L using the closure of Φ:

$$L(s) := \{\Psi \in Cl(\Phi) | \mathcal{M}_s \models \Psi\}$$

We now build the state set A of the pseudo-model from the distinct sets of L, meaning we formulate an equivalence relation on $\mathcal{M}$ which merges two states iff they have the same labeling. The transition relation of $\mathcal{A}$ is also defined using $\mathcal{M}$. If there are two states s and t in $\mathcal{M}$ and $s \rightarrow t$, then in $\mathcal{A}$ there is a transition between the equivalent states.

To prove that $\mathcal{A}$ is a pseudo-model for Φ we will check the two conditions discussed earlier with pseudo-models. That all sets in A are eligible is assured by construction.

1. If for some $s \in A$ and $\xi \in s$ with ξ is of the form $\mathbb{P}_{=1}(\bigcirc\Phi_1)$, $\neg\mathbb{P}_{=1}(\bigcirc\Phi_1)$, $\mathbb{P}_{>0}(\bigcirc\Phi_1)$, $\neg\mathbb{P}_{>0}(\bigcirc\Phi_1)$, $\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$, $\neg\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$, or $\mathbb{P}_{=1}(\Phi_1\mathbf{U}\Phi_2)$, then $\mathcal{M}_s^{\mathcal{A}} \models \xi$. We check each type of ξ:

   - $\mathbb{P}_{=1}(\bigcirc\Phi_1)$: For each transition $s \rightarrow t$ in $\mathcal{A}$ there is a corresponding transition $s' \rightarrow t'$ in $\mathcal{M}$ and per construction ξ holds in s'. Therefore for each successor of s' the formula Φ_1 must hold. This therefore also holds for t.
   - $\neg\mathbb{P}_{>0}(\bigcirc\Phi_1)$: Analog construction over all successors.
   - $\neg\mathbb{P}_{=1}(\bigcirc\Phi_1)$: The corresponding state s' in $\mathcal{M}$ fulfills ξ, therefore there exists a successor t' in which Φ_1 does not hold. By construction this transition is also present in $\mathcal{A}$, leading to a state t, where, due to the labeling, Φ_1 does not hold.
   - $\mathbb{P}_{>0}(\bigcirc\Phi_1)$: Analog construction using the existence of a suitable successor.
   - $\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$: Since $\mathcal{M}_s \models \xi$ there exists a path $\pi = s'_0, ..., s'_n$ such that $\mathcal{M}_{s'_n} \models \Phi_2$ and for all $0 \leqslant i \leqslant n$ $\mathcal{M}_{s'_i} \models \Phi_1$. By construction for every state t in $\mathcal{A}$ that is equivalent to a state in π it holds that $\Phi_1 \in t$. The transitions along the path are also present in $\mathcal{A}$, alongside with the final state s_n.
   - $\neg\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$: Proof by contraposition. Assume that $\mathcal{M}_s^{\mathcal{A}} \models \mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$. Then there exists a path $\pi = s_0, ..., s_n$ starting in s as described above. We show by induction over the path length n that there exists a state that is not eligible.

     Induction base ($n = 0$): Then $\Phi_2 \in s_0 = s$. Because of completeness, from $\neg\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2) \in s$ it follows that $\neg\Phi_2 \in s$. Then $s = s_0$ is not eligible.

     Induction step ($n \rightarrow n+1$): Since $\neg\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2) \in s_n$, there are two options: Either the formula is rejected directly by having $\neg\Phi_1 \in s_n$ and $\neg\Phi_2 \in s_n$, but then per definition of the path $\Phi_1 \in s_n$ - s_n is not eligible. Otherwise $\neg\mathbb{P}_{>0}(\bigcirc\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)) \in s_n$. Therefore it must hold that $\neg\mathbb{P}_{>0}(\Phi_1\mathbf{U}\Phi_2) \in s_{n+1}$. We can apply the induction hypothesis.

