The Satisfiability Problem for a Quantitative Fragment of PCTL

Miroslav Chodil[a], Antonín Kučera[a]

[a]*Masaryk University, Botanická 68a, Brno, 60200, Czech Republic*

Abstract

We propose a generic method for proving the decidability of the finite satisfiability problem for PCTL fragments and demonstrate its applicability in several non-trivial examples.

Keywords: satisfiability, probabilistic temporal logics, PCTL

1. Introduction

Probabilistic CTL (PCTL) [1] is a temporal logic applicable to discrete-time probabilistic systems with Markov chain semantics. PCTL is obtained from the standard CTL (see, e.g., [2]) by replacing the existential/universal path quantifiers with the probabilistic operator $P(\Phi) \bowtie r$. Here, Φ is a path formula, $\bowtie$ is a comparison such as $\geq$ or $<$, and r is a numerical constant. A formula $P(\Phi) \bowtie r$ holds in a state s if the probability of all runs initiated in s satisfying Φ is $\bowtie$-bounded by r. The *satisfiability problem for PCTL*, asking whether a given PCTL formula has a model, is a long-standing open question in probabilistic verification resisting numerous research attempts.

Unlike CTL and other non-probabilistic temporal logics, PCTL does not have the small model property guaranteeing the existence of a bounded-size model for every satisfiable formula. In fact, one can easily construct satisfiable PCTL formulae without *any* finite model (see, e.g., [3]). Hence, the PCTL satisfiability problem is studied in two basic variants: (1) *finite satisfiability*, where we ask about the existence of a finite model, and (2) *general satisfiability*, where we ask about the existence of an unrestricted model.

Email address: `kucera@fi.muni.cz` (Antonín Kučera)

For the *qualitative fragment* of PCTL, where the range of admissible probability constraints is restricted to $\{=0, >0, =1, <1\}$, both variants of the satisfiability problem are **EXPTIME**-complete, and a finite description of a model for a (finitely) satisfiable formula is effectively constructible [4, 3]. Unfortunately, the underlying proof techniques are not applicable to general PCTL with unrestricted probability constraints such as ≥ 0.25 or <0.7.

To solve the *finite* satisfiability problem for some PCTL fragment, it suffices to establish a computable upper bound on the number of states of a model for a finite-satisfiable formula of the fragment[1]. At first glance, one is tempted to conjecture the existence of such a bound for the whole PCTL because there is no apparent way how a finite-satisfiable PCTL formula φ can "enforce" the existence of $F(\varphi)$ distinct states in a model of φ, where F grows faster than any computable function. Interestingly, this conjecture is *provably wrong* in a slightly modified setting where we ask about finite PCTL satisfiability in a *subclass* of Markov chains $\mathcal{M}^k$ where every state has at most $k \geq 2$ immediate successors (the k is an arbitrarily large fixed constant). This problem is *undecidable* and hence no computable upper bound on the size of a finite model in $\mathcal{M}^k$ exists [3] (see [5] for a full proof). So far, all attempts at extending the undecidability proof of [3] to the class of unrestricted Markov chains have failed; it is not yet clear whether the obstacles are invincible.

Regardless of the ultimate decidability status of the (finite) PCTL satisfiability, the study of PCTL fragments brings important insights into the structure and expressiveness of PCTL. The existing works [6, 7] identify several fragments where every (finitely) satisfiable formula has a model of bounded size and specific shape. In [7], it is shown that every formula φ of the *bounded fragment* of PCTL, where the validity of φ in a state s depends only on a bounded prefix of a run initiated in s, has a bounded-size tree model. In [6], several PCTL fragments based on F and G operators are studied. For each of these fragments, it is shown that every finitely satisfiable formula has a bounded-size model where every non-bottom SCC is a singleton. It is also shown that there are finitely satisfiable PCTL formulae

[1] Although there are uncountably many Markov chains with n states, the edge probabilities can be represented symbolically by variables, and the satisfiability of a given PCTL formula in a Markov chain with n states can then be encoded in the existential fragment of first-order theory of the reals. This construction is presented in Appendix A.

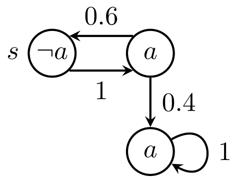

Figure 1: A Markov chain M such that $s \models \varphi$.

without a model of this shape. An example of such a formula is

$$\varphi \quad \equiv \quad G_{=1}\big(F_{\geq 0.5}(a \wedge F_{\geq 0.2}\neg a) \vee a\big) \;\wedge\; F_{=1}\, G_{=1}\, a \;\wedge\; \neg a \tag{1}$$

In [6], it is shown that φ is finitely satisfiable[2], but every finite model of φ has a non-bottom SCC with at least two states, such as the Markov chain M of Fig. 1.

