Masaryk University
Faculty of Informatics

The satisfiability problem for probabilistic temporal logics

Master's Thesis

Miroslav Chodil

Brno, Spring 2019

This is where a copy of the official signed thesis assignment and a copy of the Statement of an Author is located in the printed version of the document.

Declaration

Hereby I declare that this paper is my original authorial work, which I have worked out on my own. All sources, references, and literature used or excerpted during elaboration of this work are properly cited and listed in complete reference to the due source.

Miroslav Chodil

Advisor: prof. RNDr. Antonín Kučera, Ph.D.

Acknowledgements

I would like to express my sincerest gratitude to my advisor, prof. RNDr. Antonín Kučera, Ph.D., for his invaluable advice and patience.

Abstract

In this thesis, we study the satisfiability problem for Probabilistic CTL (PCTL), which is a probabilistic extension of the well-known temporal logic CTL. PCTL is a popular formalism for specifying properties of stochastic systems. The satisfiability problem for PCTL is a longstanding open problem interesting from a theoretical perspective, although its solution would also have various practical applications, such as consistency checking for specifications.

We give an overview of known results. We find two (closely related) criteria that would imply decidability of the finite satisfiability problem for qualitative PCTL extended with bounded until formulae. We show that bounded PCTL extended by allowing a single use of $G^{=1}$ does not have the small model property. We show that the finite and general satisfiability problems are decidable for a fragment of PCTL that consists of conjunctions of formulae of form $P(\varphi \mathrm{U} \psi) \in I$ where I is an open interval and φ, ψ are qualitative PCTL formulae.

Keywords

pctl, probabilistic ctl, satisfiability, probabilistic temporal logics

Contents

1 Introduction **1**

2 Preliminaries **3**
2.1 *Sets* 3
2.2 *Sequences* 3
2.3 *Probability spaces* 5
2.4 *Markov chains* 6
2.5 *PCTL* 8
2.6 *Bounded until formulae* 14

3 The satisfiability problem for probabilistic temporal logics **19**
3.1 *Related non-probabilistic logics* 20
3.2 *PCTL* 20
3.3 *Decidable fragments of PCTL* 23
3.4 *Further related work* 29

4 Qualitative PCTL extended with bounded until formulae **31**

5 Bounded PCTL extended with a single use of $\mathrm{G}^{=1}$ **45**

6 Formulae of form $\bigwedge_{i \in I} P(\varphi_i \mathrm{U} \psi_i) \in (l_i, r_i)$ where φ_i and ψ_i are qualitative PCTL formulae **57**
6.1 *Types of runs* 59
6.2 *Three model-preserving qPCTL constructions* 64
6.3 *Pseudo-models for qPCTL formulae* 69
6.4 *Decidability of finite and general satisfiability* 74

7 Conclusion **85**

Bibliography **87**

1 Introduction

Probabilistic CTL [1] (further called PCTL) is a probabilistic extension of the well-known branching-time temporal logic CTL [2]. PCTL is a popular formalism for specifying properties of stochastic systems. In PCTL, the universal and existential path quantifiers from CTL are replaced with a probabilistic operator, which allows it to quantify over probability of runs that satisfy some path formula. PCTL state formulae are built from atomic propositions, negations, conjunctions, and the probabilistic operator $P(\Phi) \oplus q$, where Φ is a PCTL path formula, q is a rational constant between 0 and 1, and $\oplus$ is one of $\{<, \leq, \geq, >, =, \neq\}$. PCTL path formulae are of form $\mathrm{X}\,\varphi$ ("next") or $\varphi \,\mathrm{U}\, \psi$ ("until"), where φ and ψ are PCTL state formulae. The usual $\mathrm{G}\,\varphi$ and $\mathrm{F}\,\varphi$ path formulae which express that φ is always satisfied and that φ is satisfied at some point in future (respectively) can be derived from until formulae. We use $\varphi \,\mathrm{U}^{\oplus q}\, \psi$ as a shortcut for $P(\varphi \,\mathrm{U}\, \psi) \oplus q$, and similarly for $\mathrm{X}\,\varphi$. A PCTL path formula $a \,\mathrm{U}\, b$ expresses that at some point a state satisfying b is reached, and until then all states had to satisfy a. A PCTL state formula $\mathrm{G}^{=1}(\mathit{sent} \implies \mathrm{F}^{=1}\, \mathit{received})$ expresses that a state satisfying *received* is reached almost surely from states that satisfy *sent*. PCTL formulae are evaluated over Markov chains with states labeled by atomic propositions.

In this thesis, we study the satisfiability problem for PCTL, which asks whether a given PCTL state formula φ has a (finite) model, that is, whether there exists a (finite) Markov chain M and a labeling on its states such that some state of M satisfies φ. The satisfiability problem for PCTL (further shortened to satisfiability) is a longstanding open problem that is interesting from a purely theoretical standpoint, although its solution would also have various practical applications, such as consistency checking of specifications.

Most results about PCTL are concerned with its fragments. A recent work [3] considers satisfiability for various fragments of PCTL based on operators F and G. Decidability of satisfiability, both in its general and finite version, has been shown for a fragment of PCTL called bounded PCTL [4]. Qualitative PCTL is a fragment of PCTL obtained by restricting all probabilistic constraints to 0 and 1. Decidability of satisfiability for qualitative PCTL and also some of its variants has been

shown in [5, 6, 7, 8]. In [8], it is shown that the problem of determining whether a PCTL formula has a model with branching degree at most k, where $k \geq 2$ is an arbitrary but fixed constant, is highly undecidable (k is a fixed parameter, so for each k we obtain a different problem). In the same paper, it is also shown that whenever a PCTL formula φ has a model, it must also have a model with branching degree at most $|\varphi| + 2$. We give a more thorough overview of known results in Chapter 3.

This thesis is organized as follows.

In chapter 2, we introduce definitions, notation, and some standard lemmas used throughout this thesis.

In chapter 3, we give an overview of existing results about PCTL and related logics, with added focus on PCTL fragments relevant to this thesis.

In chapter 4, we find two (closely related) criteria that would imply decidability of the finite satisfiability problem for qualitative PCTL extended with bounded until formulae.

In chapter 5, we show that extending bounded PCTL by allowing a single use of $\mathrm{G}^{=1}$ breaks the small model property.

In chapter 6, we show that the finite and general satisfiability problems are decidable for a fragment of PCTL that consists of conjunctions of formulae of form $P(\varphi \,\mathrm{U}\, \psi) \in I$ where I is an open interval and φ, ψ are qualitative PCTL formulae.

2 Preliminaries

In this chapter, we introduce definitions, notation, and a few lemmas used throughout this thesis. The definitions are taken primarily from [9], [10], and [1].

2.1 Sets

We use $\mathbb{N}, \mathbb{Q}, \mathbb{R}$ to denote the set of natural numbers with zero, rational numbers, and real numbers, respectively. Intervals of real numbers are also denoted using the standard notation. For example:

- $[0, 1]$ is the set $\{x \in \mathbb{R} \mid 0 \leq x \leq 1\}$,
- $(0, 1]$ is the set $\{x \in \mathbb{R} \mid 0 < x \leq 1\}$.

The empty set is denoted by $\emptyset$. If S is a set, then its cardinality is denoted by $|S|$, and its power set is denoted by $\mathcal{P}(S)$. The cardinality of $\mathbb{N}$ is also denoted by ω. When we say countable, we mean either finite or countably infinite. Domain and image of a function f are denoted by $Dom(f)$ and $Im(f)$, respectively. A set T is a partition of a set S if $\bigcup T = S$ and elements of T are pairwise disjoint (we allow $\emptyset \in T$).

2.2 Sequences

We shall now define some basic notions about sequences of elements. Our definitions are standard, and are only included to provide a proper basis for further notions.

Definition 2.2.1 (Infinite sequence)**.** An infinite sequence of elements of a set S is a function $\pi \colon \mathbb{N} \to S$.

Definition 2.2.2 (Finite sequence)**.** A finite sequence of elements of a set S is a function $\pi \colon \mathbb{N}_k \to S$, where $\mathbb{N}_k$ is the set of the first k natural numbers for some $k \in \mathbb{N}$. We consider the empty function a finite sequence, called the empty sequence ϵ.

Definition 2.2.3 (Sequence)**.** We define a sequence to mean either a finite or an infinite sequence.

Definition 2.2.4 (Length of a sequence)**.** Length of a sequence π is defined as the cardinality of its domain, and is denoted by $len(\pi)$.

In particular, $len(\epsilon) = 0$ and length of every infinite sequence is ω.

Definition 2.2.5 (S^ω, S^*, S^n)**.** Let S be a set. Define S^ω and S^* to be the sets of all infinite and all finite sequences of elements of S, respectively. Also for every $n \in \mathbb{N}$, define S^n to be the set of all sequences of elements of S with length n.

Note that $S^* = \bigcup_{n\in\mathbb{N}} S^n$. If S is countable, then $|S^n| \leq \omega$ for every $n \in \mathbb{N}$, which means that S^* is a countable union of countable sets, and therefore must be countable.

While sequences are formally defined as functions, we will want a more convenient way to talk about them:

Convention. When we say $s_0 s_1 s_2 \ldots$ is an infinite sequence, we mean a sequence π where $\pi(i) = s_i$ for all the specified $i \in \mathbb{N}$. A completely analogous notation is used for finite sequences.

We intuitively interpret $\pi(0)$ as the first element of the sequence, and generally $\pi(i)$ as the $i+1$-th element of the sequence. In line with this, we adopt the following convention:

Convention. Let π be a sequence. We say π *starts with* s if $\pi(0) = s$, we say π *goes through* s or π *contains* s or π *reaches* s if $s \in Im(\pi)$. When π is finite, we also say π *ends with* s if $\pi(len(\pi) - 1) = s$. We extend this notation to a finite sequences and sets of elements in an analogous way.

Definition 2.2.6 (Prefix of a sequence)**.** A sequence p is a prefix of a sequence π if and only if p is finite and $p \subseteq \pi$.

This definition matches our intuition about prefixes of sequences: p being a subset of π simply means that if p is defined for some i, then π is also defined for i and has the same value. Since prefixes are *finite* sequences, p must be the sequence of the first k elements of π for some $k \in \mathbb{N}$.

Definition 2.2.7 ($Pref(\pi)$ and $Pref(S)$)**.** Let π be a sequence. Define $Pref(\pi)$ as the set of all prefixes of π. Also if S is a set of sequences, define $Pref(S) = \bigcup_{\pi\in S} Pref(\pi)$.

2.3 Probability spaces

The semantics of PCTL formulae requires some basic notions from probability theory. We start with σ-algebras:

Definition 2.3.1 (σ-algebra)**.** Let Ω be a set. A σ-algebra over Ω is a set $\mathcal{F} \subseteq \mathcal{P}(\Omega)$ such that:

- $\Omega \in \mathcal{F}$,
- if $A \in \mathcal{F}$, then its complement $\Omega \setminus A \in \mathcal{F}$,
- $\mathcal{F}$ is closed under countable unions: if $\{A_i\}_{i \in I} \subseteq \mathcal{F}$ is a collection of sets indexed by a countable set I, then also $\bigcup_{i \in I} A_i \in \mathcal{F}$.

The set Ω is usually called a *sample space*, and its elements represent individual indivisible outcomes of a probabilistic experiment. The elements of $\mathcal{F}$ are called *events* or *measurable sets*. We can now define measures on σ-algebras. However, in this thesis we will only need a specific kind of measures called probability measures:

Definition 2.3.2 (Probability measure)**.** Let $\mathcal{F}$ be a σ-algebra over a set Ω. A function $\mathbb{P} \colon \mathcal{F} \to [0, 1]$ is called a probability measure on $\mathcal{F}$ if it satisfies the following conditions:

- $\mathbb{P}(\Omega) = 1$,
- $\mathbb{P}$ is countably additive: if $\{A_i\}_{i \in I} \subseteq \mathcal{F}$ is a collection of pairwise disjoint sets indexed by a countable set I, then $\mathbb{P}(\bigcup_{i \in I} A_i) = \sum_{i \in I} \mathbb{P}(A_i)$.

Intuitively, the second condition means that whenever we can partition a measurable set into countably many measurable sets, the measure of the original set is equal to the sum of measures of the constituent sets. Finally, we consolidate our definitions:

Definition 2.3.3 (Probability space)**.** A probability space is a triple $(\Omega, \mathcal{F}, \mathbb{P})$ where $\mathcal{F}$ is a σ-algebra over Ω and $\mathbb{P}$ is a probability measure on $\mathcal{F}$.

Probability spaces are also frequently called *probability triples* or *probability measure spaces*. We state some basic theorems about them for completeness, and use them without reference:

Lemma 2.3.1. *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Then the following holds:*

(*i*) $\mathbb{P}(\emptyset) = 0$,

(*ii*) $\mathbb{P}(A) = 1 - \mathbb{P}(\Omega \setminus A)$ *for every* $A \in \mathcal{F}$,

(*iii*) $A \subseteq B$ *implies* $\mathbb{P}(A) \leq \mathbb{P}(B)$ *for every* $A, B \in \mathcal{F}$.

2.4 Markov chains

PCTL formulae are evaluated over Markov chains. Specifically, we will use countable-space, discrete-time, time-homogenous Markov chains. This allows our definition of Markov chains to be less general, but more convenient to work with:

Definition 2.4.1 (Markov chain)**.** A Markov chain is a tuple (S, P) where S is a *countable* set and $P\colon S \times S \to [0, 1]$ is a function such that $\Sigma_{t \in S} P(s, t) = 1$ for every $s \in S$.

The set S is called a *state space* and is interpreted as the set of states the chain can be in. The probabilities $P(s, t)$ are called *transition probabilities*. The value $P(s, t)$ represents the probability of the chain going from state s to state t.

Convention. When we speak about the size of a Markov chain (S, P), we are always referring to the cardinality of S.

Consider a unary function ρ_s obtained by fixing the first argument of P to a specific s, that is, $\rho_s(t) = P(s, t)$ for every $t \in S$. Using the definition of $P(s, t)$, clearly $\Sigma_{t \in S} \rho_s(t)$ is equal to 1. We call such functions *probability distributions*:

Definition 2.4.2 (Probability distribution)**.** Let A be a countable set. A function $\rho\colon A \to [0, 1]$ is called a probability distribution over A if $\Sigma_{a \in A} \rho(a) = 1$.

Since ρ_s is a probability distribution for every $s \in S$, it is easy to see that regardless of which state the chain is in, there always exists a state to which the chain transitions with probability strictly greater than 0. In other words, the chain can never get "stuck" in some state. Another concept we will need are *runs*, which represent possible sequences of states a Markov chain can go through:

Definition 2.4.3 (Run)**.** A run $\pi \in S^\omega$ in a Markov chain (S, P) is an infinite sequence of elements of S such that $P(\pi(i), \pi(i+1)) > 0$ for all $i \in \mathbb{N}$.

We will use sets of runs as sample spaces. However, these sample spaces will not consist of all runs in some Markov chain, but rather runs starting in a specific state. To that end, we define $Run^M(p)$ as the set of runs in M starting with a prefix p:

Definition 2.4.4 ($Run^M(p)$)**.** Let $M = (S, P)$ be a Markov chain and let $p \in S^*$. Define $Run^M(p)$ in the following way:

$$Run^M(p) = \{\pi \in S^\omega \mid \pi \text{ is a run in } M \text{ and } p \in Pref(\pi)\}.$$

Sets of form $Run^M(p)$ are often called *basic cylinders* or just *cylinders*. Including M as a formal parameter in all of our definitions would be cumbersome, so we adopt the following convention:

Convention. We assume M is clear from context and not include it as a formal parameter in our defined notions. Otherwise, we explicitly indicate it using subscript, superscript, function argument, or similar. For example, we usually say π *is a run* instead of saying π *is a run in* M, or $Run(p)$ instead of $Run^M(p)$.

We also define *paths*, which give us a convenient way to talk about runs and their prefixes:

Definition 2.4.5 (Paths)**.** Let $M = (S, P)$ be a Markov chain. A path is either a run or a prefix of a run. For each $p \in S^*$, define

- $Paths_{fin}(p)$ to be the set of all finite paths that start with p.
- $Paths(p) = Paths_{fin}(p) \cup Run(p)$.

Also define $Paths_{fin} = Paths_{fin}(\epsilon)$ and $Paths = Paths(\epsilon)$.

Definition 2.4.6 (Predecessor and successor set)**.** Let (S, P) be a Markov chain and let $s \in S$. We use ${}^\bullet s$ and $s^\bullet$ to denote the predecessor set and the successor set of s, respectively:

$${}^\bullet s = \{t \in S \mid P(t, s) > 0\},$$

$$s^\bullet = \{t \in S \mid P(s, t) > 0\}.$$

Using successor sets, we can now define branching degree in Markov chains:

Definition 2.4.7 (Branching degree)**.** Let $M = (S, P)$ be a Markov chain. Branching degree of a state $s \in S$ is defined as the cardinality of its successor set $s^\bullet$. Branching degree of M is defined as the maximum of branching degrees of states of M, if it exists.

Definition 2.4.8 (s^*)**.** Let (S, P) be a Markov chain and $s \in S$ one of its states. Define $s^* = \{t \in S \mid \exists \pi \in Run(s) \text{ such that } \pi \text{ reaches } t\}$.

Definition 2.4.9 (Strongly connected set)**.** Let (S, P) be a Markov chain and $T \subseteq S$. We call T a strongly connected set if and only if for every $s, t \in T$ we have $t \in s^*$.

Definition 2.4.10 (Strongly connected component (SCC))**.** Let (S, P) be a Markov chain and $\emptyset \subset T \subseteq S$. We call T a strongly connected component (SCC) if and only if T is a maximal strongly connected subset of S.

Definition 2.4.11 (Bottom strongly connected component (BSCC))**.** Let (S, P) be a Markov chain and $\emptyset \subset T \subseteq S$. We call T a bottom strongly connected component (BSCC) if and only if T is an SCC and for every $t \in T, s \in S$ we have that $P(t, s) > 0$ implies $s \in T$.

2.5 PCTL

We are now ready to define the syntax of PCTL formulae:

Definition 2.5.1 (PCTL syntax)**.** Let AP be a countably infinite set. The abstract syntax of *PCTL state formulae* is defined as follows:

$$\varphi ::= a \mid \varphi \wedge \varphi \mid \neg\varphi \mid P(\Phi) \oplus q$$

where $a \in AP, \oplus \in \{<, \leq, \geq, >, =, \neq\}$, $q \in [0, 1] \cap \mathbb{Q}$, and Φ is a *PCTL path formula*. *PCTL path formulae* are defined in the following way:

$$\Phi ::= \mathrm{X}\,\varphi \mid \varphi\,\mathrm{U}\,\varphi$$

As a sidenote, PCTL originally also includes bounded until path formulae of form $\varphi \, \mathrm{U}_n \, \psi$ where $n \in \mathbb{N}$, so our definition is technically a restriction (which is more practical for this thesis).

Intuitively, PCTL state formulae talk about properties of states and PCTL path formulae talk about properties of runs. Throughout this thesis, we often refer to PCTL state formulae and PCTL path formulae simply as *state formulae* and *path formulae*, respectively. Also if we say φ is a formula, we mean a state formula. To easily distinguish between state and path formulae, we adopt the following convention:

Convention. State formulae are usually denoted by lowercase letters of the Greek alphabet, such as φ or ψ. Path formulae are typically denoted by uppercase letters of the Greek alphabet, such as Φ or Ψ.

Models of PCTL formulae are Markov chains whose states are labeled with elements of AP. These elements are called *atomic propositions*, and represent basic indivisible statements about states. The labeling function $v\colon S \to \mathcal{P}(AP)$ is called a *valuation* and assigns to each state in S a set of atomic propositions it satisfies. We will lose nothing by always assuming we work with one fixed AP that contains all the propositions we will ever need:

Convention. We assume AP is a fixed set given by context, unless explicitly stated otherwise.

We will now move on to define the semantics of PCTL formulae. In order to do that, we need to formally define what it means for a state s to satisfy a formula of the form $P(\Phi) \oplus q$. We will use a standard construction from measure theory, which will give us a σ-algebra $\mathcal{F}_s$ over $Run(s)$ and a probability measure $\mathbb{P}_s$ on $\mathcal{F}_s$:

Definition 2.5.2 ($\mathcal{F}_s$ and $\mathbb{P}_s$)**.** Let $M = (S, P)$ be a Markov chain and let $s \in S$. Define $\mathcal{F}_s$ to be the unique σ-algebra over $Run(s)$ such that:

- $Run(p) \in \mathcal{F}_s$ for every $p \in Paths_{fin}(s)$,
- $\mathcal{F}_s$ is the smallest such algebra (when ordered by set inclusion).

Also define $\mathbb{P}_s$ to be the unique probability measure on $\mathcal{F}_s$ such that for every $s_0 s_1 s_2 \ldots s_n \in Paths_{fin}(s)$ we have:

$$\mathbb{P}_s(Run(s_0 s_1 s_2 \ldots s_n)) = P(s_0, s_1) \cdot P(s_1, s_2) \cdot \ldots \cdot P(s_{n-1}, s_n).$$

We will not go over a proof of existence of $\mathcal{F}_s$ or $\mathbb{P}_s$. Instead, we refer the reader to standard texts on the theory of Markov chains, such as [11]. We can now define the semantics of PCTL formulae:

Definition 2.5.3 (PCTL semantics)**.** Let AP be a countably infinite set, $M = (S, P)$ a Markov chain, $v\colon S \to \mathcal{P}(AP)$ a valuation function. Let $s \in S$ and ψ be a PCTL state formula. Define the satisfaction relation $\models$ recursively, depending on the form of ψ:

- $M, s \models^v a$ iff $a \in v(s)$,
- $M, s \models^v \neg\varphi$ iff not $M, s \models^v \varphi$,
- $M, s \models^v \varphi_1 \wedge \varphi_2$ iff $M, s \models^v \varphi_1$ and $M, s \models^v \varphi_2$,
- $M, s \models^v P(\Phi) \oplus q$ iff $\mathbb{P}_s(\{\pi \in Run(s) \mid M, \pi \models^v \Phi\}) \oplus q$.

For a path formula Φ and a run π, we define $\models$ in the following way:

- $M, \pi \models^v \mathrm{X}\,\varphi$ iff $M, \pi(1) \models^v \varphi$,
- $M, \pi \models^v \varphi_1 \,\mathrm{U}\, \varphi_2$ iff there exists $i \in \mathbb{N}$ such that $M, \pi(i) \models^v \varphi_2$ and $M, \pi(k) \models^v \varphi_1$ for all $k \in \mathbb{N}$ such that $k < i$.

It is not too hard to show that sets of runs PCTL can "talk about" (called *PCTL events*) are measurable. We will not show a proof of that (see e.g. [9]), although its main idea is that for every path formula Φ and a state s, the set of runs starting in s that satisfy Φ can be expressed as a countable union of basic cylinders, and therefore must be in $\mathcal{F}_s$.

Lemma 2.5.1 (PCTL events are measurable)**.** *Let (S, P) be a Markov chain and v a valuation on its states. Suppose $s \in S$ and Φ is a PCTL path formula. Then $\{\pi \in Run(s) \mid M, \pi \models^v \Phi\} \in \mathcal{F}_s$.*

We further streamline our notation:

Convention. Same as with a Markov chain M, we often assume that the valuation v is clear from context, and is therefore left out in favor of brevity. For example $M, s \models^v \varphi$ may now become $s \models \varphi$.

Convention. Whenever we say $\{s \models \Phi\}$, we formally mean the set $\{\pi \in Run(s) \mid \pi \models \Phi\}$. When used as an argument to $\mathbb{P}_s$, we also drop the braces around it and the subscript s, which is clear from the argument itself. For example $\mathbb{P}_s(\{\pi \in Run(s) \mid \pi \models \Phi\})$ may be written as $\mathbb{P}(s \models \Phi)$.

Convention. When we say (S, P, v) is a Markov chain, we mean that (S, P) is a Markov chain and $v\colon S \to \mathcal{P}(AP)$ is a valuation on its states. This convention extends to properties of Markov chains and valuations in the obvious way. For example, if we say (S, P, v) is a finite Markov chain and s is one its states, we obviously mean that (S, P) is a finite Markov chain and $s \in S$.

We will now define some standard convenience notions:

Definition 2.5.4 ($\mathrm{X}^{\oplus q}\,\varphi_1$ and $\varphi_1\,\mathrm{U}^{\oplus q}\,\varphi_2$)**.** Let φ_1, φ_2 be a state formulae. Define $\mathrm{X}^{\oplus q}\,\varphi_1$ and $\varphi_1\,\mathrm{U}^{\oplus q}\,\varphi_2$ to be syntactic shortcuts for formulae $P(\mathrm{X}\,\varphi_1) \oplus q$ and $P(\varphi_1\,\mathrm{U}\,\varphi_2) \oplus q$, respectively.

Definition 2.5.5 ($\top$ and $\bot$)**.** Define $\top \equiv a \vee \neg a$ and $\bot \equiv a \wedge \neg a$ for an arbitrary but fixed $a \in AP$.

Definition 2.5.6 (F and G)**.** For a state formula φ, define $\mathrm{F}\,\varphi \equiv \top\,\mathrm{U}\,\varphi$ and $P(\mathrm{G}\,\varphi) = q$ to be a syntactic alias for $P(\mathrm{F}\,\neg\varphi) = 1 - q$ where $q \in [0, 1] \cap \mathbb{Q}$. Extend this notation to other $\oplus$ analogously. Intuitively, $\mathrm{G}\,\varphi \equiv \neg(\mathrm{F}\,\neg\varphi)$, but that is not a PCTL path formula. Still, we use $\pi \models \mathrm{G}\,\varphi$ as a shortcut for $\pi(i) \models \varphi$ for all $i \in \mathbb{N}$.

Definition 2.5.7 ($P(\Phi) \in I$)**.** Let Φ be a path formula and suppose $I = (a, b)$ for some $a, b \in [0, 1] \cap \mathbb{Q}$ such that $a < b$. Define $P(\Phi) \in I$ to be the formula $P(\Phi) > a \wedge P(\Phi) < b$.

Definition 2.5.8 ($s^{\bullet}_{\varphi}$)**.** Let $M = (S, P, v)$ be a Markov chain. For every state formula φ and $s \in S$, define $s^{\bullet}_{\varphi} = \{t \in s^{\bullet} \mid t \models \varphi\}$.

Definition 2.5.9 ($\widehat{\varphi}$)**.** Let φ be a state formula. If φ is of form $\neg\psi$, then we define $\widehat{\varphi} \equiv \psi$, otherwise $\widehat{\varphi} \equiv \neg\psi$.