- $\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$: Proof by contraposition. Assume that $\mathcal{M}_s^{\mathcal{A}} \models \neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. By Lemma 1 we know that there are two options. Either there exists a finite path $\pi = s_0, ..., s_n$ starting in s with $\neg\Phi_1 \in s_n$ and $\neg\Phi_2 \in s_n$. Then we know that $\mathcal{M}_{s_n} \models \mathbb{P}_{=0}(\Phi_1 \mathbf{U} \Phi_2)$. Therefore all states along the path necessarily fulfill $\mathbb{P}_{<1}(\Phi_1 \mathbf{U} \Phi_2)$. The initial transition $s_0' \to s_1'$ is a contradiction to $\mathcal{M}_{s_0} \models \mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. Otherwise there has to be a BSCC B of $\mathcal{A}$ such that $\neg\Phi_2 \in T$ for every $T \in B$ and a finite path $\pi = s_0, ..., s_n$ leading from s to B ($s_n \in B$). By construction for every state t in B the equivalent state t' in the Markov chain $\mathcal{M}$ it holds that $\mathcal{M}_t' \models \mathbb{P}_{=0}(\Phi_1 \mathbf{U} \Phi_2)$. Since that last state of π, the state s_n is also in B, we can apply the same argument as for the finite path.

2. This leaves the second condition regarding the formula $\xi = \neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$.

   We divide this proof into two parts. We show that if Φ is finitely–satisfiable, then the simple condition is sufficient, i.e. $\mathcal{M}_s^{\mathcal{A}} \not\models \mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$. Since in this case by definition $\mathcal{M}$ is a finite Markov chain, we again apply Lemma 1. As $\mathcal{M}_s \models \neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$ there exists a finite path starting in s that either leads to a single state refuting the property (e.g. fulfilling both $\neg\Phi_1$ and $\neg\Phi_2$) or a BSCC. We conclude that this finite path also exists in $\mathcal{A}$. Both conditions can be mimicked in $\mathcal{A}$ because both the final state and a BSCC also exists there.

   But this only works for the finite Markov chains. In the more general case where the model may be infinite, we know that since $\mathcal{M}_s \not\models \mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$ there are two possibilities. If there exists a finite path leading to a state which refutes $\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$, this path also exists in $\mathcal{A}$, as in the finite variant. If on the other hand a witness needs to be used, the proof is not as easy.

   Firstly, we observe that the amount of eligible states in $\mathcal{A}$ is bounded from above by $2^{|Cl(\Phi)|}$ and therefore finite. We conclude that even if an infinite amount of states of $\mathcal{M}$ is necessary to prove that ξ holds, in $\mathcal{A}$ this will always be equivalent to a finite number of states and therefore must include a loop that is traversed infinitely often. We construct our witness from $\mathcal{M}$ by selecting all those states which are traversed infinitely often, which is equivalent to the loop. Additionally, we select all those transitions which are traversed infinitely often. This forms the witness $\mathcal{B} = (B, \hookrightarrow)$.

   For now we have to proof that the witness as defined above is valid. The three conditions as given with the definition of a witness are checked in the following.

   Firstly, as $\mathcal{B}$ is built from a loop, it is strongly connected by construction.

   The same applies to the condition that $\neg\Phi_2$ holds in every state, as by construction we only viewed states s for which $\neg\Phi_2 \in s$.

   The last condition states that for every $t \in B$ and for every $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \in t$ it holds that $\mathcal{M}_t^{\mathcal{B}} \models \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$. We prove the statement by contraposition.

   We assume that $\mathcal{M}_t^{\mathcal{B}} \not\models \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$. Again, we apply Lemma 1.

   For the state t to refute the formula, there are two options.

(a) There exists a finite path $\pi = t_0, t_1, ..., t_n$ in $\mathcal{B}$ starting in $t = t_0$ such that $\neg\xi_2 \in t_i$ for every $0 \leqslant i \leqslant n$ and $\neg\xi_1 \in t_n$. Since $B \subseteq A$ and $\hookrightarrow \subseteq \rightarrow$, this path also exists in $\mathcal{A}$, therefore t is not eligible.

(b) For every state s in B it holds that $\xi_1 \in s$ and $\neg\xi_2 \in s$. We show that in that case $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \notin t$.

For $\mathcal{B}$ to be a witness, the probability mass of all runs through $\mathcal{B}$ must be greater than zero. Suppose there exists a state $u \in B$ such that the probability mass of all runs starting in u is zero. But $\mathcal{B}$ is strongly connected, therefore every state is reachable from every state - if u would exist, then the probability mass of the whole witness must equal zero, which is a contradiction.