Our contribution. A crucial step towards solving the finite satisfiability problem for PCTL is understanding the role of non-bottom SCCs. Intuitively, if a given PCTL formula φ enforces a model with a non-bottom SCC, then the top SCC must achieve some "progress" in satisfying φ, and successor SCCs are required to satisfy only "simpler" formulae. If this progress is effectively measurable, there is a computable upper bound on the model's size. We start by elaborating this intuition into a workable set of conditions defining *effectively progressive PCTL fragments* in Section 3. This entails a generic method for proving the decidability of the finite satisfiability problem for PCTL fragments.

The method applies to the fragments considered in [6, 7] and also to other fragments not covered by the existing results (see Section 3.1). Then, we design an abstract class of *loop progressive fragments* in Section 4, where the progress is achievable by a finite loop with one exit state. We show that every loop progressive fragment is effectively progressive, and establish the **2-EXPSPACE** upper complexity bound for the finite satisfiability problem. Furthermore, we give three examples of loop progressive fragments (two of them subsume the formula φ defined by (1) above).

In our constructions, we had to address fundamental issues specific to quantitative PCTL. The basic observation behind the small model property

[2] In [6], the formula φ has the same structure but uses qualitative probability constraints.

proofs for non-probabilistic temporal logics (and also *qualitative* PCTL) is that the satisfaction of a given formula in a given state s is determined by the satisfaction of φ and its subformulae in the successor states of s. This does not hold for quantitative PCTL. For example, knowing whether immediate successors of a state s satisfy the formula $\mathrm{F}_{\geq 0.2}\,\varphi$ does not necessarily allow to determine the satisfaction of $\mathrm{F}_{\geq 0.2}\,\varphi$ in s. We need a *precise probability* of satisfying the path formula $\mathrm{F}\,\varphi$ in the successors of s. Clearly, it makes no sense to filter a model according to the satisfaction of infinitely many formulae of the form $\mathrm{F}_{\geq r}\,\varphi$. In our proofs, we use a method for extending the set of "relevant formulae" so that it remains bounded and still captures the crucial properties of states.

Related work. The satisfiability problem for (non-probabilistic) CTL is known to be **EXPTIME**-complete [8]. The same upper bound is valid also for a richer logic of the modal μ-calculus [9, 10]. The probabilistic extension of CTL (and also CTL*) was initially studied in its qualitative form [11, 12, 4]. The satisfiability problem is shown decidable in these works. A precise complexity classification of general and finite satisfiability and a construction of (a finite description of) a model are given in [3]. In the same paper, it is also shown that the satisfiability and the finite satisfiability problems are undecidable when the class of admissible models is restricted to Markov chains with a k-bounded branching degree, where $k \geq 2$ is an arbitrary constant. A variant of the bounded satisfiability problem, where transition probabilities are restricted to $\{\frac{1}{2}, 1\}$, is proven **NP**-complete in [13]. The decidability of finite satisfiability for various quantitative PCTL fragments is established in the works [6, 7] discussed above.

The *model-checking* problem for PCTL has been studied both for finite Markov chains (see, e.g., [14, 15, 16, 17]) and for infinite Markov chains generated by probabilistic pushdown automata and their subclasses [18, 19, 20]. PCTL formulae have also been used as *objectives* in Markov decision processes (MDPs) and stochastic games, where the players controlling nondeterministic states strive to satisfy/falsify a given PCTL formula. Positive decidability results exist for finite MDPs and qualitative PCTL formulae [21]. For quantitative PCTL and finite MDPs, the problem becomes undecidable [22]. Let us note that the aforementioned undecidability results for the (finite) PCTL satisfiability problem in subclasses of Markov chains with bounded branching degree follow by utilizing proof techniques of [22].

2. Preliminaries

We use $\mathbb{N}$, $\mathbb{Q}$, $\mathbb{R}$ to denote the sets of non-negative integers, rational numbers, and real numbers, respectively. We use the standard notation for writing intervals of real numbers, e.g., $[0, 1)$ denotes the set of all $r \in \mathbb{R}$ such that $0 \leq r < 1$. For a set A, we use $|A|$ to denote the cardinality of A.

The logic PCTL [1] is a probabilistic version of Computational Tree Logic [2] obtained by replacing the existential and universal path quantifiers with the probabilistic operator $P(\Phi) \bowtie r$, where Φ is a path formula, $\bowtie$ is a comparison, and $r \in [0, 1]$ is a constant.