Definition 2.5.10 ($sub(\psi)$)**.** Let ψ be a PCTL state formula. The set of its subformulae $sub(\psi)$ is defined recursively, in the usual way:

- $sub(a) = \{a\}$,
- $sub(\neg\varphi) = \{\neg\varphi\} \cup sub(\varphi)$,
- $sub(\varphi_1 \wedge \varphi_2) = \{\varphi_1 \wedge \varphi_2\} \cup sub(\varphi_1) \cup sub(\varphi_2)$,
- $sub(\mathrm{X}^{\oplus q}\,\varphi) = \{\mathrm{X}^{\oplus q}\,\varphi\} \cup sub(\varphi)$,

- $sub(\varphi_1 \, \mathrm{U}^{\oplus q} \, \varphi_2) = \{\varphi_1 \, \mathrm{U}^{\oplus q} \, \varphi_2\} \cup sub(\varphi_1) \cup sub(\varphi_2)$.

We will frequently use the following standard lemma about finite Markov chains (proof e.g. in [9]):

Lemma 2.5.2. *Let $M = (S, P, v)$ be a finite Markov chain and let $s \in S$. Then runs in $Run(s)$ that do not reach some BSCC have measure 0, that is, a BSCC is reached almost surely. Moreover, runs in $Run(s)$ that reach a BSCC reach every state in it infinitely often almost surely.*

The previous lemma can be used to show the following standard result (proof e.g. in [8]):

Lemma 2.5.3. *Let $M = (S, P, v)$ be a Markov chain, $s \in S$ be state, and $\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2$ be a state formula. Suppose $s \not\models \varphi_1 \, \mathrm{U}^{=1} \, \varphi_2$. Then one of the following holds:*

(i) There exists a finite path $s_0 s_1 s_2 \ldots s_n$ such that $s_0 = s$ and $s_n \models \neg\varphi_1$ and $s_i \models \neg\varphi_2$ for all $i \in \mathbb{N}$ such that $0 \leq i \leq n$.

(ii) $\mathbb{P}_s(A) > 0$, where A is the set of all runs $\pi \in Run(s)$ such that $\pi(i) \models \varphi_1 \wedge \neg\varphi_2$ for all $i \in \mathbb{N}$. Moreover, if S is finite, then this is equivalent to existence of a BSCC α and a finite path p that starts in s and ends in some state of α, while $t \models \varphi_1 \wedge \neg\varphi_2$ whenever $t \in \alpha$ or p goes through t.

Occasionally, we will show diagrams that represent Markov chains. This is done in the obvious way: circles represent states and edges between them are labeled with their transition probabilities. Edges labeled with 0 are tacitly left out. Circles are labeled with states they represent or formulae they satisfy - which is which should be clear from the context or will be explicitly stated. Figure 2.1 shows an example of a simple Markov chain. We may also sketch parts of finite or infinite Markov chains with a suggestive use of ellipses. Figure 2.2 shows an example of an infinite Markov chain.

When showing a model of some PCTL formula φ, we use a double circle to highlight a state that satisfies φ. Figure 2.3 shows a simple model of $wait \, \mathrm{U}^{\geq 0.8} \, ok$.

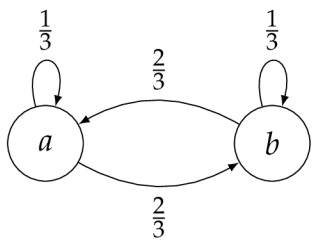

Figure 2.1: A simple finite Markov chain with two states, one of which satisfies a while the other one satisfies b.

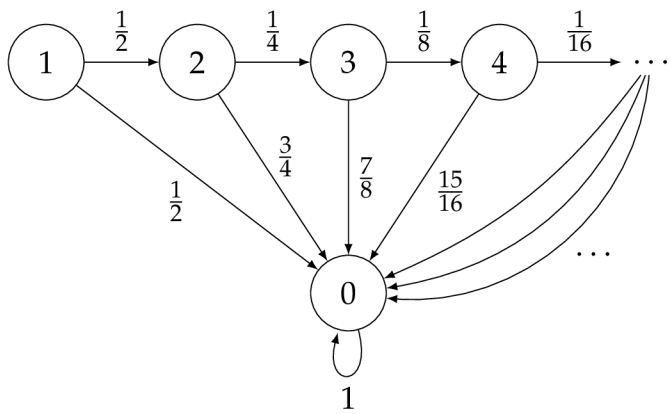

Figure 2.2: An infinite Markov chain with states labeled by natural numbers.

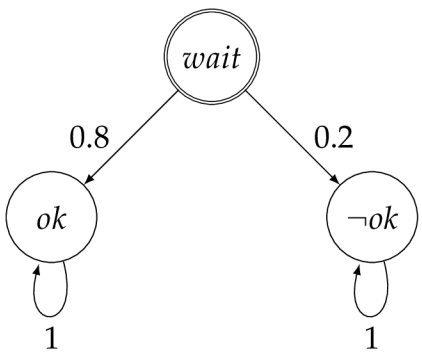

Figure 2.3: A model of $wait\ \mathrm{U}^{\geq 0.8}\ ok$.

2.6 Bounded until formulae

In chapters 4 and 5 we discuss fragments of PCTL that contain *bounded until* formulae. A *bounded until* path formula $\varphi_1 \,\mathrm{U}_n\, \varphi_2$ is constructed by adding a step bound $n \in \mathbb{N}$ to a standard until path formula $\varphi_1 \,\mathrm{U}\, \varphi_2$. A run satisfies $\varphi_1 \,\mathrm{U}_n\, \varphi_2$ if it satisfies $\varphi_1 \,\mathrm{U}\, \varphi_2$ within its first $n+1$ states, which corresponds to n steps from its first state. From a formal perspective, extending PCTL with bounded until formulae amounts to adding another option in the abstract syntax definition of path formulae, so now they would be defined as follows:

$$\Phi ::= \mathrm{X}\,\varphi \mid \varphi \,\mathrm{U}\, \varphi \mid \varphi \,\mathrm{U}_n\, \varphi$$

where $n \in N$.

The semantics of bounded until formulae is defined in the expected way:

Definition 2.6.1 ($\varphi_1 \,\mathrm{U}_n\, \varphi_2$ semantics)**.** Let AP be a countably infinite set and $M = (S, P, v)$ a Markov chain. For a bounded until path formula $\varphi_1 \,\mathrm{U}_n\, \varphi_2$ and a run π, define

$$\begin{aligned} M, \pi \models^v \varphi_1 \,\mathrm{U}_n\, \varphi_2 \text{ iff } &\exists i \in \mathbb{N} \text{ such that } i \leq n \text{ and } \pi(i) \models \varphi_2 \\ &\text{and } \pi(k) \models \varphi_1 \text{ for all } k \in \mathbb{N} \text{ such that } k < i. \end{aligned}$$

Convention. We seamlessly extend all notation and conventions we introduced for PCTL formulae to also apply to bounded until formulae. For example, $\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2$ refers to $P(\varphi_1 \,\mathrm{U}_n\, \varphi_2) \oplus q$, and $\mathrm{F}_n\, \varphi$ refers to $\top \,\mathrm{U}_n\, \varphi$.

We frequently use the following standard lemma:

Lemma 2.6.1. *Let $M = (S, P, v)$ be a Markov chain. Let $s \in S$ be a state and $\varphi_1 \,\mathrm{U}_n\, \varphi_2$ a path formula. Then the following holds:*

(i) If $s \models \varphi_2$, then $\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) = 1$.

(ii) If $s \not\models \varphi_2$ and also at least one of $s \not\models \varphi_1$ and $n = 0$ holds, then $\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) = 0$.

(iii) Otherwise $s \not\models \varphi_2$ and $s \models \varphi_1$ and $n > 0$, and we have

$$\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) = \sum_{t \in s^\bullet} P(s,t) \cdot \mathbb{P}(t \models \varphi_1 \,\mathrm{U}_{n-1}\, \varphi_2).$$

Given that our use of Lemma 2.6.1 is extensive, we include its full proof. This proof takes up the remainder of this section (with the exception of one corollary at the end), and can be safely skipped. Cases (i) and (ii) of Lemma 2.6.1 are easily seen to be true, the main idea for case (iii) is that there is a certain correspondence between runs satisfying $\varphi_1 \mathrm{U}_n \varphi_2$ that start in s and runs satisfying $\varphi_1 \mathrm{U}_{n-1} \varphi_2$ that start in successors of s.

Before the proof itself, we will show one simple lemma, which requires the following definition:

Definition 2.6.2 $(p \dashv \varphi_1 \mathrm{U}_n \varphi_2)$**.** Let $M = (S, P, v)$ be a Markov chain. Suppose $\varphi_1 \mathrm{U}_n \varphi_2$ is a path formula and $p \in S^*$ is a finite sequence of states. We say $p \dashv \varphi_1 \mathrm{U}_n \varphi_2$ if and only if

- p is a finite non-empty path,
- $len(p) - 1 \leq n$,
- $p(len(p) - 1) \models \varphi_2$,
- $p(k) \models \varphi_1 \wedge \neg\varphi_2$ whenever $k \in \mathbb{N}$ and $k < len(p) - 1$.

Intuitively, $p \dashv \varphi_1 \mathrm{U}_n \varphi_2$ expresses that p is a minimal "witness" for a cylinder of runs that satisfy $\varphi_1 \mathrm{U}_n \varphi_2$.

Lemma 2.6.2. *Let $M = (S, P, v)$ be a Markov chain. Suppose $\varphi_1 \mathrm{U}_n \varphi_2$ is a path formula and $s \in S$ a state. Then the family of basic cylinders $\{Run(sp) \mid p \in S^*$ such that $sp \dashv \varphi_1 \mathrm{U}_n \varphi_2\}$ partitions $\{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$.*

Proof. Define $R = \{Run(sp) \mid p \in S^*$ such that $sp \dashv \varphi_1 \mathrm{U}_n \varphi_2\}$. By definition, for every run π it holds that $\pi \in \{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$ if and only if $\pi \in Run(s)$ and $\pi \models \varphi_1 \mathrm{U}_n \varphi_2$, which is equivalent to π having a prefix sp such that $sp \dashv \varphi_1 \mathrm{U}_n \varphi_2$. Then $\bigcup R = \{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$.

Suppose there exists π such that $\pi \in Run(sp)$ and $\pi \in Run(sp')$ for some $p, p' \in S^*$ such that $sp \dashv \varphi_1 \mathrm{U}_n \varphi_2$ and $sp' \dashv \varphi_1 \mathrm{U}_n \varphi_2$, but at the same time $p \neq p'$. Since both sp and sp' are prefixes of π but are not equal to each other, one must be prefix of the other and they must have different lengths. Without loss of generality, assume sp is a prefix of sp'. Then $sp(len(sp) - 1) \models \varphi_2$, which contradicts $sp' \dashv \varphi_1 \mathrm{U}_n \varphi_2$. $\square$

We are now ready to prove Lemma 2.6.1:

Proof of Lemma 2.6.1. Cases (i) and (ii) follow immediately. Suppose case (iii) applies. For every $t \in S$, define

$$E_t = \bigcup_{\substack{p \in S^* \\ stp \dashv \varphi_1 \mathrm{U}_n \varphi_2}} Run(stp).$$

First, we will show that the family of sets of runs $\{E_t \mid t \in s^\bullet\}$ is a partition of the set $\{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$. Suppose $\pi \in \bigcup\{E_t \mid t \in s^\bullet\}$. Then $\pi \in E_t$ for some $t \in s^\bullet$, which means $\pi \in Run(stp)$ for some $p \in S^*$ such that $stp \dashv \varphi_1 \mathrm{U}_n \varphi_2$. Because stp is a prefix of π, it must satisfy $\varphi_1 \mathrm{U}_n \varphi_2$ and simultaneously be in $Run(s)$, and so $\pi \in \{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$. Suppose $\pi \in \{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$. Clearly π starts with s, and there exists some $p' \in S^*$ such that sp' is a prefix of π and $sp' \dashv \varphi_1 \mathrm{U}_n \varphi_2$. Here we use that $s \not\models \varphi_2$, which implies $s \not\dashv \varphi_1 \mathrm{U}_n \varphi_2$, and so $p' \neq \epsilon$. It follows that sp' must be equal to some stp such that $t \in S$ and $p \in S^*$. Then $\pi \in Run(stp)$ for some $t \in S$ and $p \in S^*$ such that $stp \dashv \varphi_1 \mathrm{U}_n \varphi_2$, and we can conclude $\pi \in E_t$. Moreover $t \in s^\bullet$, because $\pi(0) = s$ and $\pi(1) = t$ and π is a run. Then $\pi \in E_t$ for some $t \in s^\bullet$, and so $\pi \in \bigcup\{E_t \mid t \in s^\bullet\}$. Suppose π is a run such that $\pi \in E_t$ and $\pi \in E_{t'}$ for some distinct $t, t' \in s^\bullet$. Then $\pi(1) = t$ and $\pi(1) = t'$, but that contradicts $t \neq t'$.

We have shown that $\{E_t \mid t \in s^\bullet\}$ is a partition of $\{s \models \varphi_1 \mathrm{U}_n \varphi_2\}$, so we can use countable additivity of measures to get

$$\mathbb{P}(s \models \varphi_1 \mathrm{U}_n \varphi_2) = \mathbb{P}(\bigcup_{t \in s^\bullet} E_t) = \sum_{t \in s^\bullet} \mathbb{P}(E_t).$$

What remains is to show $\mathbb{P}(E_t) = P(s,t) \cdot \mathbb{P}(t \models \varphi_1 \mathrm{U}_{n-1} \varphi_2)$ for every $t \in s^\bullet$. Suppose $t \in s^\bullet$, and for convenience define

$$R = \{Run(stp) \mid p \in S^* \text{ and } stp \dashv \varphi_1 \mathrm{U}_n \varphi_2\}.$$

We know that $E_t = \bigcup R$ by its definition, but it also holds that R partitions E_t. This is done by the same argument as in the proof of Lemma 2.6.2, where a run π being in distinct $Run(stp), Run(stp') \in R$ implies at least one of $stp \dashv \varphi_1 \mathrm{U}_n \varphi_2$ and $stp' \dashv \varphi_1 \mathrm{U}_n \varphi_2$ does not hold. Then we can once again use countable additivity of probability measures and get that $\mathbb{P}(E_t)$ must be equal to

$$\sum_{\substack{p \in S^* \\ stp \dashv \varphi_1 \mathrm{U}_n \varphi_2}} \mathbb{P}(Run(stp)),$$

which in turn is clearly equal to

$$\sum_{\substack{p \in S^* \\ stp \dashv \varphi_1 \, \mathrm{U}_n \, \varphi_2}} P(s,t) \cdot \mathbb{P}(Run(tp)).$$

Furthermore, the term $P(s,t)$ is a constant, so we can factor it out of the sum and get

$$P(s,t) \cdot \sum_{\substack{p \in S^* \\ stp \dashv \varphi_1 \, \mathrm{U}_n \, \varphi_2}} \mathbb{P}(Run(tp)).$$

For every $p \in S^*$, we can easily show that $stp \dashv \varphi_1 \, \mathrm{U}_n \, \varphi_2$ is in fact equivalent to $tp \dashv \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2$. Suppose we have $stp \dashv \varphi_1 \, \mathrm{U}_n \, \varphi_2$. Since s cannot be the last state in stp because of t, we can simply "remove" it and still get $tp \dashv \varphi_1 \, \mathrm{U}_n \, \varphi_2$. Since tp contains one fewer state, we also have $tp \dashv \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2$. Now suppose $tp \dashv \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2$. Then "adding" s in front of tp causes no problems because $t \in s^\bullet$ and $s \models \varphi_1 \wedge \neg\varphi_2$, provided we also increase the step bound from $n-1$ to n. This allows us to change the condition in the subscript of the sum in our expression and get

$$P(s,t) \cdot \sum_{\substack{p \in S^* \\ tp \dashv \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2}} \mathbb{P}(Run(tp)).$$

Now consider $\mathbb{P}(t \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2)$. By applying Lemma 2.6.2 and then countable additivity of probability measures we get

$$\begin{aligned} \mathbb{P}(t \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2) &= \mathbb{P}\Big(\bigcup_{\substack{p \in S^* \\ tp \dashv \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2}} Run(tp) \Big) \\ &= \sum_{\substack{p \in S^* \\ tp \dashv \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2}} \mathbb{P}(Run(tp)). \end{aligned}$$

In conclusion,

$$\mathbb{P}(s \models \varphi_1 \, \mathrm{U}_n \, \varphi_2) = \sum_{t \in s^\bullet} \mathbb{P}(E_t) = \sum_{t \in s^\bullet} P(s,t) \cdot \mathbb{P}(t \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

□

Lemma 2.6.1 has a trivial corollary shown below, for which we include a proof because it introduces the type of an argument we will use frequently in Chapter 5.

Corollary 2.6.3. *Let $M = (S, P, v)$ be a Markov chain, $s \in S$ a state, and $\mathrm{F}_1\, \varphi$ a path formula. Suppose $s \not\models \varphi$. Then*

$$\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = \sum_{t \in s^\bullet_\varphi} P(s,t).$$

Proof. By the definition of F we have $\mathrm{F}_1\, \varphi \equiv \top\, \mathrm{U}_1\, \varphi$. Since $s \not\models \varphi$, and simultaneously both $s \models \top$ and $1 > 0$, the third case of Lemma 2.6.1 applies:

$$\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = \sum_{t \in s^\bullet} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_0\, \varphi).$$

The set $s^\bullet$ can be partitioned into $s^\bullet_\varphi$ and $s^\bullet_{\neg\varphi}$, which means

$$\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = \sum_{t \in s^\bullet_\varphi} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_0\, \varphi) + \sum_{t \in s^\bullet_{\neg\varphi}} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_0\, \varphi).$$

For every $t \in s^\bullet_\varphi$ we have $\mathbb{P}(t \models \mathrm{F}_0\, \varphi) = 1$ by the first case of Lemma 2.6.1, and for every $t \in s^\bullet_{\neg\varphi}$ we have $\mathbb{P}(t \models \mathrm{F}_0\, \varphi) = 0$ by the second case of Lemma 2.6.1. $\square$

3 The satisfiability problem for probabilistic temporal logics

In this chapter, we formally introduce the satisfiability problem for PCTL and similar logics. We also present a high-level overview of related work, with added focus on PCTL fragments that are relevant to this thesis.

Before talking about any specific logic, we introduce the relevant problems and properties in general for some probabilistic logic $\mathcal{L}$ given by context:

Definition 3.0.1. The (*finite*) *satisfiability problem* is to decide whether a given formula φ has a (finite) model, that is, if there exists a (finite) Markov chain (S, P, v) and $s \in S$ such that $s \models \varphi$.

We often shorten the *satisfiability problem* to just *satisfiability*. Formulae that have a model are called *satisfiable*, those that have a finite model are called *finitely satisfiable*. When we say *general satisfiability* we simply mean satisfiability, but this is often used in juxtaposition to finite satisfiability.

Definition 3.0.2. The (*finite*) *model checking problem* is to decide whether a given (finite) Markov chain is a model of a given formula.

Size of a formula φ is denoted by $|\varphi|$. We do not care about the exact definition of $|\varphi|$, but always assume that there exists a computable function which outputs $|\varphi|$ for every input formula φ.

When we say X *is bounded by* Y, we mean that there exists a computable function that takes Ys as inputs and outputs upper bounds for Xs. When Y is clear from context, we often simply say X *is bounded*.

Definition 3.0.3. The *small model property* holds if for every satisfiable formula φ, there exists a finite model of size bounded by $|\varphi|$.

Definition 3.0.4. The *bounded satisfiability problem* is to decide whether for a given formula φ and a bound $n \in \mathbb{N}$, there exists a Markov chain (S, P, v) of size at most n such that $s \models \varphi$ for some $s \in S$.

3.1 Related non-probabilistic logics

In this section, we briefly review results for related non-probabilistic logics. For CTL, the satisfiability problem is *EXPTIME*-complete [12], the model checking problem can be solved in $O(|M|\cdot|\varphi|)$ [13] where $|M|$ is the size of the model, and it possesses the small model property [12]. For the modal μ-calculus, the satisfiability problem is *EXPTIME*-complete [14, 15], the model checking problem is in *NP*∩*co-NP* [16], and it also has the small model property [17]. The complexity of deciding satisfiability for fragments of CTL [18] and the modal μ-calculus [19] has also been investigated. For CTL*, the satisfiability problem was shown to be *2EXPTIME*-complete [20], while the model checking problem is in *PSPACE* [21, 13] (similarly to LTL [22]).

3.2 PCTL

In this section, we review results about PCTL. The model checking problem has received treatment for both finite and infinite Markov chains (see e.g. [1, 23, 24, 25, 26, 27]). For finite Markov chains, it is solvable in polynomial time [1]. On the other hand, the satisfiability problem is not even known to be decidable, and has remained open for decades. Examples of satisfiable PCTL formulae and their models are shown on figures throughout this section.

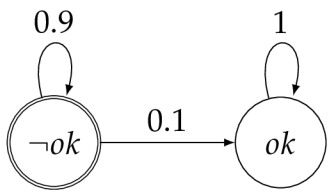

Figure 3.1: A model of $\mathrm{F}^{=1}\, ok$. Runs that do not reach the state that satisfies *ok* have measure 0 because it is a BSCC.

The satisfiability problem for PCTL appears to be fundamentally more complicated than in the case of its predecessor CTL. This is suggested, among other reasons, by the fact that PCTL does not have the small model property. This is in contrast to CTL, and is caused by existence of PCTL formulae that only have infinite models, as argued,

for example, in [28] for a formula of a PCTL fragment called *qualitative PCTL*. We reproduce this argument at the beginning of Chapter 5. Still, the small model property can be used to show decidability of satisfiability for PCTL fragments, because it reduces satisfiability to bounded satisfiability with a computable bound, which is decidable by the following folk knowledge:

Theorem 3.2.1. *The bounded satisfiability problem for PCTL extended with bounded untils can be encoded as an SMT problem in a decidable theory, and is therefore decidable.*

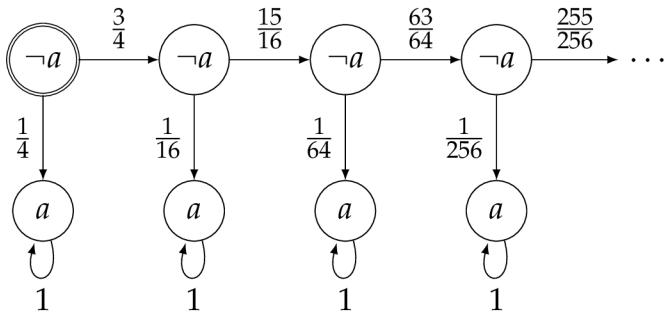

Figure 3.2: A model of $G^{>0}\,\neg a$, since the exponentially increasing transition probabilities of states of the single run that satisfies $G\,\neg a$ are sufficient for it to have a positive measure. Every state satisfies $X^{>0}\,a$, which means $G^{=1}(X^{>0}\,a)$ is satisfied as well. This Markov chain is also a model of $G^{>0}(\neg a \wedge X^{>0}\,a)$, which, unlike the previous two formulae, has no finite models (see Chapter 5).

The small model property, however, is not a necessary condition for satisfiability to be decidable. This is evidenced by qualitative PCTL, where the satisfiability problem can be decided in both its finite and general version, despite the possibility of strictly infinite models (see Subsection 3.3.3).

While most results about PCTL are concerned with its fragments, there are, however, also some results about the whole PCTL. In [8], it is shown that the problem of determining whether a PCTL formula has a model with a branching degree at most k, where $k \geq 2$ is an arbitrary but fixed constant, is highly undecidable. This was done by a reduction from the strategy synthesis problem for $1\frac{1}{2}$-games with

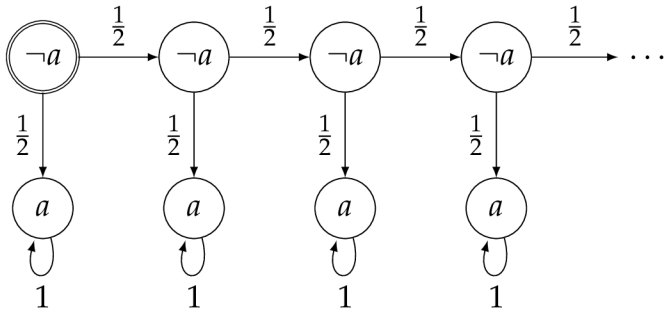

Figure 3.3: A model of $\mathrm{G}^{=0}\,\neg a$, since the only run that satisfies $\mathrm{G}\,\neg a$ has measure 0.

PCTL objectives, which was shown to be highly undecidable in [29]. On the other hand, [8] also shows that whenever a PCTL formula φ has a model, it must also have a model with a branching degree at most $|\varphi| + 2$. This was done by unfolding the model into an infinite tree, and then pruning successors of each state with a standard geometric theorem (shown below as Theorem 3.2.2, proof e.g. in [8]), while also ensuring that satisfaction of no until formula is deferred indefinitely.

Theorem 3.2.2. *Let $n \in \mathbb{N}$ and T be a countable set of vectors in $\mathbb{R}^n$. Suppose $v \in \mathbb{R}^n$ is a convex combination of vectors in T. Then there exists $I \subseteq T$ such that $|I| \leq n + 1$ and v is a convex combination of vectors in I.*

Theorem 3.2.2 is also used to prune branches in [4] and [3], which we talk about in the following section. We will now outline how branch pruning for a single state was carried out in [8].

Suppose $M = (S, P, v)$ is a model of a state formula ψ. Let F be the set of state and path subformulae of ψ. For every $s \in S$, define $\vec{s} \in \mathbb{R}^{|F|}$ to be a vector indexed by F such that for every state formula $\varphi \in F$ we have

$$\vec{s}(\varphi) = \begin{cases} 1 & \text{if } s \models \varphi \\ 0 & \text{otherwise} \end{cases},$$

and for every path formula $\Phi \in F$ we have

$$\vec{s}(\Phi) = \mathbb{P}(s \models \Phi).$$

We will now prune successors for some state $s \in S$. Let N be the subset of F restricted to X path formulae and U path formulae $\varphi_1 \,\mathrm{U}\, \varphi_2$ such

that $s \models \varphi_1 \wedge \neg\varphi_2$. Observe that for each formula in N, the following "local consistency" conditions[1] are satisfied:

$$\vec{s}(X\varphi) = \sum_{t \in s^\bullet} P(s,t) \cdot \vec{t}(\varphi),$$

$$\vec{s}(\varphi_1 \,\mathrm{U}\, \varphi_2) = \sum_{t \in s^\bullet} P(s,t) \cdot \vec{t}(\varphi_1 \,\mathrm{U}\, \varphi_2).$$

We can reinterpret $\vec{s}$ as a vector in $\mathbb{R}^{|N|}$, and for each $t \in s^\bullet$ define $\bar{t} \in \mathbb{R}^{|N|}$ to be a vector indexed by N such that

- $\bar{t}(X\varphi) = \vec{t}(\varphi)$ for $X\varphi \in N$,
- $\bar{t}(\varphi_1 \,\mathrm{U}\, \varphi_2) = \vec{t}(\varphi_1 \,\mathrm{U}\, \varphi_2)$ for $\varphi_1 \,\mathrm{U}\, \varphi_2 \in N$.