We conclude that there exists a set of runs starting in t with non-zero probability for which ξ_2 never holds. By that we know that $\mathcal{M}_t^{\mathcal{B}} \not\models \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$ and therefore $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \notin t$.

We now use Theorem 4 in our proof of Theorem 3. As Φ is satisfiable, we know that there exists a pseudo-model $\mathcal{A}' = (A', \rightarrow')$ for Φ.

Here we use the algorithm and its structure to show that $\mathcal{A}'$ is an invariant of the inner loop from lines 9 to 35. The algorithm iteratively removes states which break the model. We therefore assume that $A' \subseteq A$.

We examine each inner if-condition and show, that no element in $\mathcal{A}'$ could be removed from $\mathcal{A}$:

- Line 11, formula $\mathbb{P}_{>0}(\bigcirc\Phi_1)$: If $\xi = \mathbb{P}_{>0}(\bigcirc\Phi_1)$ and $\mathcal{M}_s^{\mathcal{A}'} \models \xi$, then there exists a transition $s \rightarrow t$ in $\mathcal{A}'$ with $\Phi_1 \in t$. Since $\mathcal{A}' \subseteq \mathcal{A}$, t also exists in $\mathcal{A}$. Therefore $\mathcal{M}_s^{\mathcal{A}} \models \xi$.

- Line 11, formula $\neg\mathbb{P}_{=1}(\bigcirc\Phi_1)$: If $\xi = \neg\mathbb{P}_{=1}(\bigcirc\Phi_1)$ and $\mathcal{M}_s^{\mathcal{A}'} \models \xi$, then there exists a transition $s \rightarrow t$ in $\mathcal{A}'$ with $\Phi_1 \notin t$. Since $\mathcal{A}' \subseteq \mathcal{A}$, t also exists in $\mathcal{A}$. Therefore $\mathcal{M}_s^{\mathcal{A}} \models \xi$.

- Line 11, formula $\mathbb{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2)$: If $\xi = \mathbb{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2)$ and $\mathcal{M}_s^{\mathcal{A}'} \models \xi$, then there exists a path $\pi = s_0, ..., s_n$ in $\mathcal{A}'$ with $\Phi_1 \in s_i$ for all $0 \leqslant i \leqslant n$ and $\Phi_2 \in s_n$. Since $\mathcal{A}' \subseteq \mathcal{A}$, π also exists in $\mathcal{A}$. Therefore $\mathcal{M}_s^{\mathcal{A}} \models \xi$.

- Line 14, formula $\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$: Suppose there exists a counterexample, i.e. a BSCC B where for every $s \in B$ we have that $\Phi_1 \in s$, $\neg\Phi_2 \in s$ and $\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2) \in s$. Assume that s is a state of B and is a part of $\mathcal{A}'$: $s \in A'$. Because $\mathcal{A}'$ is a valid pseudo–model and $\xi \in s$, we know that there must exist a path π starting in s and leading to a state t for which $\Phi_2 \in t$. This is a contradiction to the assumption that s is part of a BSCC where $\neg\Phi_2 \in u$ for every $u \in B$. This path would also exist in $\mathcal{A}$.

- Line 21, formula $\neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$: First assume that $\mathcal{A}'$ is simple. Since $\xi = \neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$ and $\xi \in s$ for some $s \in A'$, there are two possibilities: Either there exists a finite path $\pi = s_0, ..., s_n$ starting in s on which $\Phi_1 \in s_i$ for all

$0 \leqslant i \leqslant n$ and $\neg\Phi_1 \in s_n$ and $\neq \Phi_2 \in s_n$, or a reachable BSCC B in which for all states $t \in B$ it holds that $\Phi_1 \in t$ and $\neg\Phi_2 \in t$. The path leading to the BSCC B is equivalent to the first option in the sense that for paths the last state fulfills $\mathbb{P}_{=0}(\Phi_1 \mathbf{U} \Phi_2) = \neg\mathbb{P}_{>0}(\Phi_1 \mathbf{U} \Phi_2)$. Therefore the required path for the condition in line 22 exists and also exists in $\mathcal{A}$.

Now $\mathcal{A}'$ must not necessarily be simple. If a witness is required in $\mathcal{A}'$, then there exists a path leading from s to the witness. Because of $\mathcal{A}' \subseteq \mathcal{A}$ this path also exists in $\mathcal{A}$. By the same argument the witness is embedded in $\mathcal{A}$.