In full PCTL, the syntax of path formulae is based on the X, U, and $\mathrm{U}^{\leq k}$ ('next', 'until', and 'bounded until') operators. In this paper, we consider a simplified variant of PCTL based on F and G operators.

Definition 1 (PCTL). *Let AP be a set of atomic propositions. The syntax of PCTL state and path formulae is defined by the following abstract syntax equations:*

$$\begin{array}{lcl} \varphi & ::= & a \mid \neg a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid P(\Phi) \trianglerighteq r \\ \Phi & ::= & \mathrm{F}\,\varphi \mid \mathrm{G}\,\varphi \end{array}$$

Here, $a \in AP$, $\trianglerighteq \in \{\geq, >\}$, and $r \in [0, 1]$, where the trivial probability constraints '≥ 0' and '>1' are syntactically forbidden.

Since the formula Φ in the probabilistic operator $P(\Phi) \trianglerighteq r$ is always of the form $\mathrm{F}\,\varphi$ or $\mathrm{G}\,\varphi$, we often write just $\mathrm{F}_{\trianglerighteq r}\ \varphi$ and $\mathrm{G}_{\trianglerighteq r}\ \varphi$ instead of $P(\mathrm{F}\,\varphi) \trianglerighteq r$ and $P(\mathrm{G}\,\varphi) \trianglerighteq r$, respectively. The probability constraint '≥ 1' is usually written as '$=1$'. The sets of all state and path subformulae of a given state formula φ are denoted by $sub(\varphi)$ and $psub(\varphi)$, respectively. Note that every formula in $psub(\varphi)$ is of the form $\mathrm{F}\,\psi$ or $\mathrm{G}\,\psi$. We also use $AP(\varphi)$ to denote the set of all atomic propositions occurring in φ.

Observe that the negation is applicable only to atomic propositions, and the comparison ranges only over $\{\geq, >\}$. This causes no loss of generality because negations can be pushed inside, and formulae such as $\mathrm{F}_{\leq r}\ \varphi$ and $\mathrm{G}_{<r}\ \varphi$ are equivalent to $\mathrm{G}_{\geq 1-r}\ \neg\varphi$ and $\mathrm{F}_{>1-r}\ \neg\varphi$, respectively.

The encoding size of a given PCTL formula φ is denoted by $\|\varphi\|$. When PCTL formulae are given as input to algorithms, we assume that their probability bounds are rational and written as irreducible fractions of binary integers. The original definition of PCTL [1] (and also Definition 1) admits arbitrary probability bounds in the interval $[0, 1]$. This is convenient because we sometimes consider PCTL formulae where the probability bounds correspond to probabilities of certain events in unspecified Markov chains.

PCTL formulae are interpreted over Markov chains where every state s is assigned a subset $v(s) \subseteq AP$ of atomic propositions valid in s.

Definition 2 (Markov chain). *A Markov chain is a triple $M = (S, P, v)$, where S is a finite or countably infinite set of* states, *$P \colon S \times S \to [0, 1]$ is a function such that $\sum_{t \in S} P(s, t) = 1$ for every $s \in S$, and $v \colon S \to 2^{AP}$ is a* valuation.

A *path* in M is a finite sequence $w = s_0 \dots s_n$ of states where $n \geq 0$ and $P(s_i, s_{i+1}) > 0$ for all $i < n$. A *run* in M is an infinite sequence $\pi = s_0 s_1 \dots$ of states such that every finite prefix of π is a path in M. We also use $\pi(i)$ to denote the state s_i of π.

A *strongly connected component (SCC)* of M is a maximal $U \subseteq S$ such that, for all $s, t \in U$, there is a path from s to t. Given two SCCs U and V, we say that V is a *successor/predecessor* of U if $U \neq V$ and there exists a path from/to a state of U to/from a state of V. A *bottom SCC (BSCC)* is a SCC without successors, and a *top SCC* is a SCC without predecessors. The successor relation is a strict partial order and determines the standard directed acyclic graph (DAG) of SCCs.

For every path $w = s_0 \dots s_n$, let $Run(w)$ be the set of all runs starting with w, and let $\mathbb{P}(Run(w)) = \prod_{i=0}^{n-1} P(s_i, s_{i+1})$. To every state s, we associate the probability space $(Run(s), \mathcal{F}_s, \mathbb{P}_s)$, where $\mathcal{F}_s$ is the σ-field generated by all $Run(w)$ where w starts in s, and $\mathbb{P}_s$ is the unique probability measure obtained by extending $\mathbb{P}$ in the standard way (see, e.g., [23]).