The reason why we redefined $\vec{s}$ and defined $\bar{t}$ is that we can now restate the local conditions in the following way:

$$\vec{s} = \sum_{t \in s^\bullet} P(s,t) \cdot \bar{t}$$

This shows that $\vec{s} \in \mathbb{R}^{|N|}$ is a convex combination of vectors from $\{\bar{t} \in \mathbb{R}^{|N|} \mid t \in s^\bullet\}$. Then we can apply Theorem 3.2.2 to show $\vec{s}$ is also a convex combination of vectors corresponding to a bounded subset of its successors. It can be shown that modifying transition probabilities for successors of s in the corresponding way does not break satisfaction of ψ. We emphasize that this only shows part of the process in the context of general PCTL models, but this procedure is, for example, sufficient for treating finite models of bounded PCTL formulae (see [4]).

3.3 Decidable fragments of PCTL

In this section, we review decidability and some structural results for various fragments of PCTL. Generally, three strategies have been used to obtain decidability results for PCTL fragments:

1. Lemma 2.6.1, which we prove, states similar conditions for bounded untils, and the general idea why these conditions hold may be intuitively seen there.

1. Establish the small model property, which gives us a computable upper bound on size of small models for each formula, which reduces satisfiability to bounded satisfiability with a computable bound.

2. Show that every model induces a certain *pseudo-model* (or *pre-model*) and each pseudo-model can be turned back into a model. Then show existence of pseudo-models is decidable, typically by showing that their size is bounded by the size of the formula.

3. Methods based on sound and complete axiomatizations that are decidable.

The first approach is used for example in [4] and [3], and we also use it in Chapter 4. The second approach is used for example in [6] and [8], and we use it in Chapter 6. The third approach has been used for example in [7], although we will not deal with this approach in this thesis.

In the following three subsections, we will focus on specific PCTL fragments that have received attention.

3.3.1 Various fragments based on F, G

A recent work [3] considers the satisfiability problem for various fragments of PCTL based on operators F and G. Finite and general satisfiability is investigated, accompanied by various structural results. The strategy used in [3] to obtain decidability results is to reduce satisfiability to bounded satisfiability. For finite models, this is done primarily by exploiting interaction of operators F and G with BSCCs to build models of bounded size. For infinite models, the situation may be more complicated. Sometimes, it is shown that every model contains states that in some sense "simulate" BSCCs, and bounded models can be constructed from their (not necessarily direct) successors. Another strategy used is to show that every model can be transformed into a tree (with back edges) of bounded height. Branching degree is then bounded with use of the branch-pruning strategy we outlined earlier, which together gives a model of bounded size.

In order to summarize results from [3], we first need to introduce the notation it uses to name fragments. Fragments are named by the

constructs they allow in their abstract syntax definition, while $a, \neg a, \wedge$ are always included and left out for brevity. When denoting names of fragments, we assume $\oplus \in \{\geq, >\}$ and $q \in [0,1] \cap \mathbb{Q}$, and also that 1 stands for ≥ 1 and q stands for $\oplus q$. We use $\mathrm{G}^x(list)$ to denote fragment of $list$ where the topmost operator is G^x. Finally $\mathrm{F}^{q/1}$ denotes F^q with the restriction that inside G only $q = 1$ can be used.

We will now briefly summarize results in [3]. Finite satisfiability is decidable for $\mathrm{G}^q(\mathrm{F}^q, \mathrm{G}^q, \vee)$, and both general and finite satisfiability are decidable for:

- $\mathrm{G}^1(\mathrm{F}^q, \mathrm{G}^1)$
- $\mathrm{G}^1(\mathrm{F}^q, \mathrm{G}^1, \vee)$
- $\mathrm{F}^q, \mathrm{G}^1$
- $\mathrm{F}^{q/1}, \mathrm{G}^1, \vee$

To give a flavor of the type of arguments employed, we outline a decidability proof for one of the simpler fragments treated:

Definition 3.3.1 ($sub(\varphi)$)**.** For the duration of this subsection, redefine $sub(\varphi)$ to not descend through negations, e.g., $\neg a \in sub(\varphi)$ does not imply $a \in sub(\varphi)$. For a formal definition, see Definition 3 in [3].

Convention. We follow [3] and impose *negation normal form* and *lower-bound-comparison normal form* on all formulae we deal with in this subsection.

The following theorem adapts Theorem 1 and Lemma 3 from Section 3.1 in [3], and demonstrates construction of bounded-size models from BSCCs:

Theorem 3.3.1. *Finite satisfiability is decidable for the fragment of formulae φ given by*

$$\varphi ::= G^{\oplus q}\psi$$
$$\psi ::= a \mid \neg a \mid \psi \wedge \psi \mid \psi \vee \psi \mid \mathrm{F}^{\oplus q}\,\psi \mid \mathrm{G}^{\oplus q}\,\psi$$

where $q \in [0,1] \cap \mathbb{Q}$, and $\oplus \in \{>, \geq\}$, and $a \in AP$.

Proof. Let $M = (S, P, v)$ be a finite Markov chain and $s \in S$ be a state such that $s \models \mathrm{G}^{\oplus q}\,\psi$. Then $\mathbb{P}(s \models \mathrm{G}\,\psi) > 0$. Then there exists some BSCC α such that all states in α satisfy ψ, for otherwise $\mathbb{P}(s \models \mathrm{G}\,\psi) = 0$ by Lemma 2.5.2. Then α is a model of $\mathrm{G}^{\oplus q}\,\psi$, and in fact every state in α satisfies $\mathrm{G}^{\oplus q}\,\psi$.

Consider formulae of form $\mathrm{F}^{\odot q}\,\psi'$ in $sub(\mathrm{G}^{\oplus q}\,\psi)$. For every such $\mathrm{F}^{\odot q}\,\psi'$, choose one state from α that satisfies ψ', if it exists. If no states were picked, we can choose an arbitrary one. Arrange these states into a circle C. Clearly C has bounded size. Using structural induction, we can prove that $M, t \models \varphi$ implies $C, t \models \varphi$ for every $t \in C$ and $\varphi \in sub(\mathrm{G}^{\oplus q}\,\psi)$ (the opposite direction does not hold for G subformulae). The cases for literals, conjunction, and disjunction are clear. Consider some $\mathrm{F}^{\odot q}\,\psi'$. Suppose some $t \in C$ satisfies $\mathrm{F}^{\odot q}\,\psi'$ in M. Then C contains a state that satisfies ψ' in C by the induction hypothesis and construction of C. Since C is a BSCC, all states must be visited almost surely, which clearly implies $C, t \models \mathrm{F}^{\odot q}\,\psi'$. Consider some $\mathrm{G}^{\odot q}\,\psi'$ and suppose $t \in C$ satisfies it in M. Then all states in α satisfy ψ' in M by Lemma 2.5.2. Since $C \subseteq \alpha$, we have $C, t \models \psi'$ for every $t \in C$ by the induction hypothesis. Then we have $C, t \models \mathrm{G}^{=1}\,\psi'$, which implies $C, t \models \mathrm{G}^{\odot q}\,\psi'$ because $\odot \in \{\geq, >\}$.

Since $\mathrm{G}^{\oplus q}\,\psi \in sub(\mathrm{G}^{\oplus q}\,\psi)$ and $M, t \models \mathrm{G}^{\oplus q}\,\psi$ for every $t \in \alpha$, there must exist a state $t \in C$ such that $C, t \models \mathrm{G}^{\oplus q}\,\psi$, and we are done. $\square$

3.3.2 Bounded PCTL

Bounded PCTL is a fragment of PCTL obtained by imposing a horizon on all until formulae (see [4] for a formal definition, syntax is also shown in Chapter 5), that is, we no longer allow path formulae of form $\varphi\,\mathrm{U}\,\psi$ and instead use formulae of form $\varphi\,\mathrm{U}_n\,\psi$, where $n \in \mathbb{N}$ serves as the step bound. We introduced bounded until formulae in Section 2.6. The satisfiability problem for bounded PCTL, both in its general and finite version, was shown to be decidable in [4]. Moreover, the same paper also provides a *NEXPTIME* algorithm for deciding it, and shows that bounded PCTL has the small model property.

We use bounded PCTL in Chapter 5, so we briefly review the approach used in [4]. Unfolding a model of some PCTL formula ψ into an infinite tree generally gives us another model of ψ. For bounded

PCTL formulae, it suffices to unfold the model into a tree of bounded depth and make all leaves absorbing. More formally, [4] shows that if M is a model of a bounded PCTL formula ψ, then M can be unfolded into a tree with absorbing leaves (i.e. $P(s, s) = 1$ for each leaf s) of depth bounded by $|\psi|$, and this tree is also a model of ψ. Intuitively, this holds because ψ can only "see" the model up to bounded depth, so states that are "deeper" no longer affect whether ψ is satisfied, and can be arbitrarily modified (for example turned into absorbing leaves). Next, branches of the tree can be pruned using Theorem 3.2.2 to obtain branching degree bounded by $|\psi|$. Together, we get a tree of bounded depth and bounded branching degree, and so the whole tree is of bounded size.

3.3.3 Qualitative PCTL

Qualitative PCTL (*qPCTL*) is a fragment of PCTL obtained by restricting all probabilistic constraints in formulae to either 0 or 1. The syntax of qPCTL is given in Chapter 6. The satisfiability problem for qPCTL and some of its variants has been treated in [5, 6, 7, 8]. We will now briefly recall those results.

Convention. For the duration of this subsection, when we say *bounded Markov chain* or *bounded model*, we mean a Markov chain such that there exists a positive constant α and every non-zero transition probability is at least α.

In [7], sound and complete proof systems were developed for qPCTL* interpreted over finite models, bounded models, and general models (i.e. arbitrary Markov chains). The authors mention that all three systems are decidable by a reduction to $S\omega S$ and a result from [30]. It is also mentioned that all three systems are *PSPACE*-hard by results in [31], and [6] shows they are even *EXPTIME*-hard. Furthermore, [6] shows that qPCTL* satisfiability is in *2-NEXPTIME* for finite, bounded, and general models.

In [5], decision procedures based on the tableaux method and complete axiomatizations were given for qPCTL interpreted over finite models and over bounded models (the paper refers to these logics as PTL_f and PTL_b, respectively). The authors also show that the satis-

fiability problem for both logics is in *EXPTIME*, and that there exist satisfiable PTL_b formulae with no finite models.

In [8], it is established that general and finite satisfiability for qPCTL are both *EXPTIME*-complete. Decidability of general and finite satisfiability is shown as well, although this also follows from the results for qPCTL*. The methods used in [8] form the basis of our decidability arguments in Chapter 4 and Chapter 6, so we show them in a bit more detail. Specifically, [8] uses the method of filtration through Fischer-Ladner closure (originally developed in [15]) adapted for qPCTL. A model M of ψ then induces a certain "quotient" (or "factor") structure of bounded size called *pseudo-model*, which however is not a Markov chain as it lacks transition probabilities and only captures "discrete" information, i.e., it has a transition relation that says whether or not a state s can transition into a state t. We want to turn each pseudo-model back into a model to show that their existence is equivalent to satisfiability. The main issue preventing us from doing so is that satisfaction of formulae of form $\neg(\varphi_1 \mathrm{U}^{=1} \varphi_2)$ is not guaranteed by mere existence of runs that do not satisfy $\varphi_1 \mathrm{U} \varphi_2$, because such runs can still have measure 0. To get around that, two ideas are used:

1. A suitable "witness" for satisfaction of each formula of form $\neg(\varphi_1 \mathrm{U}^{=1} \varphi_2)$ in the closure is identified and shown to exist[2] whenever it is needed. These witnesses are added and connected to the pseudo-model in a consistent way. If no witnesses were employed, then the pseudo-model is called *simple*.

2. If witnesses were employed, some transitions in the modified pseudo-model are "marked". A mark specifies how probability of this transition should change when the pseudo-model is unfolded into an infinite tree. Specifically, a state with distance d from the root of the tree will (collectively and uniformly) assign $1 - \frac{1}{4^{d+1}}$ to marked transitions, and the rest uniformly to unmarked ones. If all or none of the transitions are marked, uniform probability distribution is used.

2. Existence of needed witnesses is technically defined as a necessary condition for being a pseudo-model, so it was actually shown in the proof that the induced structures are pseudo-models.

If the pseudo-model is simple, we can simply endow each state with a uniform probability distribution on its successors, and the resulting Markov chain turns out to be a finite model of ψ. If not, then we unfold the modified pseudo-model in the way described. It is then shown that the probability of remaining inside a witness upon reaching it is strictly greater than 0, which ensures satisfaction of the corresponding $\neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$. This infinite tree[3] is then a model of ψ.

Simple pseudo-models give rise to finite models, and it turns out that finite Markov chains necessarily induce simple pseudo-models. Then (finite) satisfiability is equivalent to existence of (simple) pseudo-models. Finally, an algorithm is presented which decides existence of a pseudo-model (both simple and general) for ψ in time exponential in $|\psi|$. This is based on the fact that the size of the pseudo-model, if it exists, is exponential in $|\psi|$.

3.4 Further related work

The strategy synthesis problem for Markov decision processes with branching-time objectives has been examined in [32, 29, 28, 33, 34]. An extension of the modal μ-calculus called PμTL that uses a probabilistic version of X has also been studied [35, 4]. It is shown that PμTL is orthogonal to PCTL in expressive power, but is more expressive over finite Markov chains when both logics are restricted to their qualitative fragments [35]. The mentioned paper also shows that the model checking problem for PμTL can be solved by adapting classical algorithms for μTL, and therefore is in *UP*$\cap$*co-UP*. Finally, it also gives a solution to the satisfiability problem for PμTL, and shows that it is in *2EXPTIME* by a reduction to parity games. The small model property for PμTL is shown to hold in [4]. Other probabilistic logics with recursion have also been investigated [36, 37, 38].

3. Technically, it is a family of disjoint infinite trees, but nothing prevents us from taking the one tree rooted in a state that satisfies ψ.

4 Qualitative PCTL extended with bounded until formulae

In this chapter, we find two (closely related) criteria that would imply decidability of the finite satisfiability problem for a fragment of PCTL called *qualitative PCTL with bounded untils*. Qualitative PCTL with bounded untils (further referred to as qPCTL+BU) extends qualitative PCTL with quantitative bounded until formulae, that is, formulae of form $P(\varphi_1 \mathrm{U}_n \varphi_2) \oplus q$ where $n \in \mathbb{N}$ and q is no longer constrained to either 0 or 1. We can view qPCTL+BU as qualitative PCTL where we add quantitative untils, but restrict them to their bounded form. Then our motivation for considering qPCTL+BU is that solving the satisfiability problem for it may perhaps provide some useful information on how to solve the satisfiability problem in the general case of PCTL.

The main goal of this chapter is to state those two criteria and prove that they imply decidability of finite satisfiability. We emphasize that we have not succeeded in proving that either of these criteria holds for qPCTL+BU, so they only provide a possible direction for further research. The methods developed in this chapter are based on results from [8], which uses the method of filtration through Fischer-Ladner closure [15] adapted for qPCTL. Roughly speaking, the criteria we formulate say that, for a finite model M of φ, in certain circumstances, we can extend the closure from [8] and use transition probabilities derived from M to get a finite model of φ of size bounded by $|\varphi|$. Decidability then follows from decidability of bounded satisfiability for PCTL.

We start with formal definition of the syntax of qPCTL+BU:

Definition 4.0.1 (qPCTL+BU syntax)**.** Let AP be a countably infinite set. Then *qPCTL+BU state formulae* are defined as follows:

$$\varphi ::= a \mid \varphi \wedge \varphi \mid \neg\varphi \mid P(\Phi) \oplus b \mid P(\Psi) \oplus q$$

where $a \in AP$, $\oplus \in \{<, \leq, \geq, >, =, \neq\}$, $b \in \{0, 1\}$, and $q \in [0, 1] \cap \mathbb{Q}$. Both Φ and Ψ refer to *qPCTL+BU path formulae*, which are defined as follows:

$$\Phi ::= \mathrm{X}\,\varphi \mid \varphi\,\mathrm{U}\,\varphi$$

$$\Psi ::= \varphi \, \mathrm{U}_n \, \varphi$$

where $n \in \mathbb{N}$.

Semantics of qPCTL+BU is defined in the usual way.

Convention. For the duration of this chapter, state and path formulae always refer to state and path formulae of qPCTL+BU.

We recall an observation from [8], which allows us to establish the following convention:

Convention. Suppose we have a formula of form $\mathrm{X}^{\oplus q} \, \varphi$ or $\varphi \, \mathrm{U}^{\oplus q} \, \psi$ where q is either 0 or 1. Then every formula of this form must be semantically equivalent to a tautology, a contradiction, or one of $\mathrm{X}^{>0} \, \varphi$, $\mathrm{X}^{=1} \, \varphi$, $\varphi \, \mathrm{U}^{>0} \, \psi$, $\varphi \, \mathrm{U}^{=1} \, \psi$, and their negations. Since qPCTL contains negations, we can assume without loss of generality that formulae of form $\mathrm{X}^{\oplus q} \, \varphi$ or $\varphi \, \mathrm{U}^{\oplus q} \, \psi$ have $\oplus q$ which is either >0 or $=1$.

We now define some auxiliary concepts that will be used to deal with bounded until formulae.

Definition 4.0.2 ($lower(\Phi)$)**.** Let Φ be a path formula of form $\varphi_1 \, \mathrm{U}_n \, \varphi_2$. Define $lower(\Phi)$ to be the set $\{\varphi_1 \, \mathrm{U}_k \, \varphi_2 \mid k \in \mathbb{N} \text{ such that } 0 \leq k \leq n\}$.

The set $bsub(\psi)$ defined below denotes the set of path subformulae of ψ that are bounded untils, but since we use it in a proof we give a formal definition:

Definition 4.0.3 ($bsub(\psi)$)**.** Let ψ be a state formula. The set $bsub(\psi)$ is defined recursively, depending on the form of ψ:

- $bsub(a) = \emptyset$,
- $bsub(\neg\varphi) = bsub(\varphi)$,
- $bsub(\varphi_1 \wedge \varphi_2) = bsub(\varphi_1) \cup bsub(\varphi_2)$,
- $bsub(\mathrm{X}^{\oplus q} \, \varphi) = bsub(\varphi)$,
- $bsub(\varphi_1 \, \mathrm{U}^{\oplus q} \, \varphi_2) = bsub(\varphi_1) \cup bsub(\varphi_2)$,
- $bsub(\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2) = \{\varphi_1 \, \mathrm{U}_n \, \varphi_2\} \cup bsub(\varphi_1) \cup bsub(\varphi_2)$.

Definition 4.0.4 ($lbsub(\psi)$)**.** Let ψ be a state formula. Define

$$lbsub(\psi) = \bigcup_{\Psi \in bsub(\psi)} lower(\Psi).$$

We will now define the closure set we will use for filtration. It is a straightforward extension of the one used in [8], but now contains two additional rules (rule 8 and 9) and also depends on the Markov chain and the valuation used.

Definition 4.0.5 ($eCl(\psi, M, v)$)**.** Let $M = (S, P)$ be a Markov chain, v a valuation, and ψ a state formula. Then $eCl(\psi, M, v)$ denotes the extended closure of ψ with respect to M and v, and is defined to be the least set C such that $\psi \in C$ and the following conditions are satisfied:

1. $\varphi \in C$ implies $\widehat{\varphi} \in C$,
2. $\varphi_1 \wedge \varphi_2 \in C$ implies $\varphi_1, \varphi_2 \in C$,
3. $\mathrm{X}^{>0}\, \varphi \in C$ implies $\varphi \in C$,
4. $\mathrm{X}^{=1}\, \varphi \in C$ implies $\varphi \in C$,
5. $\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2 \in C$ implies $\varphi_1, \varphi_2, \mathrm{X}^{>0}(\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2) \in C$,
6. $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \in C$ implies $\varphi_1, \varphi_2, \mathrm{X}^{=1}(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in C$,
7. $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \in C$ implies $\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2 \in C$,
8. $\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2 \in C$ implies $\varphi_1, \varphi_2 \in C$,
9. $P(\varphi_1 \,\mathrm{U}_n\, \varphi_2) = q' \in C$ whenever $\varphi_1 \,\mathrm{U}_n\, \varphi_2 \in lbsub(\psi)$ and also $q' \in \{\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) \mid s \in S\}$.

Lemma 4.0.1. *Let M be a Markov chain, v a valuation, and ψ a state formula. Then $\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2 \in eCl(\psi, M, v)$ implies $\varphi_1 \,\mathrm{U}_n\, \varphi_2 \in lbsub(\psi)$.*

Proof. We will show that for every $\varphi \in eCl(\psi, M, v)$ it holds that $lbsub(\varphi) \subseteq lbsub(\psi)$. This lemma is then a simple consequence: if $\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2 \in eCl(\psi, M, v)$ then we have $lbsub(\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2) \subseteq lbsub(\psi)$, but clearly $\varphi_1 \,\mathrm{U}_n\, \varphi_2 \in lbsub(\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2)$ so $\varphi_1 \,\mathrm{U}_n\, \varphi_2 \in lbsub(\psi)$.

To prove this, we first need to note that from a formal perspective, the set $eCl(\psi, M, v)$ corresponds to a set $\cup_{i \in \mathbb{N}} C_i$, where $C_0 = \{\psi\}$ and each C_{i+1} is the union of C_i and the set of all formulae added by applying rules 1-9 on C_i. Clearly if every C_i satisfies the condition, then so does their union. We will use induction on i:

Base case Suppose $i = 0$. We know that $C_0 = \{\psi\}$ and obviously $lbsub(\psi) \subseteq lbsub(\psi)$.

Induction step Suppose $i > 0$ and the statement holds for all $k \in \mathbb{N}$ such that $k < i$. Given the induction hypothesis, we only need to show that formulae added by applying rules 1-9 on C_{i-1} do not break the condition. This is obvious for formulae added by rules 1-8, since whenever we added φ_2 because of some φ_1 from C_{i-1}, it must be that $bsub(\varphi_2) \subseteq bsub(\varphi_1)$, which means $lbsub(\varphi_2) \subseteq lbsub(\varphi_1)$, and we know $lbsub(\varphi_1) \subseteq lbsub(\psi)$ by the induction hypothesis.

Suppose $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in C_i$ because of rule 9. Then $\varphi_1 \, \mathrm{U}_n \, \varphi_2 \in lpsub(\psi)$. It follows that $\varphi_1 \, \mathrm{U}_n \, \varphi_2 \in lower(\Phi)$ for some $\Phi \in bsub(\psi)$. Clearly $\Phi \equiv \varphi_1 \, \mathrm{U}_k \, \varphi_2$ for some $k \in \mathbb{N}$ such that $k \geq n$. It is easy to see that if $\varphi_1 \, \mathrm{U}_k \, \varphi_2 \in bsub(\psi)$, then $bsub(\varphi_1) \subseteq bsub(\psi)$ and $bsub(\varphi_2) \subseteq bsub(\psi)$. An obvious consequence of that is that $bsub(\varphi_1) \cup bsub(\varphi_2) \subseteq bsub(\psi)$. We know that $bsub(\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2)$ is equal to $\{\varphi_1 \, \mathrm{U}_n \, \varphi_2\} \cup bsub(\varphi_1) \cup bsub(\varphi_2)$. It follows that $bsub(\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2) \setminus \{\varphi_1 \, \mathrm{U}_n \, \varphi_2\}$ is a subset of $bsub(\psi)$. Then all we need to show is that $lower(\varphi_1 \, \mathrm{U}_n \, \varphi_2) \subseteq lbsub(\psi)$. We know that $\varphi_1 \, \mathrm{U}_n \, \varphi_2 \in lower(\varphi_1 \, \mathrm{U}_k \, \varphi_2)$ and $\varphi_1 \, \mathrm{U}_k \, \varphi_2 \in bsub(\psi)$. Because $k \geq n$ clearly $lower(\varphi_1 \, \mathrm{U}_n \, \varphi_2) \subseteq lower(\varphi_1 \, \mathrm{U}_k \, \varphi_2)$, but also $lower(\varphi_1 \, \mathrm{U}_k \, \varphi_2) \subseteq lbsub(\psi)$ because $\varphi_1 \, \mathrm{U}_k \, \varphi_2 \in bsub(\psi)$, and we are done.

□

Corollary 4.0.2. *Let M be a Markov chain with at least one state, v a valuation, and ψ a state formula. Suppose some $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl(\psi, M, v)$ where $n > 0$. Then $\varphi_1 \, \mathrm{U}_{n-1}^{=q'} \, \varphi_2 \in eCl(\psi, M, v)$ for some $q' \in \mathbb{Q}$.*

Proof. Since $\varphi_1 \, \mathrm{U}_n \, \varphi_2 \in lbsub(\psi)$ by Lemma 4.0.1, it must be that $\varphi_1 \, \mathrm{U}_n \, \varphi_2 \in lower(\Psi)$ for some $\Psi \in bsub(\psi)$. Given that $n > 0$, we

clearly have $\varphi_1 \mathrm{U}_{n-1} \varphi_2 \in lower(\Psi)$, and so $\varphi_1 \mathrm{U}_{n-1} \varphi_2 \in lbsub(\psi)$. Pick an arbitrary $s \in S$, which must exist because M has at least one state. Denote q' the value of $\mathbb{P}(s \models \varphi_1 \mathrm{U}_{n-1} \varphi_2)$. Then by rule 9 in the definition of $eCl(\psi, M, v)$ we have $\varphi_1 \mathrm{U}_n^{=q'} \varphi_2 \in eCl(\psi, M, v)$. □

We are now ready to prove the main theorem of this chapter. This proof is based on ideas in the proof of Theorem 3.7 from [8], while rule 9 in the definiton of $eCl(\psi, M, v)$ and our specific choice of probabilities in the newly constructed model ensure that bounded until formulae are preserved as well.

Theorem 4.0.3. *Let ψ be a state formula. Suppose $M = (S, P)$ is a finite Markov chain, v is a valuation, and $u \in S$ is a state such that $u \models \psi$. Then ψ has a finite model $(\mathfrak{S}, P^*, v^*)$ such that $|\mathfrak{S}| \leq |\mathcal{P}(eCl(\psi, M, v))|$.*

Proof. We will apply the method of filtration through $eCl(\psi, M, v)$. For the duration of this proof, we will use eCl as a shortcut for the whole $eCl(\psi, M, v)$. We will construct a new model $M^* = (\mathfrak{S}, P^*, v^*)$. Define an equivalence relation on S, where s is equivalent to t if and only if the sets $\{\varphi \in eCl \mid M, s \models^v \varphi\}$ and $\{\varphi \in eCl \mid M, t \models^v \varphi\}$ are equal to each other (they satisfy the same formulae from eCl). For every $s \in S$, denote $[s] = \{t \in S \mid t \text{ is equivalent to } s\}$. Now we define $\mathfrak{S} = \{[s] \mid s \in S\}$. Clearly $|\mathfrak{S}|$ is finite and $|\mathfrak{S}| \leq |\mathcal{P}(eCl)|$.

It is easy to see that eCl contains all atomic propositions mentioned in ψ, so all states in a single equivalence class must satisfy the same atomic propositions (when considering those mentioned in ψ). Then we can define v^* in the obvious way: each equivalence class has exactly those atomic propositions from ψ which are satisfied in every one of its states, while those not mentioned in ψ can be arbitrary.