Now all that remains is the initialization of $\mathcal{A}$ from line 1 to line 8. Both A and $\rightarrow$ are initialized to their maximal extent. There exists no state that could be in $\mathcal{A}'$ but not in this initial form of $\mathcal{A}$, because per construction $\mathcal{A}' \subset \{s \in Cl(\Phi) | s \text{ is eligible}\}$. All transitions removed in lines 4 and 7 cannot be in $\mathcal{A}'$ since it would immediately invalidate the model.

We established that $\mathcal{A}' \subseteq \mathcal{A}$ after the algorithm terminates. Since there exists a state $s \in A'$ for which $\Phi \in s$, this state also exists in $\mathcal{A}$ and $\mathcal{A}$ is a model for Φ.

Theorem 5. *The Algorithm terminates within a time that is exponential in* $|\Phi|$.

We first identify sub-problems that need to be solved in the algorithm. In many instances PCTL model checking has to be performed on finite-state models. This is polynomial in the length of Φ and the size of $\mathcal{A}$.

The computation of a witness for a formula $\neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2) \in s$ from some $s \in A$ can be done as follows.

1. Remember that a witness contains only states in which $\neg\Phi_2$ holds. We first filter the set A using this condition: Let $B := \{s \in A | \neg\Phi_2 \in s\}$.

2. The second condition for a witness requires the states to be strongly connected. We therefore split B into m disjoint strongly connected components $\mathcal{B}_i := (B_i, \hookrightarrow_i)$ with $s \hookrightarrow_i t$ iff $s, t \in B_i$ and $s \rightarrow t$.

3. Now for the third condition we need to check whether in each strongly connected component all formulas of the form $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$ are fulfilled. To do this we first identify the set of all states C such that for each $s \in C$ with $s \in B_i$ for some i there exists a formula $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$ and $\mathcal{M}_s^{\mathcal{B}_i} \not\models \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$.

4. Now we set $B := B \backslash C$. Since this may invalidate the second condition, we iterate these steps again. Should the set C be empty, the computation converged and is finished.

The aforementioned steps can be performed in time polynomial in the length of Φ and $\mathcal{A}$, as step 1 is linear in $|A|$, step 2 is quadratic in $|A|$, step 3 is polynomial in $|\xi_1|$, $|\xi_2|$ and $|A|$.

To show that this computation is correct, we first proof that no state of a witness in $\mathcal{A}'$ can be deleted in step 2. Let s be a state of a witness $\mathcal{B} = (B, \hookrightarrow)$ for a formula $\neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$.

Proof by contraposition - we assume that s is deleted in step 2. Hence there exists a formula $\xi = \mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$ such that $\xi \in s$ but $\mathcal{M}_s^{\mathcal{B}_i} \not\models \xi$ for some i. By Lemma 1 we know that there are two possibilities.

Firstly we assume that there exists a finite path that acts as a counterexample to ξ. The idea is that by construction and eligibility we know that if $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \in s$, then either $\xi_2 \in s$ or $\mathbb{P}_{=1}(\bigcirc \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)) \in s$. The first case ($\xi_2 \in s$) immediately results in a contradiction since clearly $\mathcal{M}_s^{\mathcal{B}_i} \models \xi$. In the second case we see that every successor of s must fulfill the formula ξ. If there would exist a path which contradicts ξ in $\mathcal{B}_i$, it also exists in $\mathcal{A}'$ and defies the eligibility of s.

In the second case, we assume that there exists a reachable BSCC in which $\xi_1, \neg\xi_2 \in s$ for every state t of the BSCC. Note that $\mathcal{B}_i$ is by construction the only reachable BSCC from s. This implies that this witness and BSCC is also a witness for $\neg\xi$, which contradicts the assumption that $\mathcal{B}_i$ was a witness to begin with.

By construction we find that every $\mathcal{B}_i$ we constructed in this fashion is a correct witness for a formula $\neg\mathbb{P}_{=1}(\Phi_1 \mathbf{U} \Phi_2)$, since the three requirements of witnesses are fulfilled.

We can now combine the Theorem 2, which states that if the algorithm returns a pseudo-model, it is a valid one, and the Theorem 3, which states that if a formula is satisfiable, then the algorithm will return a pseudo-model, to proof that this algorithm returns a pseudo-model for Φ if and only if the formula Φ is satisfiable. Additionally we have shown that the algorithm terminates in time which is exponential in $|\Phi|$.