The *validity* of a PCTL state/path formula for a given state/run of M is defined inductively as follows:

$$
\begin{array}{lll}
s \models a & \text{iff} & a \in v(s), \\
s \models \neg a & \text{iff} & a \notin v(s), \\
s \models \varphi_1 \wedge \varphi_2 & \text{iff} & s \models \varphi_1 \text{ and } s \models \varphi_2, \\
s \models \varphi_1 \vee \varphi_2 & \text{iff} & s \models \varphi_1 \text{ or } s \models \varphi_2, \\
s \models P(\Phi) \trianglerighteq r & \text{iff} & \mathbb{P}_s(\{\pi \in Run(s) \mid \pi \models \Phi\}) \trianglerighteq r, \\
\pi \models \mathrm{F}\,\varphi & \text{iff} & \pi(i) \models \varphi \text{ for some } i \in \mathbb{N}, \\
\pi \models \mathrm{G}\,\varphi & \text{iff} & \pi(i) \models \varphi \text{ for all } i \in \mathbb{N}.
\end{array}
$$

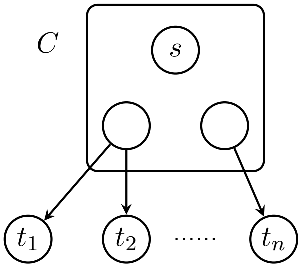

Figure 2: The structure of a finite model of φ.

We say that M is a *model* of φ if $s \models \varphi$ for some state s of M. A PCTL formula φ is *valid* if $\neg\varphi$ does not have a model. We say that φ *implies* ψ if the implication $\varphi \to \psi$ is a valid formula.

The *(finite) PCTL satisfiability problem* is the question of whether a given PCTL formula has a (finite) model.

Remark 1. *In this paper, we often apply notions and constructions defined for a single formula to finite sets of formulae. Such a set X should then be formally understood as a* conjunction *of its elements. For example, $s \models X$ means that $s \models \bigwedge_{\varphi \in X} \varphi$.*

3. Effectively Progressive PCTL Fragments

In this section, we introduce a general technique for establishing the decidability of the finite satisfiability problem for PCTL fragments.

As we already noted in Section 1, one sufficient condition implying the decidability of the finite satisfiability problem for a given PCTL fragment $\mathcal{L}$ is the existence of a computable function $B : \mathcal{L} \to \mathbb{N}$ such that every finitely satisfiable $\varphi \in \mathcal{L}$ has a model with at most $B(\varphi)$ states. This follows directly from the next (folklore) proposition:

Proposition 1. *Let φ be a PCTL formula and $n \in \mathbb{N}$. The problem of whether φ has a model with at most n states is decidable is space polynomial in $\|\varphi\|$ and n.*

For the sake of completeness, a proof of Proposition 1 is given in Appendix A.

Consider a PCTL formula φ with a finite model M. Let s be a state of M such that $s \models \varphi$, and let C be the SCC containing s. Clearly, all SCCs of M that are not reachable from s can be removed from M without influencing the validity of φ is s. Hence, we can assume that M has only *one* top SCC C, and C contains a state s satisfying φ. Furthermore, we can assume that M has been chosen so that the height of the corresponding DAG of SCCs (called the *DAG-height of M* in the sequel) is *minimal*, and we use $\mathcal{H}(\varphi)$ to denote this height.

In our next theorem, we show that all finitely satisfiable PCTL formulae φ such that $\mathcal{H}(\varphi) = 0$ have a model of bounded size. This entails the decidability of the problem of whether $\mathcal{H}(\varphi) = 0$ for a given PCTL formula.

Theorem 2. *For every PCTL formula φ we have that $\mathcal{H}(\varphi) = 0$ iff φ has a strongly connected model with at most $2^{|AP(\varphi)|}$ states.*

Proof. Let φ be a PCTL formula such that $\mathcal{H}(\varphi) = 0$, and let $s \models \varphi$ where s is a state of a strongly connected Markov chain M. For every state t of M, let $AP(\varphi, t)$ be the set of all $a \in AP(\varphi)$ valid in t. Consider a Markov chain M' where $\{AP(\varphi, t) \mid t \text{ is a state of } M\}$ is the set of states and the transitions are defined arbitrarily so that M' is strongly connected. Furthermore, every state α of M' satisfies exactly those propositions of $AP(\varphi)$ that occur in α. By a straightforward induction on the structure of ξ, we obtain that $t \models \xi$ implies $AP(\varphi, t) \models \xi$ for every state t of M and every $\xi \in \mathit{sub}(\varphi)$ (here, we use the standard result of finite Markov chain theory saying that a run initiated in an arbitrary state of a strongly connected finite Markov chain visits all states with probability one). In particular, $AP(\varphi, s) \models \varphi$. □