The function P^* is defined as follows:

$$P^*(T, K) = \sum_{t \in T, k \in K} \frac{P(t, k)}{|T|} \text{ for every } T, K \in \mathfrak{S}.$$

Before we show that M^* is a model of ψ, we need to show that M^* is a Markov chain. Clearly $P^*(T, K)$ is always non-negative. We also need to show that for every $T \in \mathfrak{S}$ it holds that

$$\sum_{K \in \mathfrak{S}} P^*(T, K) = 1.$$

By substituting the definition of $P^*(T,K)$, we get

$$\sum_{K \in \mathfrak{S}} P^*(T,K) = \sum_{K \in \mathfrak{S}} \sum_{t \in T, k \in K} \frac{P(t,k)}{|T|}.$$

It is easy to see that the denominator $|T|$ can be factored out of both sums, and so the right side of the previous equality is equal to

$$\frac{1}{|T|} \cdot \sum_{K \in \mathfrak{S}} \sum_{t \in T, k \in K} P(t,k).$$

We can reorder terms in the inner sum by t to get

$$\frac{1}{|T|} \cdot \sum_{K \in \mathfrak{S}} \sum_{t \in T} \sum_{k \in K} P(t,k),$$

which can be further reordered to obtain

$$\frac{1}{|T|} \cdot \sum_{t \in T} \sum_{K \in \mathfrak{S}} \sum_{k \in K} P(t,k).$$

Observe that these sums intuitively express that for each $t \in T$, we go through all equivalence classes $K \in \mathfrak{S}$ and sum all transition probabilities between t and states in K. We know that $\mathfrak{S}$ is a partition of the original set of states S, so we count $P(t,k)$ exactly once for each $k \in S$. In other words, it holds that $\sum_{K \in \mathfrak{S}} \sum_{k \in K} P(t,k)$ is equal to $\sum_{k \in S} P(t,k)$, which must be equal to 1 by the definition of Markov chains. Then

$$\frac{1}{|T|} \cdot \sum_{t \in T} \sum_{K \in \mathfrak{S}} \sum_{k \in K} P(t,k) = \frac{1}{|T|} \cdot \sum_{t \in T} \sum_{k \in S} P(t,k) = \frac{1}{|T|} \cdot \sum_{t \in T} 1 = 1,$$

giving us our initial goal of proving that $\sum_{K \in \mathfrak{S}} P(T,K) = 1$, and so M^* is a Markov chain.

Before we continue with the proof, we will make the following observation: for every $T, K \in \mathfrak{S}$ we have $P^*(T,K) > 0$ if and only if $P(t,k) > 0$ for some $t \in T$ and $k \in K$. This can be inferred directly from the definition of P^*, and will be continuously used throughout this proof.

Now we will prove that M^* is a model of ψ. To do so, we will show that for every $\varphi \in eCl$ and every $s \in S$ we have

$$M, s \models^v \varphi \text{ if and only if } M^*, [s] \models^{v^*} \varphi.$$

This is done by structural induction on formulae. We only show the more interesting cases. The "qualitative" cases essentially follow the proof of Theorem 3.7 from [8], and only differ in minor technical details, such as equivalence classes being represented by sets of states and not formulae:

- $\mathrm{X}^{=1}\,\varphi$. Suppose $M, s \models^v \mathrm{X}^{=1}\,\varphi$ and suppose $P([s],[t]) > 0$ for some $[t] \in \mathfrak{S}$. By our observation, there exists $s' \in [s]$ and $t' \in [t]$ such that $P(s', t') > 0$. Since $[s'] = [s]$ and $\mathrm{X}^{=1}\,\varphi \in eCl$ and $M, s \models^v \mathrm{X}^{=1}\,\varphi$, we know that $M, s' \models^v \mathrm{X}^{=1}\,\varphi$. Then $M, t' \models^v \varphi$. Clearly $\varphi \in eCl$, and so the induction hypothesis implies $M^*, [t] \models^{v^*} \varphi$. Since this holds for every $[t] \in \mathfrak{S}$ such that $P^*([s],[t]) > 0$, we have $M^*, [s] \models^{v^*} \mathrm{X}^{=1}\,\varphi$. The opposite direction can be done by contradiction. Suppose $M^*, [s] \models^{v^*} \mathrm{X}^{=1}\,\varphi$ but not $M, s \models^v \mathrm{X}^{=1}\,\varphi$. Then $M, s \models^v \neg\,\mathrm{X}^{=1}\,\varphi$. Clearly $\neg\,\mathrm{X}^{=1}\,\varphi \in eCl$, so by the case for $\neg\,\mathrm{X}^{=1}\,\varphi$ we get $M^*, [s] \models^{v^*} \neg\,\mathrm{X}^{=1}\,\varphi$, a contradiction. The case for $\neg\,\mathrm{X}^{>0}\,\varphi$ is similar.

- $\varphi_1\,\mathrm{U}^{>0}\,\varphi_2$. Suppose $M, s \models^v \varphi_1\,\mathrm{U}^{>0}\,\varphi_2$. Then there exists a path $s_0 s_1 \ldots s_n$ such that $s_0 = s$ and $M, s_n \models^v \varphi_2$ and for all $i \in \mathbb{N}$ such that $0 \leq i < n$ we have $M, s_i \models^v \varphi_1$. By our observation, there exists a path $[s_0][s_1]\ldots[s_n]$ and also $[s_0] = [s]$. Since $\varphi_1, \varphi_2 \in eCl$, apply the induction hypothesis to get $M^*, [s_n] \models^{v^*} \varphi_2$ and also $M^*, [s_i] \models^{v^*} \varphi_1$ for every $i \in \mathbb{N}$ such that $0 \leq i < n$. It follows that $M^*, [s] \models^{v^*} \varphi_1\,\mathrm{U}^{>0}\,\varphi_2$. The other direction can be done by contradiction. Suppose $M^*, [s] \models^{v^*} \varphi_1\,\mathrm{U}^{>0}\,\varphi_2$ but not $M, s \models^v \varphi_1\,\mathrm{U}^{>0}\,\varphi_2$. Then $M, s \models^v \neg(\varphi_1\,\mathrm{U}^{>0}\,\varphi_2)$ and clearly $\neg(\varphi_1\,\mathrm{U}^{>0}\,\varphi_2) \in eCl$, which by the case for $\neg(\varphi_1\,\mathrm{U}^{>0}\,\varphi_2)$ implies that $M^*, [s] \models^{v^*} \neg(\varphi_1\,\mathrm{U}^{>0}\,\varphi_2)$, a contradiction. The cases for $\mathrm{X}^{>0}\,\varphi$ and $\neg\,\mathrm{X}^{=1}\,\varphi$ are similar.

- $\neg(\varphi_1\,\mathrm{U}^{>0}\,\varphi_2)$. Suppose $M, s \models^v \neg(\varphi_1\,\mathrm{U}^{>0}\,\varphi_2)$. We cannot use the case for $\varphi_1\,\mathrm{U}^{>0}\,\varphi_2$ here, as our reasoning would be circular. If $M^*, [s] \models^{v^*} (\varphi_1\,\mathrm{U}^{>0}\,\varphi_2)$, then there would have to be a path

$[s_0][s_1]\dots[s_n]$ such that $[s] = [s_0]$ and $M^*, [s_n] \models^{v^*} \varphi_2$ and for every $i \in \mathbb{N}$ such that $0 \leq i < n$ we have $M^*, [s_i] \models^{v^*} \varphi_1$. We will show that this situation cannot happen. We use an induction on the length of the path n. Suppose $n = 0$. Then $M^*, [s] \models^{v^*} \varphi_2$, but $\varphi_2 \in eCl$, and so by the "outer" induction hypothesis we have $M, s \models^v \varphi_2$, which necessarily implies $M, s \models^v \varphi_1 \mathrm{U}^{>0} \varphi_2$, a contradiction. For the induction step, we know that $M, s \models^v \neg(\varphi_1 \mathrm{U}^{>0} \varphi_2)$, and suppose there exists such a path $[s][s_1]\dots[s_n]$. Clearly $M, s \models^v \neg\varphi_1 \wedge \neg\varphi_2$ or $M, s \models^v \varphi_1 \wedge \neg\varphi_2$. The first case gives $M^*, [s] \not\models^{v^*} \varphi_1$ by the "outer" induction hypothesis, which is a contradiction with properties of the path. In the second case, it is easy to see that $M, s \models^v \neg \mathrm{X}^{>0}(\varphi_1 \mathrm{U}^{>0} \varphi_2)$ and also $\neg \mathrm{X}^{>0}(\varphi_1 \mathrm{U}^{>0} \varphi_2) \in eCl$. We know that $P^*([s], [s_1]) > 0$, so there exist $s_0 \in [s]$ and $s_1 \in [s_1]$ such that $P(s_0, s_1) > 0$. Since $s_0 \in [s]$, we have that $M, s_0 \models^v \neg \mathrm{X}^{>0}(\varphi_1 \mathrm{U}^{>0} \varphi_2)$ and so $M, s_1 \models^v \neg(\varphi_1 \mathrm{U}^{>0} \varphi_2)$. Now observe that the state s_1 and the path $[s_1][s_2]\dots[s_n]$ contradict the "inner" induction hypothesis.

- $\varphi_1 \mathrm{U}^{=1} \varphi_2$. Suppose $M, s \models^v \varphi_1 \mathrm{U}^{=1} \varphi_2$ but at the same time $M^*, [s] \models^{v^*} \neg(\varphi_1 \mathrm{U}^{=1} \varphi_2)$. Using Lemma 2.5.3, there are two options:

  (i) There exists a path $[s_0][s_1]\dots[s_n]$ such that $[s] = [s_0]$ and $M^*, [s_n] \not\models^{v^*} \varphi_1$ and for all $i \in \mathbb{N}$ such that $0 \leq i \leq n$ we have $M^*, [s_i] \not\models^{v^*} \varphi_2$. Since $\varphi_1, \varphi_2 \in eCl$, we can use the induction hypothesis to get $M, s_n \models^v \neg\varphi_1 \wedge \neg\varphi_2$ for all $s_n \in [s_n]$. Then for each $s_n \in [s_n]$ we have $M, s_n \models^v \neg(\varphi_1 \mathrm{U}^{>0} \varphi_2)$. Observe that for every fixed $0 \leq i \leq n$ we have $M, s_i \models^v \neg\varphi_2$ for all $s_i \in [s_i]$, and by a straightforward induction on j we can show that for all $0 \leq j \leq n$ it holds that for all $s_{n-j} \in [s_{n-j}]$ we have $M, s_{n-j} \models^v \neg(\varphi_1 \mathrm{U}^{=1} \varphi_2)$, which contradicts $M, s \models^v \varphi_1 \mathrm{U}^{=1} \varphi_2$.

  (ii) There is a BSCC α in M^* such that for every $[t] \in \alpha$ we have $M^*, [t] \models^{v^*} \neg\varphi_2$ and a path $[s_0][s_1]\dots[s_n]$ such that $[s] = [s_0]$ and $[s_n] \in \alpha$ and for every $i \in \mathbb{N}$ such that $0 \leq i \leq n$ it holds that $M^*, [s_i] \models^{v^*} \neg\varphi_2$. First, since α is a BSCC observe that for every $[t] \in \alpha$ it must be that $M^*, [t] \models^{v^*} \neg(\varphi_1 \mathrm{U}^{>0} \varphi_2)$, which obviously implies

$M^*, [t] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$. This also applies to $[s_n]$ since it is in α. Similarly as in (i), we can use induction to show that $M, s \models^v \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$.

The other direction follows from the case for $\neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$.

- $\neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$. Suppose $M, s \models^v \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$ but at the same time $M^*, [s] \models^{v^*} \varphi_1 \, \mathrm{U}^{=1} \, \varphi_2$. Using Lemma 2.5.3, there are two options:

  (i) There exists a path $s_0 s_1 \ldots s_n$ such that $s = s_0$ and $M, s_n \models^v \neg\varphi_1$ and $M, s_i \models^v \neg\varphi_2$ for all $i \in \mathbb{N}$ such that $0 \leq i \leq n$. Using our observation and the induction hypothesis, there must exist a path $[s_0][s_1] \ldots [s_n]$ such that $[s] = [s_0]$ and $M^*, [s_n] \models^{v^*} \neg\varphi_1$ and $M^*, [s_i] \models^{v^*} \neg\varphi_2$ for all $i \in \mathbb{N}$ such that $0 \leq i \leq n$. But then $M^*, [s] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$.

  (ii) There exists a BSCC α such that for all $t \in \alpha$ we have $M, t \models^v \neg\varphi_2$ and a path $s_0 s_1 \ldots s_n$ such that $s = s_0$ and $s_n \in \alpha$ and $M, s_i \not\models^v \varphi_2$ for every $i \in \mathbb{N}$ such that $0 \leq i \leq n$. Since for every $t \in \alpha$ we have $M, t \models^v \neg\varphi_2$ and α is a BSCC, it follows that $M, t \models^v \neg(\varphi_1 \, \mathrm{U}^{>0} \, \varphi_2)$ for every $t \in \alpha$. With little work we can show that if $\neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2) \in eCl$, then also $\neg(\varphi_1 \, \mathrm{U}^{>0} \, \varphi_2) \in eCl$. Using the case for $\neg(\varphi_1 \, \mathrm{U}^{>0} \, \varphi_2)$, we now have $M^*, [t] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{>0} \, \varphi_2)$ for every $t \in \alpha$. This also applies to $[s_n]$ since $[s_n] \in \alpha$. Consider the path $[s_0][s_1] \ldots [s_n]$ (note that $[s_0] = [s]$), which must exist by our observation. One can easily see that $\varphi_2 \in eCl$, so the induction hypothesis yields $M^*, [s_i] \not\models^{v^*} \varphi_2$ for each $i \in \mathbb{N}$ such that $0 \leq i \leq n$. Obviously $M^*, [s_n] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{>0} \, \varphi_2)$ means that $M^*, [s_n] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$. By applying a straightforward induction on j, we can show that for each $[s_{n-j}]$ we have $M^*, [s_{n-j}] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$, and consequently $M^*, [s] \models^{v^*} \neg(\varphi_1 \, \mathrm{U}^{=1} \, \varphi_2)$.

What remains is to show that the "quantitative" cases with bounded untils also hold, that is, for every $s \in S$ and every $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl$, it holds that $M, s \models^v \varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2$ if and only if $M^*, [s] \models^{v^*} \varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2$. We

will show the following stronger statement: for every $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl$ and $s \in S$, we have

$$\mathbb{P}(M, s \models^v \varphi_1 \, \mathrm{U}_n \, \varphi_2) = \mathbb{P}(M^*, [s] \models^{v^*} \varphi_1 \, \mathrm{U}_n \, \varphi_2).$$

This is done by induction on the step bound n in $\varphi_1 \, \mathrm{U}_n \, \varphi_2$. We now have two nested inductions. Clearly $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl$ implies $\varphi_1, \varphi_2 \in eCl$. Therefore, whenever we encounter states that do or do not satisfy φ_1 and φ_2, we can always apply the "outer" induction hypothesis. Now we proceed with the induction:

Base case. Let $n = 0$ and $s \in S$ and $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl$. If $M, s \models^v \varphi_2$, then $M^*, [s] \models^{v^*} \varphi_2$ by the "outer" induction hypothesis. Using case (i) of Lemma 2.6.1 we see that both measures are equal to 1. Otherwise we have $M, s \not\models^v \varphi_2$ and $M^*, [s] \not\models^{v^*} \varphi_2$ and also $n = 0$, so the case (ii) of Lemma 2.6.1 applies, which means both measures must be equal to 0.

Induction step. Before we start, we adopt some conventions so that our notation does not get out of hand. First, we primarily use lowercase letters such as t, k to denote states of S instead of the usual s. Then uppercase letters such as T, K are typically used to denote their respective equivalence classes $[t], [k]$. Since it is usually clear whether we are talking about M or M^* and v or v^*, we do not show them whenever possible. With this in mind, we carry on with the induction:

Suppose $n > 0$ and the statement holds for all $i \in \mathbb{N}$ such that $i < n$. Let $s \in S$ and $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl$. Set $T = [s]$. The cases when $s \models \varphi_2$ or $s \models \neg\varphi_2 \wedge \neg\varphi_1$ are done similarly as in the base case by using the "outer" induction hypothesis together with Lemma 2.6.1. Otherwise $s \models \neg\varphi_2 \wedge \varphi_1$. Now apply the "outer" induction hypothesis, and since we know that $n > 0$, we can apply case (iii) of Lemma 2.6.1 to get (recall $[s] = T$):

$$\mathbb{P}([s] \models \varphi_1 \, \mathrm{U}_n \, \varphi_2) = \sum_{K \in T^\bullet} P^*(T, K) \cdot \mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Since every $K \in \mathfrak{S}$ such that $K \notin T^\bullet$ has $P^*(T, K) = 0$, the right side of the previous equality is equal to

$$\sum_{K \in \mathfrak{S}} P^*(T, K) \cdot \mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Now we use the definition of P^* and also factor $\frac{1}{|T|}$ out of both sums and get

$$\frac{1}{|T|} \sum_{K \in \mathfrak{S}} \left(\sum_{t \in T, k \in K} P(t,k) \right) \cdot \mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Since multiplication is distributive over addition, the parentheses around the "inner" sum are not necessary:

$$\frac{1}{|T|} \sum_{K \in \mathfrak{S}} \sum_{t \in T, k \in K} P(t,k) \cdot \mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Similarly to as we have done before, we split the "inner" sum in two and then equivalently reorder them so that the sum ranging over T is the "outermost" one:

$$\frac{1}{|T|} \sum_{t \in T} \sum_{K \in \mathfrak{S}} \sum_{k \in K} P(t,k) \cdot \mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

We know that M has at least one state (the one that satisfies ψ), and also that $\varphi_1 \, \mathrm{U}_n^{\oplus q} \, \varphi_2 \in eCl$ for $n > 0$. Then by Corollary 4.0.2 we have $P(\varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2) = q' \in eCl$ for some $q' \in \mathbb{Q}$. That implies $\mathbb{P}(k \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2) = \mathbb{P}([k] \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2)$ for every $k \in S$ by the "inner" induction hypothesis. Observe that in our expression, whenever we see some $P(t,k) \cdot \mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2)$, it must be that $[k] = K$ because $k \in K$. Then we can replace each $\mathbb{P}(K \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2)$ by the equivalent $\mathbb{P}(k \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2)$ and get

$$\frac{1}{|T|} \sum_{t \in T} \sum_{K \in \mathfrak{S}} \sum_{k \in K} P(t,k) \cdot \mathbb{P}(k \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Since $\mathfrak{S}$ is a partition of S and we go through each $K \in \mathfrak{S}$ and each $k \in K$ exactly once, the two "inner" sums can be equivalently replaced by a sum that ranges over all states in S:

$$\frac{1}{|T|} \sum_{t \in T} \sum_{k \in S} P(t,k) \cdot \mathbb{P}(k \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Clearly for all $k \notin t^\bullet$ we have $P(t,k) = 0$, so the expression above is equivalent to

$$\frac{1}{|T|} \sum_{t \in T} \sum_{k \in t^\bullet} P(t,k) \cdot \mathbb{P}(k \models \varphi_1 \, \mathrm{U}_{n-1} \, \varphi_2).$$

Recall that $T = [s]$ and $s \models \neg\varphi_2 \wedge \varphi_1$ and also $\varphi_1, \varphi_2 \in eCl$. Then for every $t \in T$, we have $t \models \neg\varphi_2 \wedge \varphi_1$ and we can apply Lemma 2.6.1 to get

$$\mathbb{P}(t \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) = \sum_{k \in t^\bullet} P(t,k) \cdot \mathbb{P}(k \models \varphi_1 \,\mathrm{U}_{n-1}\, \varphi_2).$$

Since the variable t ranges strictly over T in our expression, we can use this fact to further reduce it to

$$\frac{1}{|T|} \sum_{t \in T} \mathbb{P}(t \models \varphi_1 \,\mathrm{U}_n\, \varphi_2).$$

What remains is to show that $\mathbb{P}(t \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) = \mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2)$ for every $t \in T$, which would allow us to finally conclude that

$$\mathbb{P}([s] \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) = \ldots = \frac{1}{|T|} \sum_{t \in T} \mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2) =$$
$$\frac{1}{|T|} \cdot \Big(|T| \cdot \mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2)\Big) = \mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2).$$

We know that $\varphi_1 \,\mathrm{U}_n^{\oplus q}\, \varphi_2 \in eCl$. Then by Lemma 4.0.1 we have $\varphi_1 \,\mathrm{U}_n\, \varphi_2 \in lpsub(\psi)$. Set q' to be the value of $\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_n\, \varphi_2)$. Then $P(\varphi_1 \,\mathrm{U}_n\, \varphi_2) = q' \in eCl$ by rule 9 in the definition of eCl. Obviously s satisfies $P(\varphi_1 \,\mathrm{U}_n\, \varphi_2) = q'$. Recall that T is defined to be the equivalence class of s, which means that if $t \in T$, then also t must satisfy $P(\varphi_1 \,\mathrm{U}_n\, \varphi_2) = q'$, and we are done.

□

In the following theorem, we show that criteria (ii) and (iii) bound sizes of small models. If we could prove they hold for every finitely-satisfiable formula, then finite satisfiability would reduce to bounded satisfiability with computable bounds. We emphasize that we have not succeeded in doing so.

Theorem 4.0.4. *Let ψ be a state formula. Then the following are equivalent:*

(i) *ψ has a finite model of size bounded by $|\psi|$.*

(ii) *ψ has a finite model (S, P, v) such that both its branching degree and cardinality of $\{P(s,t) \mid s, t \in S\}$ are bounded by $|\psi|$.*

(*iii*) *ψ has a finite model (S, P, v) such that cardinality of $\{\mathbb{P}(s \models \Psi) \mid s \in S$ and $\Psi \in lbsub(\psi)\}$ is bounded by $|\psi|$.*

Proof. Clearly $(i) \implies (ii)$ holds. We will now show $(ii) \implies (iii)$. Our goal is to prove that the model from (ii) also satisfies the condition in (iii). Obviously $|lbsub(\psi)|$ is bounded by $|\psi|$, so there are only boundedly many choices for $\Psi \in lbsub(\psi)$. Then it is sufficient to show that the cardinality of $\{\mathbb{P}(s \models \Psi) \mid s \in S\}$ is bounded by $|\psi|$ for every $\Psi \in lbsub(\psi)$. Let $\varphi_1 \,\mathrm{U}_n\, \varphi_2 \in lpsub(\psi)$. By using a straightforward induction on i, we will show that for every $i \in \mathbb{N}$, the cardinality of $\{\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_i\, \varphi_2) \mid s \in S\}$ is bounded by $|\psi|$. Suppose $i = 0$. Then $\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_0\, \varphi_2)$ is either 0 or 1 for every $s \in S$ by Lemma 2.6.1. Now the induction step. Suppose $i > 0$. Using Lemma 2.6.1, we see that for every $s \in S$, the value of $\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_i\, \varphi_2)$ is either 0 or 1, or we have

$$\mathbb{P}(s \models \varphi_1 \,\mathrm{U}_i\, \varphi_2) = \sum_{t \in s^\bullet} P(s,t) \cdot \mathbb{P}(t \models \varphi_1 \,\mathrm{U}_{i-1}\, \varphi_2).$$

We will show that there is a set of bounded cardinality which contains all possible values of sums of this form. First, note that every $P(s,t)$ is by definition a member of $\{P(s,t) \mid s,t \in S\}$, which has bounded cardinality by (ii). Also, every $\mathbb{P}(t \models \varphi_1 \,\mathrm{U}_{i-1}\, \varphi_2)$ is by definition a member of $\{\mathbb{P}(t \models \varphi_1 \,\mathrm{U}_{i-1}\, \varphi_2) \mid t \in S\}$, which must have bounded cardinality by the induction hypothesis. Therefore, there are only boundedly many possibilities for the value of each summand. We emphasize that neither of the mentioned sets are affected by the choice of s, so the possibilities stay fixed. We know that the model has a bounded branching degree by (ii), which implies that each sum of this form has only boundedly many summands, and we are done.

Finally, we show $(iii) \implies (i)$. We know that ψ has a finite model (S, P, v) by (iii). Theorem 4.0.3 shows ψ also has a finite model of size at most $|\mathcal{P}(eCl(\psi, M, v))|$. We want to show that $|\mathcal{P}(eCl(\psi, M, v))|$ is bounded by the cardinality of $\{\mathbb{P}(s \models \Psi) \mid s \in S$ and $\Psi \in lbsub(\psi)\}$, which we know must be bounded by $|\psi|$ by (iii).

We will use eCl as a shortcut for $eCl(\psi, M, v)$. It is easy to see that application of rule 9 always adds exactly the same formulae. In other words, applying rules 1-9 one or more times before applying rule 9 has no effect on the formulae added by rule 9. Then eCl contains ψ and formulae added by rule 9, and then all formulae added by repeated application of rules 1-8.

Recall the definition of rule 9, which adds a formula $P(\varphi_1 \mathrm{U}_n \varphi_2) = q$ whenever $\varphi_1 \mathrm{U}_n \varphi_2 \in lbsub(\psi)$ and $q \in \{\mathbb{P}(s \models \varphi_1 \mathrm{U}_n \varphi_2) \mid s \in S\}$. There are only boundedly many options for $\varphi_1 \mathrm{U}_n \varphi_2 \in lpsub(\psi)$. Also, there are only boundedly many options for q, because for each $\varphi_1 \mathrm{U}_n \varphi_2 \in lpsub(\psi)$, the set $\{\mathbb{P}(s \models \varphi_1 \mathrm{U}_n \varphi_2) \mid s \in S\}$ is clearly contained within the set $\{\mathbb{P}(s \models \Psi) \mid s \in S \text{ and } \Psi \in lbsub(\psi)\}$, which is bounded by $|\psi|$ by (iii). Therefore, rule 9 can only add boundedly many formulae into eCl.

It is not hard to see that repeated application of rules 1-8 starting on a set that contains only ψ and formulae added by rule 9 can add at most boundedly many new formulae. In conclusion, the total number of formulae in eCl is bounded by $|\psi|$, and we are done. $\square$

5 Bounded PCTL extended with a single use of $\mathrm{G}^{=1}$

In this chapter, we show that extending *bounded PCTL* by allowing a single use of $\mathrm{G}^{=1}$ breaks the small property. Bounded PCTL is a fragment of PCTL obtained by imposing a step horizon on all until formulae. We introduced bounded PCTL in Subsection 3.3.2. Recall that [4] shows that every satisfiable bounded PCTL formula ψ has a finite model of size bounded by $|\psi|$, which means bounded PCTL has the small model property. Since having this property reduces satisfiability to bounded satisfiability, which is decidable, one might wonder if we can extend bounded PCTL in some simple way and still retain it. A natural direction is to consider X or U restricted to their qualitative version, since we know that satisfiability is decidable for qualitative PCTL (qPCTL).