Theorem 6. *If there is a pseudo-model $\mathcal{A}$ for Φ, then there is a marked graph $\mathcal{G} = (G, \hookrightarrow, L, M)$ whose size is exponential in $|\Phi|$, a valuation function $L : G \to 2^{AP}$ and the set of marked transitions $M \subseteq \hookrightarrow$ and a state $s \in G$ such that $\mathcal{M}_s^{\mathcal{G}} \models^L \Phi$. Moreover, if $\mathcal{A}$ is simple, then $L = \varnothing$ and Φ has a finite-state model whose size is exponential in $|\Phi|$.*

We prove this theorem by providing a construction for the marked graph given a pseudo-model $\mathcal{A} = (A, \to)$. If $\mathcal{A}$ is a simple pseudo-model, then we set $\mathcal{G} = (A, \to, L, \varnothing)$, with the labeling given by $\mathcal{A}$, i.e. $L(s) := \{ap \in AP | ap \in s\}$ for every $s \in A$.

As there are no marked transitions, the markov chain induced by the marked graph is equivalent to the transition system formed by $\mathcal{A}$, therefore all formulas formulas of a state in A also hold on the equivalent state in $\mathcal{G}$.

If on the other hand $\mathcal{A}$ is not simple, we apply the following construction. Since $\mathcal{A}$ is not simple, witnesses are employed. Let $\mathcal{B}_i = (B_i, \rightsquigarrow)$, $1 \leqslant i \leqslant m$ be all used witnesses such that for every $s \in A$ and every $\neg\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \in s$ there is a suitable $\mathcal{B}_i$ and a suitable finite path leading from s to $\mathcal{B}_i$. Because there are at most $|\Phi|$ such formulas requiring a witness, the number of witnesses required is bounded from above by $m \leqslant |A| \cdot |\Phi|$.

We construct the states of $\mathcal{G}$ by taking the disjoint union of $\mathcal{A}$ and $\mathcal{B}_i$ for $1 \leqslant i \leqslant m$. For the ease of reading we now refer to $\mathcal{A}$ as $\mathcal{B}_0$. We define the transition relation $\hookrightarrow$ of $\mathcal{G}$ as follows. For every $s, t \in A$ with $s \to t$:

- $(s,0) \hookrightarrow (t,i)$ for every $0 \leqslant i \leqslant m$ such that $t \in B_i$.
- For every $1 \leqslant i \leqslant m$ with $s,t \in B_i$ and $s \rightsquigarrow_i t$ there is a transition $(s,i) \hookrightarrow (t,i)$ and this transition is marked.
- For every $1 \leqslant i \leqslant m$ with $s \in B_i$ and $t \notin B_i$ there is a transition $(s,i) \hookrightarrow (t,0i)$.

We say that a run w *stays* at i if for all $k \in \mathbb{N}_0$ it holds that $w(k) \in \mathcal{B}_i \times i$. We say that a run w *enters* i if after some finite initial prefix of w, the run w' stays at i.

We present three preliminary observations about $\mathcal{G}$:

- For every $(s,i) \in G$ the probability of all $w \in Run((s,i))$ staying at i for $i > 0$ is at least $\frac{2}{3}$, as by construction the probability of taking a non-marked transition is bounded from above by $\sum_{k=1}^{\infty} \frac{1}{4}^k = \frac{1}{3}$.
- For every $(s,i) \in G$ the probability of all $w \in Run((s,i))$ such that w does not enter any j for $0 \leqslant j \leqslant m$, is zero, since the probability of always leaving a component converges against zero.
- For every $(s,i) \in G$ and every $\mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2) \in s$ the conditional probability of all $w \in Run((s,i))$ such that $\mathcal{M}_s^{\mathcal{G}} \models \mathbb{P}_{=1}(\xi_1 \mathbf{U} \xi_2)$ under the condition that w stays at i is equal to one because of the definition of a witness.

By induction over the structure of $\Psi \in Cl(\Phi)$, it follows that for every formula Ψ in the closure of Φ, it is fulfilled in a state (s,i) by the induced markov chain of $\mathcal{G}$ if and only if $\Psi \in s$.