Now assume $\mathcal{H}(\varphi) \geq 1$. Let M be a model of φ with one top SCC C and the DAG-height equal to $\mathcal{H}(\varphi)$ such that $s \models \varphi$ for some $s \in C$. Intuitively, the top SCC C must then achieve some "progress" in satisfying φ before a run leaves C, because otherwise C could be removed from the model and the DAG height of M would not be minimal. In other words, the initial commitment of satisfying φ in s is "transformed" into finitely many simpler commitments imposed on *C-descendants*, i.e., states $t_1, \ldots, t_n$ entered right after leaving C (see Fig. 2). We use $\mathit{Desc}(C)$ to denote the set of all C-descendants.

Now we formalize the above intuition.

Definition 3 (φ-commitment). *Let φ be a PCTL formula. A φ-*commitment *is a vector $X \in [0,1]^{AP(\varphi) \cup \mathit{psub}(\varphi)}$ where $X(a) \in \{0,1\}$ for all $a \in AP(\varphi)$.*

Intuitively, each φ-commitment X represents a finite set of PCTL formulae consisting of all

- a such that $a \in AP(\varphi)$ and $X(a) = 1$,
- $\neg a$ such that $a \in AP(\varphi)$ and $X(a) = 0$,
- $P(\Phi) \geq r$ such that $\Phi \in \mathit{psub}(\varphi)$ and $X(\Phi) = r > 0$.

Note that in the last item, we disregard the case when $r = 0$ because the constraint "≥ 0" is trivial and syntactically forbidden. Slightly abusing our notation, we use X to denote the associated set of formulae (writing, e.g., $s \models X$ or $\psi \in X$).

Observe that a φ-commitment X determines the (in)validity of all $a \in AP(\varphi)$ and implies the validity of some state subformulae of φ. In particular, some φ-commitments *imply* φ. This is illustrated in the following example:

Example 1. *Let* $\varphi \equiv \mathrm{F}_{\geq 0.4}\, a \vee \mathrm{G}_{\geq 0.3}\, b$*, and let* X *be a* φ*-commitment such that* $X(a) = 0$*,* $X(b) = 1$*,* $X(\mathrm{F}\, a) = 0.5$*, and* $X(\mathrm{G}\, a) = 0.2$*. Then* X *implies the state subformula* $\mathrm{F}_{\geq 0.4}\, a$ *(and hence also* φ*), but not the state subformula* $\mathrm{G}_{\geq 0.3}\, b$ *(clearly, if* $t \models X$*, then also* $t \models \mathrm{F}_{\geq 0.4}\, a$ *and* $t \models \varphi$*, but it may happen that* $t \not\models \mathrm{G}_{\geq 0.3}\, b$*).*

Also observe that even if φ is finitely satisfiable, some φ-commitments may not be (finitely) satisfiable.

As we already indicated, the progress in satisfying φ corresponds to decreasing the "complexity" of φ-commitments imposed on the descendants of SCCs in a model with the minimal DAG-height. In general, such progress may not be achievable for all finitely satisfiable φ-commitments in the same way. However, it suffices to achieve progress for an *eligible* subset of φ-commitments.

Definition 4 (eligible set of φ-commitments). *Let* φ *be a finitely satisfiable PCTL formula. A set* Com_φ *of* φ*-commitments is* eligible *if every element of* Com_φ *is finitely satisfiable and there exists* $X \in \mathit{Com}_\varphi$ *such that* X *implies* φ*.*

Now we define a complexity measure for Com_φ.

Definition 5 (complexity measure). *Let* φ *be a finitely satisfiable PCTL formula and* Com_φ *an eligible set of* φ*-commitments. A* complexity measure *for* Com_φ *is a function* $g : \mathit{Com}_\varphi \to \mathbb{N}$ *such that* $g(X) = 0$ *only if* $\mathcal{H}(X) = 0$*.*

The requirement $g(X) = 0$ only if $\mathcal{H}(X) = 0$ avoids the problematic case when $\mathcal{H}(X) \geq 1$ and $g(X) = 0$ (note that if $\mathcal{H}(X) \geq 1$, then each SCC C containing a state satisfying X must have a non-empty set of descendants; if we also had $g(X) = 0$, the g-value of the commitments assigned to these descendants could not be further decreased).