For example, extending bounded PCTL by allowing a single use of $\mathrm{G}^{>0}$ would break the small model property. To see this, we must first mention an observation about qualitative PCTL from [8]. Qualitative PCTL does not have the small model property, which is implied by existence of satisfiable qPCTL formulae that only have infinite models. An example of such formula is $\mathrm{G}^{>0}(\neg a \wedge \mathrm{F}^{>0}\, a)$. If it had a finite model, one could find a positive $\varepsilon \in \mathbb{Q}$ such that every state of the model that satisfies $\mathrm{F}^{>0}\, a$ must also satisfy $\mathrm{F}^{\geq \varepsilon}\, a$. It follows that the measure of all runs that satisfy $G(\neg a \wedge \mathrm{F}^{>0}\, a)$ must be equal to 0, which implies $\mathrm{G}^{>0}(\neg a \wedge \mathrm{F}^{>0}\, a)$ could not have been satisfied. However, this formula has a simple infinite model, which we show on Figure 5.1. Another example of such formula, taken again from [8], is $\mathrm{G}^{=1}(\mathrm{X}^{>0}\, a) \wedge \mathrm{G}^{>0}\, \neg a$. With this in mind, showing that bounded PCTL extended with a single use of $\mathrm{G}^{>0}$ lacks the small model property is trivial. Clearly every state that satisfies $\neg a \wedge \mathrm{X}^{>0}\, a$ must also satisfy $\neg a \wedge \mathrm{F}^{>0}\, a$, which implies that every finite model of $\mathrm{G}^{>0}(\neg a \wedge \mathrm{X}^{>0}\, a)$ is also a model of $\mathrm{G}^{>0}(\neg a \wedge \mathrm{F}^{>0}\, a)$, which we already know has no finite models.

It is not surprising that bounded PCTL extended with a single use of $\mathrm{G}^{>0}$ loses the small model property, since $G^{>0}\varphi$ technically corresponds to $\neg(\top\, \mathrm{U}^{=1}\, \neg\varphi)$, and these formulae are also responsible for more complicated structure of models in qPCTL [8]. From this

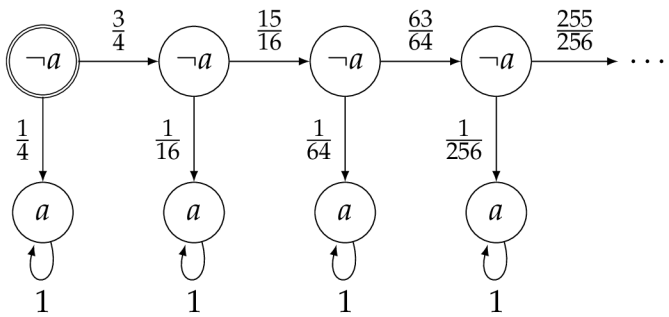

Figure 5.1: An infinite model of $G^{>0}(\neg a \wedge F^{>0}\, a)$.

perspective, $G^{=1}$ seems more amenable to preserving simple model structure, which for bounded PCTL are finite trees with absorbing leaves. Our result, then, shows another possible source of complexity, which is caused by interaction of bounded operators and a single qualitative unbounded operator $G^{=1}$.

The main goal of this chapter is to show that extending bounded PCTL with a single use of $G^{=1}$ breaks the small model property. We will do this by building a formula ψ_{inf} of that form which is satisfiable but has no finite models. Specifically, ψ_{inf} will be of form $\varphi \wedge G^{=1}\, \psi$ where φ is a boolean combination of atomic propositions and ψ is a formula where all operators are bounded.

As a sidenote, it is possible to modify ψ_{inf} so that it is of form $G^{=1}\, \psi$, but we will not do this to avoid more convoluted proofs and models. However, we will later show a model of ψ_{inf} where such modification can be easily deduced. Notice that in this modified form, ψ_{inf} could be intuitively seen as global application of a local condition. Our initial motivation for looking at formulae of this form was precisely to show that globally applying a local condition can be used to enforce infinite models.

We start by recalling the syntax of bounded PCTL:

Definition 5.0.1 (Bounded PCTL syntax)**.** Suppose AP is a countably infinite set. *Bounded PCTL state formulae* are defined as follows:

$$\varphi ::= a \mid \varphi \wedge \varphi \mid \neg\varphi \mid P(\Phi) \oplus q$$

where $a \in AP, \oplus \in \{<, \leq, \geq, >, =, \neq\}$, $q \in [0,1] \cap \mathbb{Q}$, and Φ is a *bounded PCTL path formula*. Those are defined as follows:

$$\Phi ::= \mathrm{X}\,\varphi \mid \varphi\,\mathrm{U}_n\,\varphi$$

where $n \in \mathbb{N}$.

Semantics of bounded PCTL is defined as expected. An extension of bounded PCTL which allows a single use of $G^{=1}$ is defined in the obvious way.

We will now begin the construction of ψ_{inf}. We will use four atomic propositions: a, b, c, t. We will frequently want to express that some of them are satisfied in a state and some are not, which is made more convenient by the following syntactic shortcut:

Definition 5.0.2. Let $A = \{a, b, c, t\}$ be a set of atomic propositions. For all $x_1, x_2, \ldots, x_n \in A$ define

$$\overline{x_1 \wedge x_2 \wedge \ldots \wedge x_n} \equiv x_1 \wedge x_2 \wedge \ldots \wedge x_n \wedge \bigwedge_{y \in A \setminus \{x_1, x_2, \ldots, x_n\}} \neg y.$$

For example, $\overline{a}$ is the formula $a \wedge \neg b \wedge \neg c \wedge \neg t$. The local conditions used in ψ_{inf} will use a, b, c in a somewhat "cyclic" fashion, so it will be useful to define a "successor" syntactic shortcut for them:

Definition 5.0.3. Define $S(a) \equiv b$, $S(b) \equiv c$, and $S(c) \equiv a$.

We can now define the first component of ψ_{inf} called ψ_{succ}, which expresses allowed forms of successors for each state:

Definition 5.0.4 (ψ_{succ})**.** Define

$$\psi_{succ} \equiv \bigwedge_{x \in \{a,b,c\}} \psi_{succ}(x),$$

where

$$\psi_{succ}(x) \equiv \overline{x} \implies \mathrm{F}_1^{=1}(\overline{S(x)} \vee \overline{S(S(x)) \wedge t}).$$

For example, ψ_{succ} says that if a state satisfies only a and no other proposition out of $\{a, b, c, t\}$, then all of its successors must satisfy either only b or only c, t.

The second component of ψ_{inf} is called ψ_{prob}, and it expresses the probabilistic constraints we want to apply to each state:

Definition 5.0.5 (ψ_{prob})**.** Let $q \in \mathbb{Q}$ be such that $\frac{3}{4} < q < 1$ and $\sqrt{4q-3} \in \mathbb{Q}$. Define I to be the interval $(\frac{1-\sqrt{4q-3}}{2}, \frac{1+\sqrt{4q-3}}{2})$, and define

$$\psi_{prob} \equiv \bigwedge_{x \in \{a,b,c\}} \psi_{prob}(x),$$

where

$$\begin{aligned} \psi_{prob}(x) \equiv \overline{x} \implies (\ & P(\mathrm{F}_1\, S(x)) \in I \\ & \wedge P(\mathrm{F}_1\, S(S(x))) \in I \\ & \wedge P(\mathrm{F}_2\, S(S(x))) = q). \end{aligned}$$

The definition above might raise the question of whether ψ_{prob} is in fact a well-defined formula, since the interval I may turn out to not be well-defined, or might have an endpoint strictly less than 0 or strictly greater than 1, which would cause ψ_{prob} to not be a formula from a syntactic standpoint. However, this does not happen as a consequence of Lemma 5.0.2 (shown later), and we have $0 < \frac{1-\sqrt{4q-3}}{2} < \frac{1+\sqrt{4q-3}}{2} < 1$. A more subtle syntactic caveat is that all numeric constants in a PCTL formula must be rational numbers, which is why we require q and both endpoints of I to be rational. The rationality of endpoints of I is implied by the rationality of $\sqrt{4q-3}$. Note that such values of q do exist: for example $q = \frac{13}{16}$ is certainly a rational number such that $\frac{3}{4} < q < 1$ and also $\sqrt{4q-3} = \frac{1}{2}$.

We can now to define ψ_{inf}:

Definition 5.0.6 (ψ_{inf})**.** Define

$$\psi_{inf} \equiv \overline{a} \wedge \mathrm{G}^{=1}(\psi_{succ} \wedge \psi_{prob}).$$

The crux of our proof that ψ_{inf} has no finite models lies in the fact that the local probabilistic constraints expressed by ψ_{prob} together with ψ_{succ} imply that whenever we find a state s that satisfies $\overline{x}$ for some $x \in \{a, b, c\}$, it must have a successor t that satisfies $\overline{S(x)}$ such that $\mathbb{P}(s \models \mathrm{F}_1\, S(x)) < \mathbb{P}(t \models \mathrm{F}_1\, S(S(x)))$. For finite models, the set of values of $\mathbb{P}(s \models \mathrm{F}_1\, S(x))$ for states s that satisfy $\overline{x}$ for some $x \in \{a, b, c\}$ must be finite, which means it must have a maximum, which contradicts our previous observation. On the other hand, a simple infinite model with branching degree 2 can be easily found.

First, we will establish a lemma that gives us a lower bound on a measure of certain runs for successors of certain states:

Lemma 5.0.1. *Let $M = (S, P, v)$ be a finite Markov chain, φ some state formula, and $q \in [0,1] \cap \mathbb{Q}$. Suppose $s \models P(\mathrm{F}_2\, \varphi) \geq q$ and also $s \models P(\mathrm{F}_1\, \varphi) < q$. Then there exists $r \in s^{\bullet}_{\neg\varphi}$ such that*

$$\mathbb{P}(r \models \mathrm{F}_1\, \varphi) \geq \frac{q - 1 + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}.$$

Proof. We will first show a few simple propositions:

(i) $s \models \neg\varphi$. If not, then $s \models \varphi$, which implies $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = 1$ by Lemma 2.6.1. We know that $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) < q$ from the lemma statement, so $1 < q$, which contradicts the lemma statement.

(ii) $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = \sum_{t \in s^{\bullet}_{\varphi}} P(s,t)$. Since (i) holds, we can simply use Corollary 2.6.3.

(iii) $\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) = 1 - \sum_{t \in s^{\bullet}_{\varphi}} P(s,t)$. By the definition of Markov chains $\sum_{t \in s^{\bullet}} P(s,t) = 1$. Clearly $s^{\bullet}_{\varphi}$ and $s^{\bullet}_{\neg\varphi}$ partition $s^{\bullet}$, so we can split the sum in two parts, giving us $1 = \sum_{t \in s^{\bullet}_{\varphi}} P(s,t) + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)$.

(iv) $\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) > 0$. We have $\sum_{t \in s^{\bullet}_{\varphi}} P(s,t) = \mathbb{P}(s \models \mathrm{F}_1\, \varphi)$ by (ii). By the lemma statement $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) < q$ and $q \leq 1$, which then implies $\sum_{t \in s^{\bullet}_{\varphi}} P(s,t) < 1$. We showed $\sum_{t \in s^{\bullet}_{\varphi}} P(s,t) + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) = 1$ in (iii). Then $\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) > 0$.

We are now ready to prove the lemma. First, note that $\frac{q-1+\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}$ must be defined because of (iv). Now, suppose that for every $t \in s^{\bullet}_{\neg\varphi}$, we have

$$\mathbb{P}(t \models \mathrm{F}_1\, \varphi) < \frac{q - 1 + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}.$$

For each $t \in s^{\bullet}_{\neg\varphi}$, multiply both sides by (necessarily positive) $P(s,t)$ to get inequalities of form

$$P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_1\, \varphi) < P(s,t) \cdot \frac{q - 1 + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}.$$

Now we can sum up all these inequalities together and get

$$\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_1\, \varphi) < \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) \cdot \frac{q - 1 + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}.$$

Here we had to rely on the fact that $s^{\bullet}_{\neg\varphi}$ is not empty (because of (iv)), which was not needed for the previous inequalities since they are universally quantified for $t \in s^{\bullet}_{\neg\varphi}$.

Note that $\frac{q-1+\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}$ does not depend on t, so we can factor it out and simplify the expression on the right side to

$$\frac{q - 1 + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)}{\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t)} \cdot \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t),$$

which is clearly equal to

$$q - 1 + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t),$$

which by (iii) is in turn equal to

$$q - \sum_{t \in s^{\bullet}_{\varphi}} P(s,t).$$

In conclusion, the inequality can be simplified to

$$\sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_1\, \varphi) < q - \sum_{t \in s^{\bullet}_{\varphi}} P(s,t).$$

By adding $\sum_{t \in s^{\bullet}_{\varphi}} P(s,t)$ to both sides, we get

$$\sum_{t \in s^{\bullet}_{\varphi}} P(s,t) + \sum_{t \in s^{\bullet}_{\neg\varphi}} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_1\, \varphi) < q. \tag{5.1}$$

On the other hand, a simple application of Lemma 2.6.1 gives

$$\mathbb{P}(s \models \mathrm{F}_2\, \varphi) = \sum_{t \in s^{\bullet}} P(s,t) \cdot \mathbb{P}(t \models \mathrm{F}_1\, \varphi).$$

From the lemma statement we know that $\mathbb{P}(s \models F_2\, \varphi) \geq q$, which implies

$$\sum_{t \in s^\bullet} P(s,t) \cdot \mathbb{P}(t \models F_1\, \varphi) \geq q.$$

Analogously to previous paragraphs, we can split this sum into two parts to get

$$\sum_{t \in s^\bullet_\varphi} P(s,t) \cdot \mathbb{P}(t \models F_1\, \varphi) + \sum_{t \in s^\bullet_{\neg\varphi}} P(s,t) \cdot \mathbb{P}(t \models F_1\, \varphi) \geq q.$$

The left summand $\sum_{t \in s^\bullet_\varphi} P(s,t) \cdot \mathbb{P}(t \models F_1\, \varphi)$ is equal to $\sum_{t \in s^\bullet_\varphi} P(s,t)$, because for each $t \in s^\bullet_\varphi$ we have $\mathbb{P}(t \models F_1\, \varphi) = 1$ by Lemma 2.6.1. This implies

$$\sum_{t \in s^\bullet_\varphi} P(s,t) + \sum_{t \in s^\bullet_{\neg\varphi}} P(s,t) \cdot \mathbb{P}(t \models F_1\, \varphi) \geq q,$$

which obviously contradicts inequality (5.1). □

Now we will state a straightforward numeric lemma. The probabilistic constraints in ψ_{prob} were chosen precisely so that at some point this lemma would be applicable and we will get (among other things) a strict lower bound for the lower bound stated in the previous lemma.

Lemma 5.0.2. *Suppose x, q are real numbers such that $\frac{3}{4} < q < 1$ and $x \in (\frac{1-\sqrt{4q-3}}{2}, \frac{1+\sqrt{4q-3}}{2})$. Then it must be that $x < \frac{q-1+x}{x}$ and $\frac{q-1+x}{x} \in (\frac{1-\sqrt{4q-3}}{2}, \frac{1+\sqrt{4q-3}}{2})$ and also $\frac{1+\sqrt{4q-3}}{2} < q$.*

A proof of this lemma can be done with standard methods and is not of particular interest, so we only include a sketch:

Proof sketch. Since $\frac{3}{4} < q < 1$, it is easy to see that $\frac{1-\sqrt{4q-3}}{2}$ is defined and is strictly greater than 0. Because this is the left endpoint of an open interval x is in, we have $x > 0$. This also implies $\frac{q-1+x}{x}$ is defined on the interval.

Suppose that $x = \frac{q-1+x}{x}$. This equation can be equivalently reduced to $x^2 - x - q + 1 = 0$. The discriminant of $x^2 - x - q + 1$ is equal to $\sqrt{4q-3}$, which must be positive because $\frac{3}{4} < q$. Therefore, the polynomial has two roots, namely $\frac{1-\sqrt{4q-3}}{2}$ and $\frac{1+\sqrt{4q-3}}{2}$, which

we initially chose as the endpoints of the interval for this purpose. As a sidenote, we now also must have $\frac{1-\sqrt{4q-3}}{2} < \frac{1+\sqrt{4q-3}}{2}$, so the interval is well-defined.

Now we will show that $x < \frac{q-1+x}{x}$. Since $x > 0$, we can equivalently modify the inequality to $x^2 - x + 1 - q < 0$. We have already shown that the roots of the polynomial on the left side are the endpoints of our interval. It can be easily seen that for all x strictly within these endpoints, the value of $x^2 - x + 1 - q$ is negative.

Now we will show that $\frac{q-1+x}{x} \in (\frac{1-\sqrt{4q-3}}{2}, \frac{1+\sqrt{4q-3}}{2})$. First we show that $\frac{q-1+x}{x}$ is increasing for all positive $x \in \mathbb{R}$, that is

$$\frac{q-1+x_1}{x_1} < \frac{q-1+x_2}{x_2}$$

whenever $x_1 < x_2$. We know that $0 < x_1 < x_2$ by their definitions. After multiplying both sides of the inequality by (positive) x_1x_2 it reduces to $x_2(q-1+x_1) < x_1(q-1+x_2)$. Simplifying both sides, we get $x_2q - x_2 + x_1x_2 < x_1q - x_1 + x_1x_2$, which in turn is equivalent to $x_1 - x_2 < x_1q - x_2q$. Factoring out q and dividing by the (negative) $x_1 - x_2$, we see that this is equivalent to $q < 1$, which is true by the lemma statement. Because $\frac{q-1+x}{x}$ is strictly increasing for all positive $x \in \mathbb{R}$ and $0 < \frac{1-\sqrt{4q-3}}{2}$, we can get an upper (lower) bound on its maximal (minimal) value by using x equal to the right (left) endpoint of the interval. We have already shown that when x is equal to the left and right endpoint, the value of $\frac{q-1+x}{x}$ is equal to the left and right endpoint, respectively. Since the endpoints themselves are not included, both maximal and minimal bounds are strict, and we have what we wanted.

Finally, we show that $\frac{1+\sqrt{4q-3}}{2} < q$. We already know that $\frac{1+\sqrt{4q-3}}{2}$ is defined, and with a simple reduction we can see that the inequality is equivalent to $\sqrt{4q-3} < 2q - 1$. Using standard methods, we see that this holds for either $q > 1$ or $\frac{3}{4} \leq q < 1$, ours being the latter case. Note that this also implies $\frac{1+\sqrt{4q-3}}{2} < 1$, since $q < 1$. □

The final lemma we will need states a simple consequence of ψ_{succ}, but in a way that will be convenient to use in the proof of our main theorem:

Lemma 5.0.3. *Let $M = (S, P, v)$ be a finite Markov chain and $u \in S$ a state such that $u \models \psi_{inf}$. Suppose $s \in u^*$ is a state such that $s \models \overline{x}$ for some $x \in \{a, b, c\}$. Then $t \models \overline{S(x)} \iff t \models S(x) \iff t \models \neg S(S(x))$ for every $t \in s^\bullet$.*

Proof. Set $\varphi \equiv \overline{S(x)} \vee \overline{S(S(x)) \wedge t}$. We will now follow our usual line of argument. By the definition of Markov chains, we have $1 = \sum_{t \in s^\bullet} P(s, t)$. Partition this sum depending on whether t satisfies φ or not. Suppose there exists some $t \in s^\bullet_{\neg\varphi}$. Then $\sum_{t \in s^\bullet_\varphi} P(s, t) < 1$. Because $s \models \overline{x}$, we know that s must not satisfy $S(x)$ or $S(S(x))$, which implies $s \models \neg\varphi$. Then by Corollary 2.6.3 we have $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = \sum_{t \in s^\bullet_\varphi} P(s, t)$, and so $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) < 1$. At the same time, we know that $s \models \overline{x} \wedge \psi_{succ}$, which implies $s \models P(\mathrm{F}_1\, \varphi) = 1$. Together, we see that $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) = 1$, but also that $\mathbb{P}(s \models \mathrm{F}_1\, \varphi) < 1$. In conclusion, every $t \in s^\bullet$ must satisfy $\overline{S(x)} \vee \overline{S(S(x)) \wedge t}$.

Let $t \in s^\bullet$. Clearly $t \models \overline{S(x)}$ implies $t \models S(x)$. Suppose $t \models S(x)$. If t satisfied $\overline{S(S(x)) \wedge t}$, it would not satisfy $S(x)$. Then $t \models \overline{S(x)}$, which implies $t \models \neg S(S(x))$. If $t \models \neg S(S(x))$, then clearly $t \not\models \overline{S(S(x)) \wedge t}$, and so $t \models \overline{S(x)}$. □

The reason why we require that $s \in u^*$ and not just $s \in S$ is a somewhat awkward technical detail. If s was not reachable from any state which satisfies ψ_{inf}, it would not have to satisfy ψ_{succ} and ψ_{prob}. Of course, states that are not reachable from any state that satisfies ψ_{inf} have no effect on whether ψ_{inf} is satisfied, and can be safely removed from the model. Instead of adding a convention that no such states are present in our models, we simply work with states that are reachable from a state that satisfies ψ_{inf}.

We are now ready to prove the main theorem of this chapter:

Theorem 5.0.4. *The formula ψ_{inf} is satisfiable, but only has infinite models.*

Proof. Let us recall that $q \in \mathbb{Q}$ is a rational number such that $\frac{3}{4} < q < 1$ and $\sqrt{4q - 3} \in \mathbb{Q}$, and I is the open interval $(\frac{1-\sqrt{4q-3}}{2}, \frac{1+\sqrt{4q-3}}{2})$.

We will first show that ψ_{inf} has no finite models. Suppose there exists a finite Markov chain $M = (S, P, v)$ such that $u \models \psi_{inf}$ for some state $u \in S$. Define

$$R = \{\mathbb{P}(t \models \mathrm{F}_1\, S(x)) \mid t \in u^*, x \in \{a, b, c\} \text{ and } t \models \overline{x}\}.$$

It is easy to show that R is finite but not empty: there are only finitely many choices for t and x, so the set R must be finite. Since $u \models \psi_{inf}$, we have $u \models \overline{a}$. Since $u \in u^*$ and $u \models \overline{a}$ and $a \in \{a, b, c\}$, we have $\mathbb{P}(u \models \mathrm{F}_1\, S(a)) \in R$, which ensures R is not empty.

Given that R is a finite but non-empty subset of real numbers, it must have a maximum. Set m to be the maximum of R. Pick $s \in u^*$ and $x \in \{a, b, c\}$ such that $s \models \overline{x}$ and $\mathbb{P}(s \models \mathrm{F}_1\, S(x)) = m$.

We know that $s \models \overline{x}$, but also $s \models \psi_{prob}$ because $s \in u^*$. It follows that $s \models P(\mathrm{F}_2\, S(S(x))) = q$ and $s \models P(\mathrm{F}_1\, S(S(x))) \in I$. Obviously this means $s \models P(\mathrm{F}_2\, S(S(x))) \geq q$. Moreover, since q is a rational number such that $\frac{3}{4} < q < 1$, by Lemma 5.0.2 we have $\frac{1+\sqrt{4q-3}}{2} < q$, which shows $s \models P(\mathrm{F}_1\, S(S(x))) < q$. Then by Lemma 5.0.1 there exists $r \in s^\bullet$ such that $r \models \neg S(S(x))$ and

$$\mathbb{P}(r \models \mathrm{F}_1\, S(S(x))) \geq \frac{q - 1 + \sum_{t \in s^\bullet_{\neg S(S(x))}}}{\sum_{t \in s^\bullet_{\neg S(S(x))}}}.$$

We will now show a few simple propositions:

(i) $s \not\models S(x)$. We know that $s \models \overline{x}$.

(ii) $m = \sum_{t \in s^\bullet_{\neg S(S(x))}} P(s, t)$. By definition, $m = \mathbb{P}(s \models \mathrm{F}_1\, S(x))$. Since $s \not\models S(x)$, we have $\mathbb{P}(s \models \mathrm{F}_1\, S(x)) = \sum_{t \in s^\bullet_{S(x)}} P(s, t)$ by Corollary 2.6.3. Notice that $u \models \psi_{inf}$ and $s \in u^*$ and $s \models \overline{x}$. Then by Lemma 5.0.3 we have $t \models S(x)$ if and only if $t \models \neg S(S(x))$ for every $t \in s^\bullet$, which implies $\sum_{t \in s^\bullet_{S(x)}} P(s, t) = \sum_{t \in s^\bullet_{\neg S(S(x))}} P(s, t)$.

(iii) $m \in I$. Since $s \models \overline{x} \wedge \psi_{prob}$, we have $s \models P(\mathrm{F}_1\, S(x)) \in I$. Then $\mathbb{P}(s \models \mathrm{F}_1\, S(x)) \in I$, but that is equal to m by definition.

(iv) $r \models \overline{S(x)}$. Recall that $u \models \psi_{inf}$ and $s \in u^*$ and $s \models \overline{x}$ and $r \in s^\bullet$ and $r \models \neg S(S(x))$, so we get this from Lemma 5.0.3.

(v) $\mathbb{P}(r \models \mathrm{F}_1\, S(S(x))) > m$. We already know that

$$\mathbb{P}(r \models \mathrm{F}_1\, S(S(x))) \geq \frac{q - 1 + \sum_{t \in s^\bullet_{\neg S(S(x))}}}{\sum_{t \in s^\bullet_{\neg S(S(x))}}}.$$

Since $m = \sum_{t \in s^{\bullet}_{\neg S(S(x))}} P(s,t)$ by (ii), this means

$$\mathbb{P}(r \models \mathrm{F}_1\, S(S(x))) \geq \frac{q-1+m}{m}.$$

Since $\frac{3}{4} < q < 1$ by definition and $m \in I$ by (iii), we can use Lemma 5.0.2 to get $\frac{q-1+m}{m} > m$.

Observe that $r \in u^*$ because $r \in s^{\bullet}$ and $s \in u^*$. Also by (iv) we have $r \models \overline{S(x)}$ where $S(x) \in \{a,b,c\}$, so $\mathbb{P}(r \models \mathrm{F}_1\, S(S(x))) \in R$. However, we also have $\mathbb{P}(r \models \mathrm{F}_1\, S(S(x))) > m$ by (v), which contradicts the maximality of m.

What remains is to show that ψ_{inf} is satisfiable. We will describe a model of ψ_{inf} in an informal fashion, although it should be fairly easy to see that it does satisfy ψ_{inf}. We will use a recursive definition: for $x \in \{a,b,c\}$ and $r \in \mathbb{R}$, define $node(x,r)$ to be an infinite tree rooted in a node called t with exactly two succesors t_1 and t_2 such that $P(t,t_1) = r$ and $P(t,t_2) = 1-r$. The node t satisfies $\overline{x}$. The successor t_2 satisfies $\overline{S(S(x)) \wedge t}$ and $P(t_2,t_2) = 1$. The successor t_1 is the root of $node(S(x), \frac{q-1+r}{r})$. Pick an arbitrary $r_0 \in \mathbb{R}$ such that $r_0 \in I$. We will argue that the root of $node(a,r_0)$ is a model of ψ_{inf}. We show $node(a,r_0)$ in Figure 5.2.