3 Conclusion

We have shown that for a given PCTL formula we can either find a finite description of a satisfying model or conclude its unsatisfiability in time exponential in the size of the formula. We presented an algorithm that iteratively removes parts of an abstract model–description that would violate logical permissibility. We have seen that a certain class of qualitative PCTL formulas need special attention as they may require infinite–state models to hold.

4 Bibliography

References

[1] Brázdil, T., Forejt, V., Kretinsky, J., and Kucera, A. (2008). The satisfiability problem for probabilistic ctl. In Logic in Computer Science, 2008. LICS'08. 23rd Annual IEEE Symposium on, pages 391–402. IEEE.

[2] Emerson, E. A. and Halpern, J. Y. (1982). Decision procedures and expressiveness in the temporal logic of branching time. In Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing, STOC '82, pages 169–180, New York, NY, USA. ACM.

Input: A qualitative PCTL formula Φ.

Output: A (simple) pseudo–model $\mathcal{A} = (A, \rightarrow)$ if Φ is (finite) satisfiable, *unsatisfiable* otherwise.

1: $A :=$ the set of all eligible subsets of $Cl(\Phi)$
2: $\rightarrow := A \times A$
3: **for all** $S \in A, \mathbf{P}_{=1}(\bigcirc\Phi_1) \in S$ **do**
4: delete all edges $S \rightarrow T$ where $\Phi_1 \notin T$
5: **end for**
6: **for all** $S \in A, \neg\,\mathbf{P}_{>0}(\bigcirc\Phi_1) \in S$ **do**
7: delete all edges $S \rightarrow T$ where $\Phi_1 \in T$
8: **end for**
9: **repeat**
10: **for all** $S \in A, \xi \in S$ **do**
11: **if** $\xi \equiv \mathbf{P}_{>0}(\bigcirc\Phi_1)$ or $\xi \equiv \neg\,\mathbf{P}_{=1}(\bigcirc\Phi_1)$ or $\xi \equiv \mathbf{P}_{>0}(\Phi_1\mathbf{U}\Phi_2)$ **then**
12: **if** $\mathcal{M}_s^{\mathcal{A}} \not\models \xi$ **then** $A := A\backslash\{S\}$ **end if**
13: **end if**
14: **if** $\xi \equiv \mathbf{P}_{=1}(\Phi_1\mathbf{U}\Phi_2)$ **then**
15: **for all** BSCC B of $(A, \rightarrow)$ **do**
16: **if** $\Phi_1, \neg\Phi_2, \mathbf{P}_{=1}(\Phi_1\mathbf{U}\Phi_2) \in T$ for every $T \in B$ **then**
17: $A := A\backslash B$
18: **end if**
19: **end for**
20: **end if**
21: **if** $\xi \equiv \neg\,\mathbf{P}_{=1}(\Phi_1\mathbf{U}\Phi_2)$ **then**
22: **if** there is no finite path $\pi = s_0, ..., s_n$ where $\neg\,\mathbf{P}_{>0}(\Phi_1\mathbf{U}\Phi_2) \in s_n$ and $\Phi_1, \neg\Phi_2 \in s_i$ for all $0 \leqslant i \leqslant n$ **then**
23: **if** CREATE_SIMPLE **then**
24: $A := A\backslash\{S\}$
25: **else if** there is no witness $(B, \hookrightarrow)$ for $\neg\,\mathbf{P}_{=1}(\Phi_1\mathbf{U}\Phi_2)$ in $(A, \rightarrow)$ such that there is a finite path $\pi = s_0, ..., s_n$ where $s_n \in B$ and $\neg\Phi_2 \in s_i$ for all $0 \leqslant i \leqslant n$ **then**
26: $A := A\backslash\{S\}$
27: **end if**
28: **end if**
29: **end if**
30: **repeat**
31: $\rightarrow := \rightarrow \cap (A \times A)$
32: $A := A\backslash\{S \in A | S$ has no outgoing edges$\}$
33: **until** $(A, \rightarrow)$ does not change
34: **end for**
35: **until** $(A, \rightarrow)$ does not change
36: **if** $\Phi \in S$ for some $S \in A$ **then return** $\mathcal{A} = (A, \rightarrow)$
37: **end if**
38: **return** *unsatisfiable*

Figure 1: An algorithm for constructing a (simple) pseudo–model.