So far, we have not discovered a universal complexity measure applicable to all Com_φ (this would yield the decidability of the finite satisfiability problem for the whole PCTL). Since we do not require the computability of g, the function $\mathcal{H}$ appears to be a natural candidate. However, some effectiveness assumption is unavoidable to obtain a computable upper bound on the size of a model. In our setup (see Definition 7), we need an effective upper bound on the size of

$$\min\{g(X) \mid X \in \mathit{Com}_\varphi, X \text{ implies } \varphi\}$$

which prevents using $\mathcal{H}$. Nevertheless, it is possible to tailor specific complexity measures for various PCTL fragments, as we shall see in the next sections.

The concept of assigning commitments to C-descendants is formalized as follows:

Definition 6 (C-assignment). *Let C be a SCC in some Markov chain, $X \in \mathit{Com}_\varphi$, and g a complexity measure for Com_φ. A* C-assignment *is a function $\mathcal{A} : \mathit{Desc}(C) \to \mathit{Com}_\varphi$. We say that $\mathcal{A}$ is*

- safe *for X if there exists $t \in C$ such that $t \models X$ in every Markov chain obtained from C by replacing each $u \in \mathit{Desc}(C)$ with a state satisfying $\mathcal{A}(u)$;*
- g-progressive *for X if for every $u \in \mathit{Desc}(C)$ we have that either $\mathcal{H}(\mathcal{A}(u)) = 0$ or $g(\mathcal{A}(u)) < g(X)$.*

Now we define an effectively progressive PCTL fragment and prove that the finite satisfiability problem is decidable for each such fragment.

Definition 7 (effectively progressive PCTL fragment). *A PCTL fragment $\mathcal{L}$ is* effectively progressive *if there are computable functions $c, h : \mathcal{L} \to \mathbb{N}$ such that, for every finitely satisfiable $\varphi \in \mathcal{L}$, there exist an eligible set Com_φ and a complexity measure g for Com_φ satisfying the following conditions:*

- *There is $X \in Com_\varphi$ such that X implies φ and $g(X) \leq h(\varphi)$.*
- *For every $X \in Com_\varphi$ such that $\mathcal{H}(X) > 0$ there exists a SCC C of some Markov chain satisfying the following conditions:*
  - *$|C| \leq c(\varphi)$.*
  - *There is a C-assignment which is safe and g-progressive for X.*

Using the conditions of Definition 7, we immediately obtain the existence of a finite model M for φ whose DAG-height is bounded by $h(\varphi)$ and every non-bottom SCC of M has at most $c(\varphi)$ states. Now we prove that the *branching degree* of M can also be effectively bounded, yielding a computable upper bound on the number of states of M.

We start by introducing a special φ-commitment that is also used in the following sections.

Definition 8 ($\mathcal{X}_t$ commitment). *Let φ be a PCTL formula. For every state t of a Markov chain M, let $\mathcal{X}_t$ be the φ-commitment defined as follows:*

- *For every $a \in AP(\varphi)$, we have that $\mathcal{X}_t(a)$ is either 1 or 0, depending on whether $t \models a$ or not, respectively.*
- *For every $\Phi \in psub(\varphi)$, we put $\mathcal{X}_t(\Phi) = r$, where $r = \mathbb{P}_t(\{\pi \in Run(t) \mid \pi \models \Phi\})$.*

Let φ be a PCTL formula. Furthermore, let $M = (S, P, v)$ be a finite Markov chain, C a SCC of M, and $s \in C$. Let $Desc(C, s) = \{u_1, \ldots, u_n\}$ be the set of all C-descendants of s, i.e., $Desc(C, s) = \{u \in Desc(C) \mid P(s, u) > 0\}$. Consider a Markov chain obtained by changing the set $Desc(C, s)$ into $\{v_1, \ldots, v_m\}$ where $v_1, \ldots, v_m \notin C$. That is, the function P is changed into P' so that

- $P'(x, y) = P(x, y)$ whenever $x \neq s$ or $y \notin \{u_1, \ldots, u_n, v_1, \ldots, v_m\}$;
- $\sum_{i=1}^{n} P(s, u_i) = \sum_{i=1}^{m} P'(s, v_i)$.

Now observe that for all $\Phi \in psub(\varphi)$ and $t \in C$ we have that if

$$\sum_{i=1}^{n} P(s, u_i) \cdot \mathcal{X}_{u_i}(\Phi) \quad \leq \quad \sum_{i=1}^{m} P'(s, v_i) \cdot \mathcal{X}_{v_i}(\Phi) \tag{2}$$

then also

$$\mathbb{P}_t(\pi \in \mathit{Run}(t) \mid \pi \models \Phi) \quad \leq \quad \mathbb{P}'_t(\pi \in \mathit{Run}(t) \mid \pi \models \Phi) \tag{3}$$

where $\mathbb{P}$ and $\mathbb{P}'$ are the probability measures induced by P and P', respectively. This follows immediately from the semantics of PCTL path formulae.