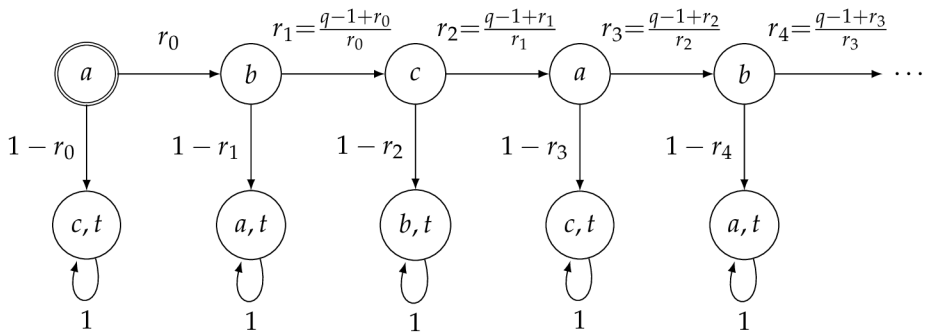

Figure 5.2: Example of a model of ψ_{inf} called $node(a,r_0)$.

Since $r_0 \in I$ and $\frac{3}{4} < q < 1$ by their respective definitions, we have $r_1 \in I$ by Lemma 5.0.2. A simple induction would then show that $r_i \in I$ for all $i \in \mathbb{N}$. Now, observe that if $r_i \in I$, then $r_i = \frac{1-\sqrt{4q-3}}{2} + \delta$ for

some $0 < \delta < \frac{1+\sqrt{4q-3}}{2} - \frac{1-\sqrt{4q-3}}{2}$. Then $1 - r_i = 1 - (\frac{1-\sqrt{4q-3}}{2} + \delta)$, which is equal to $\frac{1+\sqrt{4q-3}}{2} - \delta$. Looking at the definition of δ, we see that this implies $1 - r_i \in I$. To sum it up, we have both $r_i \in I$ and $1 - r_i \in I$ for all $i \in \mathbb{N}$. It is easy to see that this implies both r_i and $1 - r_i$ are in the open interval $(0, 1)$ (see proof of Lemma 5.0.2). Therefore, every $t \in node(a, r_0)$ that satisfies $\overline{x}$ for some $x \in \{a, b, c\}$ has exactly two successors and their transition probabilities form a probability distribution, so $node(a, r_0)$ is a Markov chain.

Clearly the root of $node(a, r_0)$ satisfies $\overline{a}$, and every state in it satisfies ψ_{succ}. We will now show that every state also satisfies ψ_{prob}. Let $t \in S$ be a state that satisfies $\overline{x}$ for some $x \in \{a, b, c\}$. Let r_i and $1 - r_i$ denote the transition probabilities to its successors t_1 and t_2, respectively. Here t_2 is the successor such that $P(t_2, t_2) = 1$. We have already shown that both r_i and $1 - r_i$ must be in I, and so t satisfies $P(\mathrm{F}_1\, S(x)) \in I$ and $P(\mathrm{F}_1\, S(S(x))) \in I$. Also by repeatedly applying Lemma 2.6.1 on t and its successors, we can show that

$$\begin{aligned}
\mathbb{P}(t \models \mathrm{F}_2\, S(S(x))) &= \sum_{k \in t^\bullet} P(t,k) \cdot \mathbb{P}(k \models \mathrm{F}_1\, S(S(x))) \\
&= (1 - r_i) \cdot \mathbb{P}(t_2 \models \mathrm{F}_1\, S(S(x))) + r_i \cdot \mathbb{P}(t_1 \models \mathrm{F}_1\, S(S(x))) \\
&= (1 - r_i) \cdot 1 + r_i \cdot \big((1 - r_{i+1}) \cdot 0 + r_{i+1} \cdot 1\big) \\
&= (1 - r_i) + r_i \cdot r_{i+1} \\
&= (1 - r_i) + r_i \cdot \frac{q - 1 + r_i}{r_i} \\
&= (1 - r_i) + q - 1 + r_i \\
&= q.
\end{aligned}$$

Then $t \models P(\mathrm{F}_2\, S(S(x))) = q$, which now gives us $t \models \psi_{prob}$, and we are finally done. □

6 Formulae of form $\bigwedge_{i \in I} P(\varphi_i \mathrm{U} \psi_i) \in (l_i, r_i)$ where φ_i and ψ_i are qualitative PCTL formulae

In this chapter, we show that both finite and general satisfiability is decidable for a fragment of PCTL that consists of conjunctions of formulae of form $P(\varphi \mathrm{U} \psi) \in I$ where I is an open interval and φ, ψ are qualitative PCTL formulae.

We will show that every satisfiable formula of this form has models which have representations (we call them *pre-models*) whose existence is decidable. We will show that every model induces a pre-model, and every pre-model induces a model. Decidability follows from the fact that every pre-model consists of a quantitative part whose existence reduces to a decidable SMT problem, and a qualitative part whose existence is decidable by results from [8].

Convention. By ζ we always denote formulae of form

$$\bigwedge_{i \in I} P(\varphi_i \mathrm{U} \psi_i) \in (l_i, r_i)$$

where I is a finite set of indices, and φ_i, ψ_i are qPCTL state formulae and (l_i, r_i) are open intervals for every $i \in I$. By $\mathcal{F}$ we always mean the set $\{\varphi_i \mathrm{U} \psi_i \mid i \in I\}$. Throughout this chapter, we assume ζ and $\mathcal{F}$ are given by context, unless specified otherwise.

This chapter is split into four sections. The last section shows our main results, and the previous three introduce the necessary concepts. To motivate the following three sections, we outline the strategy used for solving the satisfiability problem. A pre-model $(T, \mathcal{A}, \{m_x\}_{x \in T})$ for ζ is obtained from a model M of ζ in the following way:

- Classify runs based on "types". A type encodes which formulae from $\{\varphi_i \mathrm{U} \psi_i \mid i \in I\}$ are satisfied.
- Partition all runs in M by their types. If runs of type x have a positive measure m_x, then M must contain paths with certain structural properties.
- These properties can be described with a qPCTL formula. This formula is satisfied if and only if such paths exist. We will use λ to denote such formulae.

- Define T to be the set of all types that have positive measure. Define λ_x to be the relevant qPCTL formula for $x \in T$.
- The qualitative part of the pre-model consists of a representation $\mathcal{A}$ of a model of $\bigwedge_{x \in T} \lambda_x$. We use *pseudo-models* from [8]. The quantitative part of the pre-model consists of a collection of constants $\{m_x\}_{x \in T}$, and we show these m_x satisfy the required quantitative properties.

A pre-model $(T, \mathcal{A}, \{m_x\}_{x \in T})$ can be turned back into a model in the following way:

- We follow [8] and turn $\mathcal{A}$ back into a model M of $\bigwedge_{x \in T} \lambda_x$.
- M is then modified to ensure certain structural properties of paths whose existence is implied by $\bigwedge_{x \in T} \lambda_x$.
- Denote $\mathbb{P}(x)$ to be the measure of runs of type x. For every $\delta > 0$, we can modify M so that $|\mathbb{P}(x) - m_x| < \delta$ for each $x \in T$. Intuitively, we can "approximate" each m_x to arbitrary precision.
- Due to our restriction to open intervals, this ensures that ζ can be satisfied.

We start by recalling the syntax of qualitative PCTL:

Definition 6.0.1 (qPCTL syntax)**.** Let AP be a countably infinite set. Then *qPCTL state formulae* are defined as follows:

$$\varphi ::= a \mid \varphi \wedge \varphi \mid \neg\varphi \mid P(\Phi) \oplus q$$

where $a \in AP, \oplus \in \{<, \leq, \geq, >, =, \neq\}$, $q \in \{0, 1\}$, and Φ is a *qPCTL path formula*, which are defined as follows:

$$\Phi ::= \mathrm{X}\,\varphi \mid \varphi \,\mathrm{U}\, \varphi$$

We recall an observation from [8] that we also used in Chapter 4:

Convention. Suppose we have a qPCTL formula of form $\mathrm{X}^{\oplus q}\,\varphi$ or $\varphi \,\mathrm{U}^{\oplus q}\, \psi$. Every such formula must be semantically equivalent to a tautology, a contradiction, or one of $\mathrm{X}^{>0}\,\varphi$, $\mathrm{X}^{=1}\,\varphi$, $\varphi \,\mathrm{U}^{>0}\, \psi$, $\varphi \,\mathrm{U}^{=1}\, \psi$, and their negations. Since qPCTL contains negations, we can assume without loss of generality that $\oplus q$ in formulae of form $\mathrm{X}^{\oplus q}\,\varphi$ or $\varphi \,\mathrm{U}^{\oplus q}\, \psi$ is either >0 or $=1$.

We will frequently use the following standard lemma:

Lemma 6.0.1. *Let (S, P, v) be a Markov chain, p a finite path, and s its last state $p(len(p) - 1)$. Suppose $L \subseteq Run(p)$ is a measurable set. Then $\mathbb{P}(L) = \mathbb{P}(Run(p)) \cdot \mathbb{P}(\{sw \mid pw \in L\})$.*

It will be useful to extend the semantics of until formulae to finite paths, and we do it in the expected way:

Definition 6.0.2 ($p \models \varphi_1 \,\mathrm{U}\, \varphi_2$)**.** Let $M = (S, P, v)$ be a Markov chain. For $p \in Paths_{fin}$ and a PCTL path formula $\varphi_1 \,\mathrm{U}\, \varphi_2$, define

$$M, p \models^v \varphi_1 \,\mathrm{U}\, \varphi_2 \text{ iff } \exists i \in \mathbb{N} \text{ such that } i \leq len(p) - 1 \text{ and } p(i) \models \varphi_2 \text{ and } p(k) \models \varphi_1 \text{ for every } k \in \mathbb{N} \text{ such that } k < i.$$

6.1 Types of runs

In this section, we define *types* of runs and state basic lemmas used to deal with them.

Convention. We assume that $M = (S, P, v)$ is a Markov chain given by context, unless specified otherwise.

We want to classify runs based on until formulae from $\mathcal{F}$ they satisfy. We do this by defining types of runs. A type of a run records which until formulae from $\mathcal{F}$ are satisfied and which are not. For those until formulae that are not satisfied, we also want to record the reason why. If a run does not satisfy $\varphi \,\mathrm{U}\, \psi$, then it must either satisfy $\neg\psi \,\mathrm{U}(\neg\psi \wedge \neg\varphi)$ or $\mathrm{G}(\varphi \wedge \neg\psi)$. It is not hard to see that all three possibilities are mutually exclusive, but also that one of them must happen. Then we define the type of a run to be an ordered triple (x_1, x_2, x_3), where x_1, x_2, x_3 contain the corresponding formulae from $\mathcal{F}$. Formally:

Definition 6.1.1 ($type(\pi)$)**.** For every $\pi \in Paths$, define $type(\pi)$ to be (x_1, x_2, x_3) such that

- $x_1 = \{\varphi \,\mathrm{U}\, \psi \in \mathcal{F} \mid \pi \models \varphi \,\mathrm{U}\, \psi\}$,
- $x_2 = \{\varphi \,\mathrm{U}\, \psi \in \mathcal{F} \mid \pi \models \neg\psi \,\mathrm{U}(\neg\psi \wedge \neg\varphi)\}$,

- $x_3 = \{\varphi \,\mathrm{U}\, \psi \in \mathcal{F} \mid \pi \models \mathrm{G}(\varphi \wedge \neg\psi)\}$.

It follows that every type defines a partition of $\mathcal{F}$. The set of all types with respect to $\mathcal{F}$ is denoted by *type*, and is formally defined as the set of all ordered triples (x_1, x_2, x_3) such that $\{x_1, x_2, x_3\}$ is a partition of $\mathcal{F}$. Clearly $|type|$ is finite and bounded by $|\mathcal{F}|$.

Convention. We usually use x and y to denote types. If x is a type, then x_1, x_2, x_3 always refers to the corresponding elements of the ordered triple x. In other words, if $x = (\alpha, \beta, \gamma)$, then $x_1 = \alpha$, $x_2 = \beta$, and $x_3 = \gamma$. This allows us to say $x \in type$ and use x_1, x_2, x_3 without explicitly stating that $x = (x_1, x_2, x_3)$.

We will now make a few simple observations about types, which will be used later without reference. Let p be a path and $x \in type$ a type. Define

$$X = x_1 \cup \{\neg\psi \,\mathrm{U}(\neg\psi \wedge \neg\varphi) \mid \varphi \,\mathrm{U}\, \psi \in x_2\},$$

$$always \equiv \bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3} (\varphi \wedge \neg\psi).$$

Then the following holds:

- $type(p) = x$ if and only if $p \models \mathrm{G}\, always$ and $p \models \Phi$ for every $\Phi \in X$. Set $y = type(p)$. Observe that if $p \models \Phi$ for every $\Phi \in X$, then $x_1 \subseteq y_1$ and $x_2 \subseteq y_2$. Also $p \models \mathrm{G}\, always$ ensures that $x_3 \subseteq y_3$. Since both $\{x_1, x_2, x_3\}$ and $\{y_1, y_2, y_3\}$ are partitions of $\mathcal{F}$, we necessarily have $(x_1, x_2, x_3) = (y_1, y_2, y_3)$. The opposite direction clearly holds.

- if p was finite and $type(p) = x$, then $type(pw) = x$ for every run $pw \in Run(p)$ such that $w \models \mathrm{G}\, always$. This is because all formulae in X are already satisfied on p, and $w \models \mathrm{G}\, always$ ensures that $pw \models \mathrm{G}\, always$.

- if $type(p) = x$, then p must have a unique minimal-length prefix of type x. This is because each $\Phi \in X$ must be satisfied after finitely many states, and every state of p satisfies *always*. Then p has at least one prefix of type x, and among those we can take the one with minimal length, which is clearly unique.

Definition 6.1.2 ($\mathbb{P}_s(x)$)**.** For $s \in S$ and $x \in type$, define $\mathbb{P}_s(x)$ to be a shortcut for $\mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\})$. The subscript s is required here, since $\mathbb{P}(x)$ does not make it clear which state do the considered runs start in.

In the following definition, $p \rightsquigarrow x$ expresses that a finite path p is a "witness" for a set runs of type x, and in fact this set must have positive measure as a consequence of Lemma 6.0.1.

Definition 6.1.3 ($p \rightsquigarrow x$)**.** Let $x \in type$ be a type and p a finite path. We say $p \rightsquigarrow x$ if and only if $type(p) = x$ and

$$p(len(p) - 1) \models \mathrm{G}^{>0} \bigwedge_{\varphi \mathrm{U} \psi \in x_3} (\varphi \wedge \neg \psi).$$

The following lemma shows the opposite direction: whenever a set of runs of type x has a positive measure, there must be a finite path p such that $p \rightsquigarrow x$.

Lemma 6.1.1. *Let $M = (S, P, v)$ be Markov chain, $s \in S$, $x \in type$. Then $\mathbb{P}_s(x) > 0$ implies there exists $p \in Paths_{fin}(s)$ such that $p \rightsquigarrow x$.*

Proof. Define $R = \{\pi \in Run(s) \mid type(\pi) = x\}$ and $A \subseteq Paths_{fin}$ to be a set such that $p \in A$ if and only if

- $type(p) = x$,
- $type(p') = x$ implies $p' = p$ for every $p' \in Pref(p)$.

The set A describes minimal-length paths of type x. We know that every $\pi \in R$ has a unique minimal-length prefix of type x. It follows that every $\pi \in R$ has exactly one prefix p such that $p \in A$. It follows that the countable family $\{R \cap Run(p)\}_{p \in A}$ partitions R. Each $R \cap Run(p)$ is clearly measurable (R can be easily seen to be measurable). By countable additivity of measures, we get:

$$\mathbb{P}_s(R) = \mathbb{P}_s(\bigcup_{p \in A} R \cap Run(p)) = \sum_{p \in A} \mathbb{P}_s(R \cap Run(p)).$$

We know that $\mathbb{P}_s(R) > 0$, so there must exist at least one $p \in A$ such that $\mathbb{P}_s(R \cap Run(p)) > 0$. Fix such p and set $R_p = R \cap Run(p)$. We will show $p \rightsquigarrow x$. We already know $type(p) = x$, so it remains to show

$$p(len(p) - 1) \models \mathrm{G}^{>0}(\bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3} (\varphi \wedge \neg\psi)).$$

Set $t = p(len(p) - 1)$ and $Y = \{t \models \mathrm{G}(\bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3}(\varphi \wedge \neg\psi))\}$. We will show $\mathbb{P}_t(Y) > 0$. Clearly $R_p \subseteq Run(p)$, and by Lemma 6.0.1 we have

$$\mathbb{P}_s(R_p) = \mathbb{P}_s(Run(p)) \cdot \mathbb{P}_t(\{tw \mid pw \in R_p\}).$$

Since $\mathbb{P}_s(R_p) > 0$, we get $\mathbb{P}_t(\{tw \mid pw \in R_p\}) > 0$. We will show $Y = \{tw \mid pw \in R_p\}$. Suppose $\pi \in Y$. Then $\pi = tw$ for some run w such that $w \models \mathrm{G}(\bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3}(\varphi \wedge \neg\psi))$. We know that $type(p) = x$ and $w \models \mathrm{G}(\bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3}(\varphi \wedge \neg\psi))$, and so $type(pw) = x$. Clearly p is a prefix of pw, which means $pw \in R_p$, which means $tw \in \{tw \mid pw \in R_p\}$. Suppose $tw \in \{tw \mid pw \in R_p\}$. Then $type(pw) = x$, so every state in pw satisfies $\mathrm{G}(\bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3}(\varphi \wedge \neg\psi))$. Given that t is the last state of p, necessarily $tw \models \mathrm{G}(\bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3}(\varphi \wedge \neg\psi))$, which implies $tw \in Y$. $\square$

We want to encode the property $p \rightsquigarrow x$ with a qPCTL formula, which will allow us to use results from [8] to show that such paths also exist in a Markov chain with a decidable representation. Note that qPCTL (and PCTL in general) does not permit conjunctions of until path formulae, which makes construction of such formula a bit less intuitive:

Lemma 6.1.2. *Let (S, P, v) be a Markov chain, $s \in S$ one of its states, and $x \in type$ a type. Then there exists a qPCTL formula λ such that $s \models \lambda$ if and only if there exists $p \in Paths_{fin}(s)$ such that $p \rightsquigarrow x$.*

Proof. We want λ to describe some hypothetical p such that $p \rightsquigarrow x$. Then p must have $type(p) = x$. We have already argued that having $type(p) = x$ amounts to satisfying every $\Phi \in X$ and $\mathrm{G}\,always$ for

$$X = x_1 \cup \{\neg\psi \,\mathrm{U}(\neg\psi \wedge \neg\varphi) \mid \varphi \,\mathrm{U}\, \psi \in x_2\},$$

$$always \equiv \bigwedge_{\varphi \,\mathrm{U}\, \psi \in x_3} (\varphi \wedge \neg\psi).$$

Since we do not a priori know in which order are $\Phi \in X$ satisfied, λ needs to account for all possible orders. Define *ord* to be the set of all (strict total) orders on formulae of X. Then

$$\lambda \equiv \bigvee_{o \in ord} \lambda_o,$$

where each λ_o ensures that $\Phi \in X$ are satisfied in the order o and $p \models \mathrm{G}\, always$. Moreover, λ_o also requires that last state of p satisfies $\mathrm{G}^{>0}\, always$, which gives us $p \rightsquigarrow x$.

It might be possible that p satisfies more than one $\Phi \in X$ in the same state. We could reflect this in the construction λ, but it is not necessary: p will satisfy λ_o for each order o that resolved such ties in an arbitrary way.

We will describe construction of one specific λ_o. Rename each $\Phi \in X$ to reflect the order o. Set $\alpha_i \,\mathrm{U}\, \beta_i \equiv \Phi$ if $\Phi \in X$ is the i-th formula satisfied in o. We start counting at 1 and in total have $n = |X|$ formulae of form $\alpha_i \,\mathrm{U}\, \beta_i$. Set $\lambda_o \equiv \lambda_1$, and for each $i \in \mathbb{N}$ such that $1 \leq i < n$, define

$$\lambda_i \equiv (always \wedge \bigwedge_{i \leq k \leq n} \alpha_k \wedge \neg\beta_k) \,\mathrm{U}^{>0} (\beta_i \wedge \lambda_{i+1}),$$

and

$$\lambda_n \equiv (always \wedge \alpha_n \wedge \neg\beta_n) \,\mathrm{U}^{>0} (\beta_n \wedge \mathrm{G}^{>0}\, always).$$

Intuitively, it is not hard to see how λ_o expresses existence of a path p such that $p \rightsquigarrow x$ and p satisfies $\alpha_i \,\mathrm{U}\, \beta_i$ in the order o. Consider λ_1. If it is satisfied, then there exists $p \in Paths_{fin}$ that at some point satisfies β_1, and until then each of its states must satisfy *always* and also $\bigwedge_{1 \leq k \leq n} \alpha_k \wedge \neg\beta_k$. Satisfaction of $\bigwedge_{1 \leq k \leq n} \alpha_k \wedge \neg\beta_k$ ensures that all of $\alpha_k \,\mathrm{U}\, \beta_k$ can still be satisfied later, but are not satisfied sooner than $\alpha_1 \,\mathrm{U}\, \beta_1$. The state that satisfies β_1 must also satisfy λ_2. Again, λ_2 requires satisfaction of *always*, but now $\bigwedge_{1 \leq k \leq n} \alpha_k \wedge \neg\beta_k$ is relaxed to $\bigwedge_{2 \leq k \leq n} \alpha_k \wedge \neg\beta_k$ because $\alpha_1 \,\mathrm{U}\, \beta_1$ has already been satisfied and we do not care about subsequent satisfactions of $\alpha_1 \,\mathrm{U}\, \beta_1$. This continues until β_2 is satisfied, at which point λ_3 must be satisfied, and so on. At some point, all $\alpha_i \,\mathrm{U}\, \beta_i$ have been satisfied, and λ_n requires the state that satisfied β_n to also satisfy $\mathrm{G}^{>0}\, always$.

Before we continue, we need to clarify what it means for p to satisfy each $\alpha_i \,\mathrm{U}\, \beta_i$ in the order o. Define $min(p, \alpha_i \,\mathrm{U}\, \beta_i)$ to be the least k such that $p(0)p(1)...p(k) \models \alpha_i \,\mathrm{U}\, \beta_i$, if it exists. Then p satisfies each $\alpha_i \,\mathrm{U}\, \beta_i$ in the order o if and only $min(p, \alpha_i \,\mathrm{U}\, \beta_i)$ exists for every $1 \leq i \leq n$ and we can order $\alpha_i \,\mathrm{U}\, \beta_i$ by $min(p, \alpha_i \,\mathrm{U}\, \beta_i)$ and resolve all ties in a way such that the resulting order is o.

Formally, we want to prove that for every $s \in S$, it holds that $s \models \lambda_o$ if and only if there exists $p \in Paths_{fin}(s)$ such that $p \models \mathrm{G}\, always$ and $p(len(p) - 1) \models \mathrm{G}^{>0}\, always$ and p satisfies each $\alpha_i \,\mathrm{U}\, \beta_i$ in the order o. A proof of this is omitted, but can be done with a straightforward induction on k: for $k \in \mathbb{N}$ and $s \in S$, we have $s \models \lambda_{n-k}$ if and only if there exists $p \in Paths_{fin}(s)$ such that:

- $p \models \mathrm{G}\, always$,
- $p(len(p) - 1) \models \mathrm{G}^{>0}\, always$,
- $p \models \alpha_i \,\mathrm{U}\, \beta_i$ for all $n - k \leq i \leq n$, and this happens in an order that is a restriction of o to formulae $\{\alpha_i \,\mathrm{U}\, \beta_i \mid n - k \leq i \leq n\}$.

The lemma statement then follows easily.

□

6.2 Three model-preserving qPCTL constructions

In this section, we will introduce three constructions that preserve satisfaction of qPCTL formulae. We know that existence of a finite path p such that $p \rightsquigarrow x$ for a type x can be described with a qPCTL formula, and we will use these constructions to establish additional properties of such paths.

Convention. Let S, I be a sets. In the following two lemmas, we use sets of form $\mathfrak{S} = \{s_i \mid s \in S, i \in I\}$. Each s_i formally corresponds to the ordered pair (s, i), but we use s_i for convenience.

Convention. Let $\mathfrak{S}$ be as in the previous convention. When we say $s_i \in \mathfrak{S}$, we mean an s_i such that $s \in S$ and $i \in I$, unless s or i was already set to a specific state or index. For example if $3 \in I$, then $s_3 \in \mathfrak{S}$ refers to an element of $\{s_3 \mid s \in S\}$.

The first construction is used in finite Markov chains to remove cycles shorter than n for an arbitrary $n \in \mathbb{N}$.

Lemma 6.2.1 (Construction 1). *Let $M = (S, P, v)$ be a finite Markov chain. Let $n \in \mathbb{N}$. Consider the following Markov chain $M^* = (\mathfrak{S}, P^*, v^*)$:*

- $\mathfrak{S} = \{s_i \mid s \in S$ *and* $i \in \mathbb{N}$ *such that* $0 \leq i < n\}$,
- $v^*(s_i) = v(s)$ *for every* $s_i \in \mathfrak{S}$,
- P^* *is defined as follows: for each* $s_i \in \mathfrak{S}$, *the successors of* s_i *are exactly* $\{t_k \mid t \in s^\bullet$ *and* $j = i + 1 \pmod n\}$. *For every successor* t_k *of* s_i, *we set* $P^*(s_i, t_k) = P(s, t)$.

Let ψ be a qPCTL formula. Then $M^*, s_i \models^{v^*} \varphi$ if and only if $M, s \models^v \varphi$ for every $\varphi \in sub(\psi)$ and $s_i \in \mathfrak{S}$.

Proof sketch. It is not hard to see that M^* is a Markov chain. We use structural induction on formulae. The case for $\mathrm{U}^{=1}$ formulae follows from Lemma 2.5.3 and a straightforward argument about BSCCs in M^*, the other cases follow immediately. □

The second construction, in essence, does the following: for some $n \in \mathbb{N}$ and some Markov chain and its state u, create a new Markov chain by taking a disjoint union of the original one n times, and then "merge" all copies of u back into a single state in a way that preserves all qPCTL formulae.

Lemma 6.2.2 (Construction 2). *Let $M = (S, P, v)$ be a Markov chain that is either finite (not necessarily a tree) or an infinite tree, and a state $u \in S$. If M is an infinite tree, it is rooted in u. Consider $M^* = (\mathfrak{S}, P^*, v^*)$ defined as follows (we assume states of form s_i are not u):*

- $\mathfrak{S} = \{u\} \cup \{s_i \mid s \in S, i \in \mathbb{N}$ *such that* $s \neq u$ *and* $0 \leq i < n\}$,
- $v^*(u) = v(u)$ *and* $v^*(s_i) = v(s)$ *for each* $s_i \in \mathfrak{S}$,
- P^* *is defined as follows:*
  - $P^*(s_i, t_i) = P(s, t)$ *for each* $s_i, t_i \in \mathfrak{S}$,
  - $P^*(s_i, u) = P(s, u)$ *for each* $s_i \in \mathfrak{S}$,

- $P^*(u, s_i) = \frac{P(u,s)}{n}$ *for each* $s_i \in \mathfrak{S}$,
- $P^*(u, u) = P(u, u)$.