Now suppose that $t \models X$ for some $t \in C$ and a φ-commitment X. Consider the vector

$$Y \ = \ \sum_{i=1}^{n} \frac{P(s,u_i)}{P_s} \cdot \mathcal{X}_{u_i}$$

where $P_s = \sum_{i=1}^{n} P(s,u_i)$. By Carathéodory's convex hull theorem, there exist $\{v_1,\ldots,v_m\} \subseteq \{u_1,\ldots,u_n\}$ and positive coefficients $p_1,\ldots,p_m$ such that $m \leq |\mathit{psub}(\varphi)| + 1$, $\sum_{i=1}^{m} p_i = 1$, and

$$Y(\Phi) \leq \sum_{i=1}^{m} p_i \cdot \mathcal{X}_{v_i}(\Phi)$$

for every $\Phi \in \mathit{psub}(\varphi)$. Let $P'(s,v_i) = p_i \cdot P_s$. Then (2) is satisfied for all $\Phi \in \mathit{psub}(\varphi)$, and hence (3) holds for t and all $\Phi \in \mathit{psub}(\varphi)$. This means that $t \models X$ in the modified Markov chain where the C-descendants $\{u_1,\ldots,u_n\}$ of s are replaced with $\{v_1,\ldots,v_m\}$. The same procedure can be repeated for all states of C. Thus, we obtain the following theorem:

Theorem 3. *If a PCTL formula φ has a finite model M, then φ also has a model M' such that for every SCC C of M', the following conditions are satisfied:*

- *C is a SCC of M;*
- *$\mathit{Desc}(C,s)$ in M' is a subset of $\mathit{Desc}(C,s)$ in M for every $s \in C$;*
- *the size of $\mathit{Desc}(C,s)$ in M' is at most $|\mathit{psub}(\varphi)| + 1$.*

For effectively progressive PCTL fragments, Theorem 3 allows for computing an upper bound on the size of a finite model.

Theorem 4. *Let $\mathcal{L}$ be an effectively progressive PCTL fragment. Then every finitely satisfiable $\varphi \in \mathcal{L}$ has a model with at most*

$$2 \cdot (c(\varphi) \cdot |\mathit{sub}(\varphi)|)^{h(\varphi)} \cdot \max\{c(\varphi), 2^{|AP(\varphi)|}\}$$

states.

PROOF. According to Definition 7 and Theorem 3, φ has a model where every non-bottom SCC has at most $c(\varphi)$ states and hence at most $c(\varphi) \cdot (|psub(\varphi){+}1|)$ descendants. Clearly, $|psub(\varphi){+}1| \leq |sub(\varphi)|$. The DAG-height of the model is bounded by $h(\varphi)$, and hence the total number of all SCCs of the model is bounded by $2 \cdot (c(\varphi) \cdot |sub(\varphi)|)^{h(\varphi)}$. The number of states in a non-bottom SCC is bounded by $c(\varphi)$, and the number of states in a BSCC is bounded by $2^{|AP(\varphi)|}$ by Theorem 2. Thus, we obtain the presented bound on the number of states. □

3.1. Examples of Effectively Progressive PCTL Fragments

First, let us note that the quantitative PCTL fragments with decidable finite satisfiability problem studied in [6, 7] are effectively progressive.

The fragments studied in [6] are arranged into a syntactic hierarchy with two maximal elements $\mathrm{G}_q(\mathrm{F}_q, \mathrm{G}_q, \vee)$ and $\mathrm{F}_q, \mathrm{G}_1, \vee$. The fragment $\mathrm{G}_q(\mathrm{F}_q, \mathrm{G}_q, \vee)$ consists of formulae φ defined by the following abstract syntax:

$$\begin{array}{lll} \varphi & ::= & \mathrm{G}_{\triangleright r}\, \psi \\ \psi & ::= & a \;\mid\; \neg a \;\mid\; \psi_1 \wedge \psi_2 \;\mid\; \psi_1 \vee \psi_2 \;\mid\; \mathrm{F}_{\triangleright r}\, \psi \;\mid\; \mathrm{G}_{\triangleright r}\, \psi \end{array}$$

In [6], it is shown that $\mathcal{H}(\varphi) = 0$ for every finitely satisfiable formula φ of this fragment. Hence, the conditions of Definition 7 are satisfied as follows. For every finitely satisfiable φ, we put

- $c(\varphi) = h(\varphi) = 0$;
- $Com_\varphi = \{\mathcal{X}_s\}$ where s is a state of a strongly connected model of φ such that $s \models \varphi$; we also put $g(X) = 0$.