Let ψ be a qPCTL formula. Then for every $\varphi \in sub(\psi)$ and every $s_i \in \mathfrak{S}$, we have $M^*, s_i \models^{v^*} \varphi$ if and only if $M, s \models^{v} \varphi$, and the analogous statement also holds for u.

Proof. Again, we assume $s_i \in \mathfrak{S}$ is never u, and explicitly say u when we mean u. We first need to show that M^* is a Markov chain. The probability distribution on successors of every $s_i \in \mathfrak{S}$ mimics the one s had on its successors in M, except it uses t_i instead of t and the probability of going into a hypothetical "local" u_i is "rerouted" to a single "global" u. Then $\sum_{t \in \mathfrak{S}} P^*(s_i, t) = 1$ for every $s_i \in \mathfrak{S}$. As for u, we have $\sum_{t \in \mathfrak{S}} P^*(u, t) = \sum_{s_i \in \mathfrak{S}} P^*(u, s_i) + P^*(u, u)$. Partitioning $s_i \in \mathfrak{S}$ by their specific index i, we see that

$$\sum_{s_i \in \mathfrak{S}} P^*(u, s_i) = \sum_{s_0 \in \mathfrak{S}} \frac{P(u,s)}{n} + \sum_{s_1 \in \mathfrak{S}} \frac{P(u,s)}{n} + \ldots + \sum_{s_{n-1} \in \mathfrak{S}} \frac{P(u,s)}{n}.$$

We use $s_0 \in \mathfrak{S}$ to refer to elements of $\{s_0 \mid s \in S \text{ and } s \neq u\}$, which means that $\sum_{s_0 \in \mathfrak{S}} \frac{P(u,s)}{n}$ is equal to $\sum_{s \in S \setminus \{u\}} \frac{P(u,s)}{n}$. This holds for the other summands as well, so $\sum_{s_i \in \mathfrak{S}} P^*(u, s_i) = n \cdot \sum_{s \in S \setminus \{u\}} \frac{P(u,s)}{n}$, which is just $\sum_{s \in S \setminus \{u\}} P(u, s)$. We also know $P^*(u, u) = P(u, u)$, so we may conclude

$$\sum_{t \in \mathfrak{S}} P^*(u, t) = \sum_{s_i \in \mathfrak{S}} P^*(u, s_i) + P^*(u, u) = \sum_{t \in S \setminus \{u\}} P(u, t) + P(u, u) = 1.$$

We use induction on subformulae to prove the lemma statement. We split the proof into two cases:

- M is an infinite tree rooted in u. The statement follows immediately for $s_i \in \mathfrak{S}$, since if we take the subtree of M^* induced by s_i and the subtree of M induced by s, we can obtain one from the other simply by adding or dropping the subscript i from its states (more formally, we can show that s and s_i are equivalent under probabilistic bisimulation, so they are indistinguishable by any PCTL formula, namely a qPCTL formula).

We use that to prove the lemma for u. All cases follow easily, except for $\varphi_1 \mathrm{U}^{=1} \varphi_2$. In that case, we use Lemma 2.5.3. Either direction follows immediately if case (i) applies. Now consider $\implies$ if case (ii) applies. Then $\mathbb{P}(A^*) > 0$ where A^* is the set $\{\pi \in Run^{M^*}(u) \mid \pi(i) \models \varphi_1 \wedge \neg\varphi_2 \text{ for } i \in \mathbb{N}\}$. Partition A^* by the second states of runs in it, which can never be u because M^* is also a tree rooted in u. Then $\mathbb{P}(A^* \cap Run^{M^*}(ut_i)) > 0$ for some $t_i \in \mathfrak{S}$ by countable additivity and the fact that $\mathbb{P}(A^*) > 0$. Given that $A^* \cap Run^{M^*}(ut_i) \subseteq Run^{M^*}(ut_i)$, we use Lemma 6.0.1 to show $\mathbb{P}(A^* \cap Run^{M^*}(ut_i))$ is equal to

$$\mathbb{P}(Run^{M^*}(ut_i)) \cdot \mathbb{P}(\{t_i w \mid ut_i w \in A^* \cap Run^{M^*}(ut_i)\}).$$

It follows that both factors must be positive. Now consider the set $A = \{\pi \in Run^M(u) \mid \pi(i) \models \varphi_1 \wedge \neg\varphi_2 \text{ for } i \in \mathbb{N}\}$ in M. Since t_i is a successor of u in M^*, we know t is a successor of u in M by the definition of P^*. We will show $\mathbb{P}(A \cap Run^M(ut)) > 0$, which implies $\mathbb{P}(A) > 0$ since $A \cap Run^M(ut) \subseteq A$. By Lemma 6.0.1, we have $\mathbb{P}(A \cap Run^M(ut))$ equal to

$$\mathbb{P}(Run^M(ut)) \cdot \mathbb{P}(\{tw \mid utw \in A \cap Run^M(ut)\}).$$

We already know $\mathbb{P}(Run^M(ut)) > 0$. It is easy to see that both $\{tw \mid utw \in A \cap Run^M(ut)\}$ and the corresponding set in M^* are simply sets of runs that start in t and t_i (respectively) that satisfy $G(\varphi_1 \wedge \neg\varphi_2)$. Then their measures are equal to each other by our argument about subtrees induced by t_i and t, and we can conclude $\mathbb{P}(A) > 0$.

The other direction is done analogously.

- M is finite. The case for $\mathrm{U}^{=1}$ formulae follows from Lemma 2.5.3 and straightforward arguments about finite paths and BSCCs, the other cases are not hard to see.

□

The third construction is used to modify probability distribution on successors of a single state, but without altering its successor set. We do this for finite Markov chains and also Markov chains that are

infinite trees. For finite Markov chains, the construction works because satisfaction of qPCTL formulae does not depend on the exact probability distributions on successors of states, provided successor sets remain intact. This is generally not the case for infinite Markov chains. For example, if we modified the Markov chain shown on Figure 5.1 to use uniform probability distribution on every state that satisfies $\neg a$, it would no longer be a model of $\mathrm{G}^{>0}(\neg a \wedge \mathrm{F}^{>0}\, a)$. However, in our specific case of infinite trees where we only change a single state without changing its successor set, the construction still works.

Lemma 6.2.3 (Construction 3)**.** *Let $M = (S, P, v)$ be a Markov chain which is either finite (not necessarily a tree) or an infinite tree. Let $u \in S$. Consider a Markov chain $M^* = (S, P^*, v)$ such that for every $s, t \in S$ we have*

- *$P^*(s, t) > 0$ if and only if $P(s, t) > 0$,*
- *$P^*(s, t) = P(s, t)$ whenever $s \neq u$.*

Let ψ be a qPCTL formula. Then $M, s \models \varphi$ if and only if $M^*, s \models \varphi$ for every $\varphi \in sub(\psi)$ and $s \in S$.

Proof. We use structural induction on formulae. All cases follow immediately, except for subformulae of form $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2$. We will prove $M, s \not\models \varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \iff M^*, s \not\models \varphi_1 \,\mathrm{U}^{=1}\, \varphi_2$. Both directions are done analogously (observe how role of M and M^* is "symmetric"), and we show $\implies$. We use the induction hypothesis and the fact that successor sets did not change for any state. If M is finite, a simple application of Lemma 2.5.3 is sufficient. If M was an infinite tree, apply Lemma 2.5.3. If case (i) applies, the statement follows immediately. Otherwise we get case (ii) and $\mathbb{P}_s(A) > 0$, where A is the set $\{\pi \in Run(s) \mid \pi(n) \models \varphi_1 \wedge \neg\varphi_2 \text{ for all } n \in \mathbb{N}\}$. We will consider the same set in M^* as well. Clearly $Run^M(s) = Run^{M^*}(s)$. It follows that A contains the same runs in both M and M^* by the induction hypothesis. We want to show A has a positive measure in M^*.

If there is no way to reach u from s, the "subtree" induced by s in M and M^* is exactly the same. It follows that A must have the same (positive) measure in M^*. Otherwise there exists some finite path π that starts with s and ends with u. This path is unique because M is

an infinite tree. Set $n = len(\pi) + 1$. Consider the family of sets of runs $\{A \cap Run(p) \mid p \in S^n\}$. It is easy to see that this family partitions A, so by countable additivity we get

$$\mathbb{P}(A) = \mathbb{P}(\bigcup_{p \in S^n} A \cap Run(p)) = \sum_{p \in S^n} \mathbb{P}(A \cap Run(p)).$$

Since $\mathbb{P}(A) > 0$, we have $\mathbb{P}(A \cap Run(p)) > 0$ for at least one $p \in S^n$. Fix such p. We know that $p = s_1 s_2 \ldots s_n$ where each $s_i \in S$ and $s_1 = s$ and s_{n-1} is possibly u. No other s_i can be u because M is an infinite tree and we know that the only path that reaches u from s has $n - 1$ states. Clearly $A \cap Run(p) \subseteq A$, so by Lemma 6.0.1 we have

$$\mathbb{P}(A \cap Run(p)) = \mathbb{P}(Run(p)) \cdot \mathbb{P}(\{s_n w \mid pw \in A \cap Run(p)\}).$$

Since $\mathbb{P}(A \cap Run(p)) > 0$, both factors must be positive. Since there is no way to reach u from s_n, the "subtree" induced by s_n in M and M^* is exactly the same. It follows that $\mathbb{P}(\{s_n w \mid pw \in A \cap Run(p)\}$ must also be positive in M^*. If $s_{n-1} \neq u$, then $\mathbb{P}(Run(p))$ in M^* is the same as in M. Otherwise $s_{n-1} = u$ and $\mathbb{P}(Run(p))$ in M^* may be different, but is still positive because $P(s_{n-1}, s_n) > 0$ implies $P^*(s_{n-1}, s_n) > 0$, and no other transition probability was modified. Then $\mathbb{P}(A \cap Run(p)) > 0$ in M^*, which implies $\mathbb{P}(A) > 0$ in M^*. $\square$

6.3 Pseudo-models for qPCTL formulae

In this section, we introduce *pseudo-models* for qPCTL formulae, and make some observations about Markov chains they induce. Definitions and results in this section either come directly from [8] or are a simple consequence. Due to space constraints, we restrict ourselves to mere enumeration of the concepts we need and do not explain them. This may be, to some extent, mitigated by Subsection 3.3.3 where we give a high-level overview of how methods in [8] are used.

The following definitions are needed to define pseudo-models:

Definition 6.3.1. The closure of a qPCTL state formula ψ is denoted by $Cl(\psi)$ and is defined as the least set C such that $\psi \in C$ and the following conditions are satisfied:

1. $\varphi \in C$ implies $\widehat{\varphi} \in C$,
2. $\varphi_1 \wedge \varphi_2 \in C$ implies $\varphi_1, \varphi_2 \in C$,
3. $\mathrm{X}^{>0}\, \varphi \in C$ implies $\varphi \in C$,
4. $\mathrm{X}^{=1}\, \varphi \in C$ implies $\varphi \in C$,
5. $\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2 \in C$ implies $\varphi_1, \varphi_2, \mathrm{X}^{>0}(\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2) \in C$,
6. $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \in C$ implies $\varphi_1, \varphi_2, \mathrm{X}^{=1}(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in C$,
7. $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \in C$ implies $\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2 \in C$.

Definition 6.3.2. A set $S \subseteq \psi$ is called eligible if and only if for every $\varphi \in Cl(\psi)$ we have that φ or $\widehat{\varphi}$ is in S, and the following conditions hold:

- if $\varphi \in S$, then $\widehat{\varphi} \notin S$,
- if $\varphi_1 \wedge \varphi_2 \in S$, then $\varphi_1, \varphi_2 \in S$,
- if $\neg(\varphi_1 \wedge \varphi_2) \in S$, then $\widehat{\varphi_1} \in S$ or $\widehat{\varphi_2} \in S$,
- if $\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2 \in S$, then $\varphi_2 \in S$ or $\varphi_1, \mathrm{X}^{>0}(\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2) \in S$,
- if $\neg(\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2) \in S$, then we have $\widehat{\varphi_2}, \neg\, \mathrm{X}^{>0}(\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2) \in S$ or $\widehat{\varphi_1}, \widehat{\varphi_2} \in S$,
- if $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \in S$, then $\varphi_2 \in S$ or $\varphi_1, \mathrm{X}^{=1}(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in S$,
- if $\neg(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in S$, then we have $\widehat{\varphi_2}, \neg\, \mathrm{X}^{=1}(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in S$ or $\widehat{\varphi_1}, \widehat{\varphi_2} \in S$,

Definition 6.3.3. A pseudo-structure for a qPCTL formula ψ is defined to be $\mathcal{A} = (A, \rightarrow)$, where A is a set of eligible subsets of $Cl(\psi)$ and $\rightarrow \subseteq A \times A$ is a total relation (i.e. for every $a \in A$ there exists $a' \in A$ such that $a \rightarrow a'$).

Every pseudo-structure $\mathcal{A} = (A, \rightarrow)$ for a formula ψ induces a unique Markov chain $M_{\mathcal{A}} = (A, P)$ such that $a^{\bullet} = \{a' \in A \mid a \rightarrow a'\}$ for every $a \in A$, and P assigns a uniform probability distribution to each state. Moreover, to each $\varphi \in Cl(\psi)$ we associate a fresh atomic proposition $[\varphi]$ and define a valuation v on A such that $[\varphi] \in v(S)$ if and only if $\varphi \in S$. For convenience, we write:

- $A, S \models \mathrm{X}^{\oplus q}\, \varphi$ instead of $M_{\mathcal{A}}, S \models^{v} \mathrm{X}^{\oplus q}[\varphi]$,
- $A, S \models \neg\, \mathrm{X}^{\oplus q}\, \varphi$ instead of $M_{\mathcal{A}}, S \models^{v} \neg\, \mathrm{X}^{\oplus q}[\varphi]$,
- $A, S \models \varphi_1 \,\mathrm{U}^{\oplus q}\, \varphi_2$ instead of $M_{\mathcal{A}}, S \models^{v} [\varphi_1]\, \mathrm{U}^{\oplus q}[\varphi_2]$,
- $A, S \models \neg(\varphi_1 \,\mathrm{U}^{\oplus q}\, \varphi_2)$ instead of $M_{\mathcal{A}}, S \models^{v} \neg([\varphi_1]\, \mathrm{U}^{\oplus q}[\varphi_2])$,

where $S \in A$ and $\oplus q$ is >0 or $=1$.

Definition 6.3.4. Let ψ be a qPCTL formula and $\mathcal{A} = (A, \rightarrow)$ a pseudo-structure for ψ. A witness for a formula $\neg(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in Cl(\psi)$ in $\mathcal{A}$ is a pseudo-structure $\mathcal{B} = (B, \hookrightarrow)$ where $\emptyset \neq B \subseteq A$ and $\hookrightarrow \subseteq \rightarrow$ such that:

- $\mathcal{B}$ is strongly connected,
- For every $S \in B$ we have $\widehat{\varphi_2} \in S$,
- For every $S \in B$ and $\xi_1 \,\mathrm{U}^{=1}\, \xi_2 \in S$, we have $\mathcal{B}, S \models \xi_1 \,\mathrm{U}^{=1}\, \xi_2$.

Definition 6.3.5. A pseudo-model for a qPCTL formula ψ is a pseudo-structure $\mathcal{A} = (A, \rightarrow)$ for ψ such that $\psi \in T$ for some $T \in A$, and every $S \in A$ satisfies the following:

1. If $\xi \in S$ for some ξ of form $\mathrm{X}^{>0}\, \varphi_1$, $\mathrm{X}^{=1}\, \varphi_1$, $\neg\, \mathrm{X}^{>0}\, \varphi_1$, $\neg\, \mathrm{X}^{=1}\, \varphi_1$, $\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2$, $\neg(\varphi_1 \,\mathrm{U}^{>0}\, \varphi_2)$, or $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2$, then $\mathcal{A}, S \models \xi$.
2. If $\neg(\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2) \in S$, then one of the following holds:
   (a) $\mathcal{A}, S \not\models \varphi_1 \,\mathrm{U}^{=1}\, \varphi_2$,

(b) There is a witness $\mathcal{B} = (B, \hookrightarrow)$ for $\neg(\varphi_1 \mathrm{U}^{=1} \varphi_2)$ and a finite path[1] $S_0 \to S_1 \to \ldots \to S_n$ such that $S_0 = S$, $S_n \in B$, and $\widehat{\varphi_2} \in S_i$ for every $0 \leq i \leq n$.

A pseudo-model is *simple* if condition (2) is always satisfied by (a), that is, no witnesses were used.

Definition 6.3.6. Let S be a set. Define S^+ to be the set $S^* \setminus \{\epsilon\}$.

Definition 6.3.7. A marked graph is a triple $\mathcal{G} = (G, \hookrightarrow, L)$, where G is a finite set of nodes, $\hookrightarrow \subseteq G \times G$ is a total relation, and $L \subseteq \hookrightarrow$ is a set of marked transitions.

Every marked graph $\mathcal{G} = (G, \hookrightarrow, L)$ induces a unique Markov chain $M_{\mathcal{G}} = (G^+, P)$ where

- for every $w \in G^*$ and $v \in G$, $\bar{w} \in wv^\bullet$ if and only if $\bar{w} = wvu$ for some $u \in G$ such that $v \hookrightarrow u$. A transition from wv to wvu is marked if and only if $v \hookrightarrow u$ is marked.

- for every $w \in G^+$, it holds that P assigns a uniform probability distribution to successors of w if none or all of the "outgoing" transitions are marked. Otherwise, P assigns the same probability p' to all non-marked transitions, and the same probability p to all marked transitions, in a way that $\sum_{w \hookrightarrow w' \in L} p = 1 - \frac{1}{4^{len(w)}}$.

Whenever $s \in G$ and ψ is a qPCTL formula, we write $\mathcal{G}, s \models \psi$ instead of $M_{\mathcal{G}}, s \models \psi$ (valuations are given by context).

Theorem 6.3.1. *In [8], the following is shown for every qPCTL formula ψ: there exists a (simple) pseudo-model for ψ if and only if ψ is (finitely) satisfiable. Every (simple) pseudo-model can be turned into a marked graph, which induces a (finite) model of ψ. Finally, it is shown that existence of (simple) pseudo-models is decidable, which implies decidability of (finite) satisfiability of qPCTL formulae.*

1. We have defined paths in Markov chains only, so here we rely on the standard definition from graph theory. We do this whenever we talk about paths in psedo-models.

We will not show the exact construction of marked graphs from pseudo-models, but make some observations that are either explicitly mentioned or follow directly from [8]. Let $\mathcal{A} = (A, \rightarrow)$ be a pseudo-model for ψ. Then we have the following:

- If $\mathcal{A}$ was simple, the marked graph used is $(A, \rightarrow, \emptyset)$. It is then shown that $M_{\mathcal{A}}, S \models \varphi$ if and only if $\varphi \in S$ for every $\varphi \in Cl(\psi)$ and $S \in A$.

- If $\mathcal{A}$ was not simple, then every state of the constructed marked graph $\mathcal{G} = (G, \hookrightarrow, L)$ is of form (S, i) where $S \in A$ and $i \in \mathbb{N}$. It is shown that for every $(S, i) \in G$ and $\varphi \in Cl(\psi)$, it holds that $M_{\mathcal{G}}, (S, i) \models \varphi$ if and only if $\varphi \in S$. Moreover, $\mathcal{G}$ is constructed in a way that there is a state $s \in M_{\mathcal{G}}$ that satisfies ψ but does *not* correspond to a state of any witness (specifically $s = (S, 0)$ for $S \in A$ such that $\psi \in S$).

In the previous section, we defined the property $p \rightsquigarrow x$ for a type x and a finite path p, and it requires that $p(len(p) - 1)$ satisfies a formula of form $\mathrm{G}^{>0}\, \varphi$, where φ is a qPCTL formula. Formally, $\mathrm{G}^{>0}\, \varphi$ is the formula $\mathrm{F}^{<1}\, \neg\varphi$, and that is $\top\, \mathrm{U}^{<1}\, \neg\varphi$. Our convention about qPCTL formulae means that we use that formula only in form $\neg(\top\, \mathrm{U}^{=1}\, \neg\varphi)$. Therefore, we use $\mathrm{G}^{>0}\, \varphi$ as a shortcut for $\neg(\top\, \mathrm{U}^{=1}\, \neg\varphi)$. Analogously, $G^{=1}\varphi$ is a shortcut for $\neg(\top\, \mathrm{U}^{>0}\, \neg\varphi)$.

We now establish the first observation we will use:

Lemma 6.3.2 (Observation 1)**.** *Let $\mathcal{A} = (A, \rightarrow)$ be a pseudo-model for a qPCTL formula ψ, and $\mathcal{G} = (G, \hookrightarrow, L)$ the corresponding marked graph. Suppose $s \in M_{\mathcal{G}}$ and $M_{\mathcal{G}}, s \models \mathrm{G}^{>0}\, \varphi$ for some $\mathrm{G}^{>0}\, \varphi \in Cl(\psi)$. Then there exists a finite path $s_0 s_1 \ldots s_n$ in $M_{\mathcal{G}}$ such that $s_0 = s$, and $M_{\mathcal{G}}, s_i \models \varphi$ for $0 \leq i \leq n$, and one of the following holds:*

(i) $M_{\mathcal{G}}, s_n \models G^{=1}\varphi$,

(ii) *s_n corresponds to a state of a witness for $\mathrm{G}^{>0}\, \varphi$.*

Proof. Since $\mathrm{G}^{>0}\, \varphi$ is formally $\neg(\top\, \mathrm{U}^{=1}\, \neg\varphi)$, we know $M_{\mathcal{G}}, s \models \mathrm{G}^{>0}\, \varphi$ stands for $M_{\mathcal{G}}, s \models \neg(\top\, \mathrm{U}^{=1}\, \neg\varphi)$. Then $M_{\mathcal{G}}, s \not\models \top\, \mathrm{U}^{=1}\, \neg\varphi$. The state s_0 corresponds to some $(S, i) \in G$ in the marked graph $\mathcal{G}$, and we

have $M_{\mathcal{G}}, (S, i) \not\models \top \,\mathrm{U}^{=1}\, \neg\varphi$. Since $\neg(\top \,\mathrm{U}^{=1}\, \neg\varphi) \in Cl(\psi)$ implies $\top \,\mathrm{U}^{=1}\, \neg\varphi \in Cl(\psi)$, we have $\top \,\mathrm{U}^{=1}\, \neg\varphi \notin S$. The proof now continues by following the proof of Theorem 3.10 from [8], specifically the $\Longleftarrow$ direction for $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2$. Note that in [8], this is used to show that $\varphi_1 \,\mathrm{U}^{=1}\, \varphi_2 \notin S$ implies $\mathcal{G}, (S, i) \not\models \varphi_1 \,\mathrm{U}^{=1}\, \varphi_2$ for every $(S, i) \in G$, which in our specific case we already know (and in this proof we already used the whole theorem), but the way it is established also gives a proof for this lemma. □

We follow [8] and make the following remark: consider a marked graph $\mathcal{G} = (G, \hookrightarrow, L)$ and suppose $(S, i) \in G$ such that (S, i) belongs to some witness B. Then the measure of runs $\pi \in Run^{M_{\mathcal{G}}}((S, i))$ that stay in B (that is, all states they reach correspond to states of B) is at least $2/3$. This follows from the fact that "internal" transitions in witnesses are marked, and so the measure of runs $\pi \in Run^{M_{\mathcal{G}}}((S, i))$ such that π performs at least one unmarked transition is bounded by $\sum_{i=1}^{\infty} 1/4^i = 1/3$, and this can be seen from the definition of $M_{\mathcal{G}}$.

Now we state the second observation:

Lemma 6.3.3 (Observation 2)**.** *Let $\mathcal{G} = (G, \hookrightarrow, L)$ be a marked graph for ψ and suppose $(S, i) \in G$ is such that (S, i) belongs to some witness B. Then for every $\delta \in \mathbb{R}$ such that $0 < \delta < 1$, we can redefine $M_{\mathcal{G}}$ in a way that the measure of runs $\pi \in Run^{M_{\mathcal{G}}}((S, i))$ such that π performs at least one unmarked transition is bounded from above by δ, and the resulting modified $M_{\mathcal{G}}$ is still a model of ψ.*

Proof sketch. In the definition of the induced Markov chain $M_{\mathcal{G}}$, we set $\sum_{w \hookrightarrow w' \in L} p = 1 - 1/4^{len(w)}$. Replacing 4 in the previous definition by $x \in \mathbb{N}$ such that $x > 1$ changes the relevant upper bound to $\sum_{i=1}^{\infty} 1/x^i$, which is then equal to $1/(x-1)$. It is clear that we can choose x which makes $1/(x-1)$ sufficiently low. The proof that $M_{\mathcal{G}}$ is a model of ψ proceeds exactly the same way as it did previously. □

6.4 Decidability of finite and general satisfiability

In this section, we state the main results of this chapter. We will show that a formula ζ has a (finite) model if and only if there exists a (simple) *pre-model* for ζ.

We start with the definition of pre-models:

Definition 6.4.1 (Pre-model)**.** A (simple) pre-model for a formula ζ is given by $(T, \mathcal{A}, \{m_x\}_{x \in T})$ where

- $T \subseteq type$ is a set of types,
- $\mathcal{A} = (A, \rightarrow)$ is a (simple) pseudo-model for $\bigwedge_{x \in T} \lambda_x$,
- $\{m_x\}_{x \in T} \subseteq \mathbb{R}$ such that:
  1. $0 < m_x \leq 1$ for every $x \in T$,
  2. $\sum_{x \in T} m_x = 1$,
  3. for every $i \in I$ we have

$$l_i < \sum_{\substack{x \in T \\ \varphi_i \mathrm{U} \psi_i \in x_1}} m_x < r_i.$$

Recall the following definition, which we use frequently in the proofs that follow:

Definition 6.4.2 ($\mathbb{P}_s(x)$)**.** For a state s of some Markov chain and a type x, we use $\mathbb{P}_s(x)$ to denote $\mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\})$.

Definition 6.4.3 (λ_x)**.** We use λ_x to denote the formula λ for the type $x \in type$ we constructed in Lemma 6.1.2.

Theorem 6.4.1. *If ζ is a (finitely) satisfiable, then there exists (simple) pre-model for ζ.*

Proof. Suppose there exists a (finite) Markov chain (S, P, v) and $s \in S$ such that $s \models \zeta$. Define $T = \{x \in type \mid \mathbb{P}_s(x) > 0\}$. Note that T contains exactly all types that have positive measure, when considered from s. By Lemma 6.1.1, for every $x \in T$, there exists $p \in Paths_{fin}(s)$ such that $p \rightsquigarrow x$. By Lemma 6.1.2, we have $s \models \lambda_x$ for every $x \in T$. Therefore $s \models \bigwedge_{x \in T} \lambda_x$. We conclude that $\bigwedge_{x \in T} \lambda_x$ is satisfiable, and if (S, P, v) was finite, then it is also finitely-satisfiable. By Theorem 6.3.1 (which comes from [8]), there exists a pseudo-model $\mathcal{A} = (A, \rightarrow)$ for $\bigwedge_{x \in T} \lambda_x$. Moreover, $\mathcal{A}$ is simple if (S, P, v) was finite. This gives us T and $\mathcal{A}$ in the pre-model. What remains is to find suitable $\{m_x\}_{x \in T}$.

For every $x \in T$, define $m_x = \mathbb{P}_s(x)$. Clearly $0 < m_x \leq 1$ for every $x \in T$. Obviously each run $\pi \in Run(s)$ has a type, but also has only

a single type (namely $type(\pi)$). Then we can partition $Run(s)$ based on types of runs and use countable additivity of measures to get

$$\begin{aligned} 1 &= \mathbb{P}_s(Run(s)) \\ &= \mathbb{P}_s(\bigcup_{x \in type} \{\pi \in Run(s) \mid type(\pi) = x\}) \\ &= \sum_{x \in type} \mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\}) \\ &= \sum_{x \in type} \mathbb{P}_s(x). \end{aligned}$$

Since every $x \in type \setminus T$ has $\mathbb{P}_s(x) = 0$, we also have

$$\sum_{x \in type} \mathbb{P}_s(x) = \sum_{x \in T} \mathbb{P}_s(x).$$

Therefore,

$$1 = \mathbb{P}_s(Run(s)) = \ldots = \sum_{x \in type} \mathbb{P}_s(x) = \sum_{x \in T} \mathbb{P}_s(x) = \sum_{x \in T} m_x,$$

and the second condition for $\{m_x\}_{x \in T}$ is satisfied. It remains to show the third condition. Let $i \in I$. By partitioning runs in $\{s \models \varphi_i \,\mathrm{U}\, \psi_i\}$ by their types, we can show $\mathbb{P}_s(s \models \varphi_i \,\mathrm{U}\, \psi_i)$ is equal to

$$\mathbb{P}_s(\bigcup_{x \in type} \{\pi \in Run(s) \mid \pi \models \varphi_i \,\mathrm{U}\, \psi_i \text{ and } type(\pi) = x\}),$$

and by countable additivity that is equal to

$$\sum_{x \in type} \mathbb{P}_s(\{\pi \in Run(s) \mid \pi \models \varphi_i \,\mathrm{U}\, \psi_i \text{ and } type(\pi) = x\}).$$

Since $\mathbb{P}_s(x)$ for each $x \in types(\mathcal{F}) \setminus T$ is equal to 0, it follows that the sum above is equal to

$$\sum_{x \in T} \mathbb{P}_s(\{\pi \in Run(s) \mid \pi \models \varphi_i \,\mathrm{U}\, \psi_i \text{ and } type(\pi) = x\}).$$

A run π such that $type(\pi) = x$ satisfies $\varphi_i \,\mathrm{U}\, \psi_i \in \mathcal{F}$ if and only if $\varphi_i \,\mathrm{U}\, \psi_i \in x_1$ by the definition of types, so this further reduces to

$$\sum_{\substack{x \in T \\ \varphi_i \,\mathrm{U}\, \psi_i \in x_1}} \mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\}).$$

Together, we obtain

$$\mathbb{P}_s(s \models \varphi_i \mathrm{U} \psi_i) = \sum_{\substack{x \in T \\ \varphi_i \mathrm{U} \psi_i \in x_1}} \mathbb{P}_s(x) = \sum_{\substack{x \in T \\ \varphi_i \mathrm{U} \psi_i \in x_1}} m_x.$$

We know that $s \models P(\varphi_i \mathrm{U} \psi_i) \in (l_i, r_i)$, which means

$$l_i < \sum_{\substack{x \in T \\ \varphi_i \mathrm{U} \psi_i \in x_1}} m_x < r_i.$$

Therefore, the third condition for $\{m_x\}_{x \in T}$ is also satisfied, and we can conclude $(T, \mathcal{A}, \{m_x\}_{x \in T})$ is a pre-model for ζ. Moreover, if (S, P, v) was finite, then the pre-model is simple. □

Theorem 6.4.2. *If there exists a (simple) pre-model for ζ, then ζ is a (finitely) satisfiable.*

Proof. Suppose $(T, \mathcal{A}, \{m_x\}_{x \in T})$ is a (simple) pre-model for ζ. We will show ζ has a (finite) model. By its definition, $\mathcal{A}$ is a pseudo-model for $\bigwedge_{x \in T} \lambda_x$. By Theorem 6.3.1, this pseudo-model induces a model $M_{\mathcal{G}}$ of $\bigwedge_{x \in T} \lambda_x$. This model is finite if $\mathcal{A}$ was simple. We know there exists $s \in G$ such that $M_{\mathcal{G}}, s \models \bigwedge_{x \in T} \lambda_x$, and by our observation about marked graphs we may assume s does not correspond to a state of any witness. We redefine $M_{\mathcal{G}}$ to be the infinite tree induced by s. Formally, we restrict $M_{\mathcal{G}}$ to the tree given by $\pi \in G^+$ such that $\pi(0) = s$.

Before we continue, we recall that for a type x and a finite path p, $p \rightsquigarrow x$ implies its last state satisfies $\mathrm{G}^{>0}\, always$ where *always* is the formula $\bigwedge_{\varphi \mathrm{U} \psi \in x_3} (\varphi \wedge \neg \psi)$. We will use the formula *always* frequently, and assume that the corresponding x_3 is considered with respect to the type of the path.

By Lemma 6.1.2, for every $x \in T$, there exists $p_x \in Paths_{fin}(s)$ such that $p_x \rightsquigarrow x$. We extend each p_x with use of *Observation 1* (Lemma 6.3.2):

Path extension Consider p_x for $x \in T$. By the previous paragraph, we have $p_x(len(p_x) - 1) \models \mathrm{G}^{>0}\, always$. It is easy to see that $\mathrm{G}^{>0}\, always \in Cl(\bigwedge_{x \in T} \lambda_x)$. Then we use *Observation 1* to show that there exists a finite path $s_0 s_1 \ldots s_n$ such that $p_x(len(p_x) - 1)$ is s_0, $M_{\mathcal{G}}, s_i \models always$ for $0 \leq i \leq n$, and one of the following holds:

(i) $M_{\mathcal{G}}, s_n \models G^{=1} always$,

(ii) s_n corresponds to a state of a witness for $\mathrm{G}^{>0}\, always$.

Extend p_x by $s_1 s_2 \ldots s_n$. We want to show $p_x s_1 s_2 \ldots s_n \rightsquigarrow x$. By the definition of witnesses, if B is a witness for some $\mathrm{G}^{>0}\, \varphi$, then $\varphi \in S$ for every $S \in B$. It follows that *always* is satisfied in every state that corresponds to a state of the witness that contains s_n. Since the measure of runs that start in s_n and reach only states corresponding to the witness is at least $2/3$, clearly $s_n \models \mathrm{G}^{>0}\, always$. Since each s_i satisfies *always*, we can conclude $p_x s_1 s_2 \ldots s_n \rightsquigarrow x$. However, we now also have $s_n \models G^{=1} always$ or that s_n must correspond to a state of a witness for $\mathrm{G}^{>0}\, always$.

For convenenice, we redefine p_x to $p_x s_1 s_2 \ldots s_n$.

We repeat this procedure for each p_x and $x \in T$.

We now we have a collection of finite paths $\{p_x\}_{x \in T}$ such that every p_x starts with s, has the property $p_x \rightsquigarrow x$, and now also ends with a state that satisfies $G^{=1} always$ or corresponds to a state of a witness for $\mathrm{G}^{>0}\, always$.

We want to ensure that no p_x visits the same state twice. This is done by applying *Construction 1* (Lemma 6.2.1) from Section 6.2:

Construction 1 If $\mathcal{A}$ is not simple, then $M_{\mathcal{G}}$ is an infinite tree, and we do not have to do anything. Suppose $\mathcal{A}$ is simple. We know that every p_x starts in s. Since $\mathcal{A}$ was simple, no witnesses were used, which means each p_x ends with a state that satisfies $\mathrm{G}^{=1}\, always$ (for *always* with respect to x).

Since there are fintely many $x \in T$ and each p_x is finite, we can choose $n \in \mathbb{N}$ which is greater than maximum of their lengths and carry out *Construction 1* for n. Let M^* be the resulting Markov chain. Clearly M^* is finite. Consider some p_x for $x \in T$. Then M^* contains a path p'_x such that

$$p'_x(i) = (p_x(i), i \pmod{n})$$

for $0 \leq i \leq len(p) - 1$ and $len(p'_x) = len(p)$. We know that $n > len(p) - 1$ by its definition, which implies $i = i \pmod{n}$

for every $0 \leq i \leq len(p) - 1$. It follows that all states on this path are distinct. Given that this construction preserves qPCTL formulae, it is easy to see that p'_x satisfies $p'_x \rightsquigarrow x$ and its last state satisfies $\mathrm{G}^{=1}$ *always*.

For convenience, redefine each p_x to the corresponding p'_x in M^*. Also, redefine $M_{\mathcal{G}}$ to M^*. Note that $p'_x(0) = (s, 0)$ for every $x \in T$, so each p'_x starts in the same state. Set $s = (s, 0)$, so that we can again say all the considered paths start in s.

We may now assume that no path in $\{p_x\}_{x \in T}$ visits the same state twice, while the relevant properties are still satisfied.

Our next goal is to ensure that distinct paths from $\{p_x\}_{x \in T}$ share exactly only the first state. In other words, they are mutually disjoint except for their first state, and all of them start in the same state. We use *Construction 2* (Lemma 6.2.2) for $n = |T|$ and $u = s$, which essentially ensures that each p_x gets its own copy of $M_{\mathcal{G}}$, with the exception of the first state that remains shared:

Construction 2 Define $n = |T|$ and $u = s$. Note that $M_{\mathcal{G}}$ is either finite or an infinite tree rooted in u, so we can apply *Construction 2*. Call the resulting Markov chain M^*. Clearly M^* is finite if $M_{\mathcal{G}}$ was finite, and remains an infinite tree rooted in u if $M_{\mathcal{G}}$ was an infinite tree rooted in u. For each $x \in T$, pick $i \in \{0, 1, \ldots n-1\}$ in a way that no distinct $x, x' \in T$ have the same i. Consider some $x \in T$ and suppose we chose $i = k$ for x. Then M^* contains a path p' that corresponds to p_x, namely

$$p_x(0) \;\; p_x(1)_k \;\; p_x(2)_k \;\; \ldots \;\; p_x(len(p_x) - 1)_k.$$

Note that $p_x(0) = u$, which is why no subscript was used (or could have been used, since u is not copied). Also, adding the subscript k to the remaining states clearly results in a state of M^* if $M_{\mathcal{G}}$ was an infinite tree, but also when $M_{\mathcal{G}}$ was finite, because p_x has reached u as its first state and never reached it again by *Construction 1*. Therefore, p' is well-defined. Since all of states of p_x were distinct, all states of p' are distinct. Moreover, if we consider some other type $y \in T$ which necessarily has different i, the path that corresponds to p_y starts in u and shares no other states with p'.

Construction 2 preserves satisfaction of qPCTL formulae. If $M_{\mathcal{G}}$ was finite, then p' satisfies the relevant properties. Suppose $M_{\mathcal{G}}$ is an infinite tree. Preservation of qPCTL formulae clearly implies $p' \rightsquigarrow x$. Let v be the last state of p_x. If v satisfies $\mathrm{G}^{=1}$ *always*, then so does the last state of p'. Otherwise, v corresponds to a state of a witness for $\mathrm{G}^{>0}$ *always*. If $v = u$, then $v = s$, which contradicts our initial assumption that s does not correspond to a state of any witness. Then $v \neq u$. It follows that the last state of p' also corresponds to a state of a witness for $\mathrm{G}^{>0}$ *always*, since the subtree induced by v and the last state of p' is exactly the same Markov chain except for names of states. In conclusion, p' satisfies all the relevant properties even if $M_{\mathcal{G}}$ was an infinite tree.

Again, we redefine p_x to their corresponding p' in M^*, and also redefine $M_{\mathcal{G}}$ to M^*. We do not need to redefine s because all p' still start in it.

To summarize, we now have a Markov chain $M_{\mathcal{G}}$ that is finite if and only if $\mathcal{A}$ was simple and a collection of finite paths $\{p_x\}_{x \in T}$ such that:

- each p_x starts in s.
- each p_x has the property $p_x \rightsquigarrow x$.
- each p_x ends in a state that satisfies $G^{=1}$*always* or corresponds to a state of a witness for $\mathrm{G}^{>0}$ *always*.
- no p_x reaches a single state more than once.
- p_x and $p_{x'}$ start in s and share no other state for every distinct $x, x' \in T$.

We will now argue that *Construction 3* (Lemma 6.2.3) together with *Observation 2* (Lemma 6.3.3) can be used to obtain a model of ζ. We will not actually construct a model of ζ as the formal process would be unnecessarily long and technical, but instead argue a model can be constructed by modifying transition probabilities in $M_{\mathcal{G}}$. For the remainder of this proof, when we say *X is arbitrarily close to Y*, we

mean that for every $\delta \in \mathbb{R}$ such that $\delta > 0$, the construction could have been carried out in a way that $|X - Y| < \delta$.

We start with the probability distribution on successors of s. We cannot change $P(s, p_x(1))$ to m_x for each $x \in T$, because $\sum_{x \in T} m_x = 1$ and no probability would be "left" for other successors of s, which implies the successor set of s would not remain intact. However, we can change $P(s, p_x(1))$ so that it is lower but arbitrarily close to m_x. Notice that $\sum_{x \in T} m_x = 1$ implies that "approaching" each m_x from below by $P(s, p_x(1))$ gives $\sum_{x \in T} P(s, p_x(1)) < 1$, and so it follows that some probability must be "left" for other successors of s. We can simply choose uniform probability distribution for the other successors. Here we use that all $p_x(1)$ are distinct states, otherwise the modification would not be well-defined.

Since each p_x is a finite path and there are only finitely many types, we can change transition probabilities on successive states of p_x to be arbitrarily close to 1 (using 1 could alter some successor sets), except for the first two states s and $p_x(1)$, which we do not modify. Here we use that no state except for s is shared or visited more than once, otherwise our changes could conflict each other.

With these two modifications, we can get $\mathbb{P}_s(Run(p_x))$ arbitrarily close to m_x for each $x \in T$ (at the same time). We want to get each $\mathbb{P}_s(x)$ arbitrarily close to m_x.

First, we make the following observation. It is sufficient to consider arbitrarily high lower bounds for each $\mathbb{P}_s(x)$, because arbitrarily low upper bounds are implied by the fact that runs of different types partition $Run(s)$ and $\sum_{x \in T} m_x = 1$. This holds because if $\delta > 0$ is a positive constant and we ensure that $m_x - \delta \leq \mathbb{P}_s(x)$ for every $x \in T$, then also $\mathbb{P}_s(x) \leq m_x - \delta + (1 - \sum_{x \in T}(m_x - \delta))$ for every $x \in T$. Intuitively, $1 - \sum_{x \in T}(m_x - \delta)$ is the amount of "free" probability that we have not ensured belongs to a specific type, and so it gives an upper bound on how much some $\mathbb{P}_s(x)$ could have increased from the lower bound $m_x - \delta$ we have guaranteed. This argument would not work if $\sum_{x \in T} m_x$ was less than 1, because the upper bounds given would not be arbitrarily low.

Now we continue with the proof. If $M_{\mathcal{G}}$ is finite, then no extra steps are needed since $Run(p_x) \subseteq \{\pi \in Run(s) \mid type(\pi) = x\}$. If $M_{\mathcal{G}}$ is infinite, the situation is a bit more complicated. Let p_x for $x \in T$, and

set v to be its last state. If v satisfies $\mathrm{G}^{=1}\, always$, then we use Lemma 6.0.1. Otherwise v corresponds to a state of a witness for $\mathrm{G}^{>0}\, always$. By *Observation 2*, at the beginning of this construction we could have chosen some $0 < \delta < 1$ so that at this point the measure of runs $\pi \in Run(v)$ that stay in states corresponding to that witness is at least δ. Then $\mathbb{P}(v \models \mathrm{G}\, always)$ is at least δ, and so we can assume that it is arbitrarily close to 1. A simple application of Lemma 6.0.1 now yields that $\mathbb{P}_s(x)$ is arbitrarily close to m_x.

As a technical sidenote, our "a posteriori" use of *Observation 2* causes no problems, because our construction was not affected by it until now, and now we are merely saying that it was possible to pick a sufficiently high bound for runs that "start in a witness and perform no unmarked transitions" so that at this point we find ourselves with $\mathbb{P}(v \models \mathrm{G}\, always)$ that is as close to 1 as we want.

Finally, consider some $\varphi_i \,\mathrm{U}\, \psi_i$ such that $i \in I$. It is not hard to see that $\mathbb{P}_s(s \models \varphi_i \,\mathrm{U}\, \psi_i)$ is equal to

$$\sum_{\substack{x \in type \\ \varphi_i \,\mathrm{U}\, \psi_i \in x_1}} \mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\}).$$

Observe that we can partition $Run(s)$ based on types of runs and $\sum_{x \in T} \mathbb{P}_s(x)$ is arbitrarily close to $\sum_{x \in T} m_x = 1$. Then $\sum_{x \in type \setminus T} \mathbb{P}_s(x)$ is arbitrarily close to 0, which implies

$$\sum_{\substack{x \in type \setminus T \\ \varphi_i \,\mathrm{U}\, \psi_i \in x_1}} \mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\})$$

is arbitrarily close to 0. Then

$$\sum_{\substack{x \in type \\ \varphi_i \,\mathrm{U}\, \psi_i \in x_1}} \mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\})$$

is arbitrarily close to

$$\sum_{\substack{x \in T \\ \varphi_i \,\mathrm{U}\, \psi_i \in x_1}} \mathbb{P}_s(\{\pi \in Run(s) \mid type(\pi) = x\}),$$

which is arbitrarily close to

$$\sum_{\substack{x \in T \\ \varphi_i \mathrm{U} \psi_i \in x_1}} m_x.$$

Since $(T, \mathcal{A}, \{m_x\}_{x \in T})$ is a pre-model, we know

$$l_i < \sum_{\substack{x \in T \\ \varphi_i \mathrm{U} \psi_i \in x_1}} m_x < r_i,$$

which means we can ensure $P(\varphi_i \mathrm{U} \varphi_i) \in (l_i, r_i)$ is satisfied.

□

Theorem 6.4.3. *Existence of a (simple) pre-model for ζ is decidable.*

Proof sketch. For each $T \in \mathcal{P}(type)$:

1. Check if there exists a (simple) pseudo-model for $\bigwedge_{x \in T} \lambda_x$. This is decidable by results from [8]. If no such pseudo-model exists, continue with the next T.

2. Existence of $\{m_x\}_{x \in T}$ that satisfy the relevant constraints reduces to a SMT problem for a decidable theory. If such $\{m_x\}_{x \in T}$ exist, we have found a pre-model.

If no (simple) pre-model was found during this process, no (simple) pre-model exists. □

7 Conclusion

We have found two criteria that would imply decidability of the finite satisfiability problem for qualitative PCTL extended with bounded until formulae. A possible direction for future work is to attempt to prove they hold for finitely satisfiable formulae.

We have shown that bounded PCTL extended by allowing a single use of $G^{=1}$ does not have the small model property. A potential direction for future work is to consider different extensions or limited classes of models.

We have also shown that the finite and general satisfiability problems are decidable for a fragment of PCTL that consists of conjunctions of formulae of form $P(\varphi \mathrm{U} \psi) \in I$ where I is an open interval and φ, ψ are qualitative PCTL formulae. A natural direction is to extend this result to closed intervals. The authors focus on this topic in their current work.

Bibliography

[1] Hans Hansson and Bengt Jonsson. "A logic for reasoning about time and reliability". In: *Formal aspects of computing* 6.5 (1994), pp. 512–535.

[2] Edmund M Clarke and E Allen Emerson. "Design and synthesis of synchronization skeletons using branching time temporal logic". In: *Workshop on Logic of Programs*. Springer. 1981, pp. 52–71.

[3] Jan Kretínský and Alexej Rotar. "The satisfiability problem for unbounded fragments of probabilistic CTL". In: *arXiv preprint arXiv:1806.11418* (2018).

[4] Souymodip Chakraborty and Joost-Pieter Katoen. "On the satisfiability of some simple probabilistic logics". In: *2016 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE. 2016, pp. 1–10.

[5] Sergiu Hart and Micha Sharir. "Probabilistic propositional temporal logics". In: *Information and control* 70.2-3 (1986), pp. 97–155.

[6] Sarit Kraus and Daniel Lehmann. "Decision procedures for time and chance". In: *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*. IEEE. 1983, pp. 202–209.

[7] Daniel Lehmann and Saharon Shelah. "Reasoning with time and chance". In: *Information and Control* 53.3 (1982), pp. 165–198.

[8] Tomáš Brázdil et al. "The satisfiability problem for probabilistic CTL". In: *2008 23rd Annual IEEE Symposium on Logic in Computer Science*. IEEE. 2008, pp. 391–402.

[9] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT press, 2008.

[10] Robert B Ash and Catherine Doleans-Dade. *Probability and measure theory*. Academic Press, 2000.

[11] John G Kemeny, J Laurie Snell, and Anthony W Knapp. *Denumerable Markov chains: with a chapter of Markov random fields by David Griffeath*. Vol. 40. Springer Science & Business Media, 2012.

[12] E Allen Emerson and Joseph Y Halpern. "Decision procedures and expressiveness in the temporal logic of branching time". In: *Journal of computer and system sciences* 30.1 (1985), pp. 1–24.

[13] Edmund M. Clarke, E Allen Emerson, and A Prasad Sistla. "Automatic verification of finite-state concurrent systems using temporal logic specifications". In: *ACM Transactions on Programming Languages and Systems* (*TOPLAS*) 8.2 (1986), pp. 244–263.

[14] Behnam Banieqbal and Howard Barringer. "Temporal logic with fixed points". In: *Temporal Logic in Specification*. Springer. 1989, pp. 62–74.

[15] Michael J Fischer and Richard E Ladner. "Propositional dynamic logic of regular programs". In: *Journal of computer and system sciences* 18.2 (1979), pp. 194–211.

[16] E Allen Emerson, Charanjit S Jutla, and A Prasad Sistla. "On model checking for the μ-calculus and its fragments". In: *Theoretical Computer Science* 258.1-2 (2001), pp. 491–522.

[17] Dexter Kozen. "A finite model theorem for the propositional μ-calculus". In: *Studia Logica* 47.3 (1988), pp. 233–241.

[18] Orna Kupferman and Moshe Y Vardi. "An automata-theoretic approach to reasoning about infinite-state systems". In: *International Conference on Computer Aided Verification*. Springer. 2000, pp. 36–52.

[19] Thomas A Henzinger, Orna Kupferman, and Rupak Majumdar. "On the universal and existential fragments of the μ-calculus". In: *Theoretical Computer Science* 354.2 (2006), pp. 173–186.

[20] E Allen Emerson and Charanjit S Jutla. "The complexity of tree automata and logics of programs". In: [*Proceedings 1988*] *29th Annual Symposium on Foundations of Computer Science*. IEEE. 1988, pp. 328–337.

[21] E Allen Emerson and Chin-Laung Lei. "Modalities for model checking: Branching time logic strikes back". In: *Science of computer programming* 8.3 (1987), pp. 275–306.

[22] Orna Lichtenstein and Amir Pnueli. "Checking that finite state concurrent programs satisfy their linear specification". In: *Proceedings of the 12th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*. ACM. 1985, pp. 97–107.

[23] Costas Courcoubetis and Mihalis Yannakakis. "The complexity of probabilistic verification". In: *Journal of the ACM (JACM)* 42.4 (1995), pp. 857–907.

[24] Michael Huth and Marta Kwiatkowska. "Quantitative analysis and model checking". In: *Proceedings of Twelfth Annual IEEE Symposium on Logic in Computer Science*. IEEE. 1997, pp. 111–122.

[25] Kousha Etessami and Mihalis Yannakakis. "Recursive Markov chains, stochastic grammars, and monotone systems of nonlinear equations". In: *Journal of the ACM (JACM)* 56.1 (2009), p. 1.

[26] Javier Esparza, Antonín Kučera, and Richard Mayr. "Model checking probabilistic pushdown automata". In: *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.* IEEE. 2004, pp. 12–21.

[27] Tomáš Brázdil, Antonín Kučera, and Oldřich Stražovský. "On the decidability of temporal properties of probabilistic pushdown automata". In: *Annual Symposium on Theoretical Aspects of Computer Science*. Springer. 2005, pp. 145–157.

[28] Tomáš Brázdil, Vojtěch Forejt, and Antonín Kučera. "Controller synthesis and verification for Markov decision processes with qualitative branching time objectives". In: *International Colloquium on Automata, Languages, and Programming*. Springer. 2008, pp. 148–159.

[29] Tomáš Brázdil et al. "Stochastic games with branching-time winning objectives". In: *21st Annual IEEE Symposium on Logic in Computer Science (LICS'06)*. IEEE. 2006, pp. 349–358.

[30] Michael O Rabin. "Decidability of second-order theories and automata on infinite trees". In: *Transactions of the american Mathematical Society* 141 (1969), pp. 1–35.

[31] AP Sistla and EM Clarke. "The complexity of propositional linear temporal logics". In: *Proceedings of the fourteenth annual ACM symposium on Theory of computing*. ACM. 1982, pp. 159–168.

[32] Christel Baier et al. "Controller synthesis for probabilistic systems". In: *Exploring New Frontiers of Theoretical Informatics*. Springer, 2004, pp. 493–506.

[33] Antonín Kučera and Oldřich Stražovský. "On the controller synthesis for finite-state Markov decision processes". In: *International Conference on Foundations of Software Technology and Theoretical Computer Science*. Springer. 2005, pp. 541–552.

[34] Ernst Moritz Hahn, Tingting Han, and Lijun Zhang. "Synthesis for PCTL in parametric Markov decision processes". In: *Nasa formal methods symposium*. Springer. 2011, pp. 146–161.

[35] Wanwei Liu et al. "A simple probabilistic extension of modal mu-calculus". In: *Twenty-Fourth International Joint Conference on Artificial Intelligence*. 2015.

[36] Kim G Larsen, Radu Mardare, and Bingtian Xue. "Probabilistic Mu-Calculus: Decidability and Complete Axiomatization". In: *36th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science* (*FSTTCS 2016*). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2016.

[37] Matteo Mio. "Probabilistic modal μ-calculus with independent product". In: *arXiv preprint arXiv:1211.1511* (2012).

[38] Pablo Castro, Cecilia Kilmurray, and Nir Piterman. "Tractable probabilistic μ-calculus that expresses probabilistic temporal logics". In: *32nd International Symposium on Theoretical Aspects of Computer Science* (*STACS 2015*). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2015.