Note that the second condition of Definition 7 holds trivially because $g(X) = 0$ for every $X \in Com_\varphi$.

The fragment $\mathrm{F}_q, \mathrm{G}_1, \vee$ is defined as follows:

$$\varphi \;::=\; a \;\mid\; \neg a \;\mid\; \varphi_1 \wedge \varphi_2 \;\mid\; \mathrm{F}_{\triangleright r}\, \psi \;\mid\; \mathrm{G}_{=1}\, \psi$$

As it is proven in [6], every finitely satisfiable formula φ of $\mathrm{F}_q, \mathrm{G}_1, \vee$ has a model where every non-bottom SCC is a singleton, and the DAG of SCCs is a tree of height at most $H(\varphi)$ where $H(\varphi) \in \mathbb{N}$ is a computable constant. Again, the model's shape allows us to construct the functions required in Definition 7 almost trivially. For every finitely satisfiable φ, we fix a tree-like model M_φ of height at most $H(\varphi)$. For every state s of M_φ, let H_s be the DAG-height of the Markov chain obtained from M_φ by removing all SCCs that are not reachable from the SCC containing s. We put

- $c(\varphi) = 0$, $h(\varphi) = H(\varphi)$;
- $Com_\varphi = \{\mathcal{X}_s \mid s \text{ is a state of } M\}$;
- $g(\mathcal{X}_s) = \min\{H_t \mid\ t \text{ is a state of } M_\varphi \text{ such that } t \models \mathcal{X}_s\}$.

Then, the two conditions of Definition 7 are satisfied.

The bounded PCTL fragment studied in [7] contains formulae whose validity in a state s depends only on an effectively bounded prefix of runs initiated in s. Hence, every finitely satisfiable formula of this fragment has a tree-like model of an effectively bounded height. In principle, the effective progressivity of the fragment can be justified in the same way as for the $\mathrm{F}_q, \mathrm{G}_1, \vee$ fragment above. However, our definition of effectively progressive fragments (Definition 7) is tailored for a simplified variant of PCTL without the bounded until operator used in the bounded PCTL fragment. Hence, we first need to adjust Definition 7 to full PCTL syntax, which is straightforward (observe that even if the set $psub(\varphi)$ changes, the notion of φ-commitments makes good sense, and Definition 7 does not require any major modifications).

Now we give an example of an effectively progressive PCTL fragment that is *not* covered by the results of [6, 7].

Let $\mathcal{L}_1$ be a PCTL fragment defined by the following abstract syntax equations:

$$\begin{array}{lcl} \varphi & ::= & a \mid \neg a \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \mathrm{F}_{\triangleright r}\, \varphi \mid \mathrm{G}_{\triangleright r}\, \theta \\ \theta & ::= & a \mid \neg a \mid \theta_1 \wedge \theta_2 \mid \theta_1 \vee \theta_2 \mid \mathrm{G}_{\triangleright r}\, \theta \end{array}$$

Note that the syntax of $\mathcal{L}_1$ forbids using the $\mathrm{F}_{\triangleright r}$ operator inside formulae of the form $\mathrm{G}_{\triangleright r}\, \theta$. We show that $\mathcal{L}_1$ is effectively progressive. To achieve that, we need some additional definitions.

For every $\varphi \in \mathcal{L}_1$, let Com_φ be the set of all finitely satisfiable φ-commitments. Furthermore, let $c, h\colon \mathcal{L}_1 \to \mathbb{N}$ and $g\colon Com_\varphi \to \mathbb{N}$ be functions defined as follows:

$$\begin{aligned} c(\varphi) &= 1, \\ h(\varphi) &= \sum_{\mathrm{F}\,\psi \in psub(\varphi)} \langle \mathrm{F}\,\psi \rangle, \\ g(X) &= \sum_{\substack{\mathrm{F}\,\psi \in psub(\varphi) \\ X(\mathrm{F}\,\psi) > 0}} \langle \mathrm{F}\,\psi \rangle. \end{aligned}$$

Here, $\langle \Phi \rangle$ is defined for every path formula Φ of the form $\mathrm{F}\,\varphi$ or $\mathrm{G}\,\varphi$ inductively by $\langle \Phi \rangle = 1 + \sum_{\Psi \in psub(\varphi)} \langle \Psi \rangle$ (the empty sum denotes 0). Observe that $g(X) = 0$ only if X contains only formulae of the form a, $\neg a$, and $\mathrm{G}_{\rhd r}\,\psi$. One can easily show that each such $X \in \mathit{Com}_\varphi$ has a one-state model, i.e., $\mathcal{H}(X) = 0$.

For the rest of this section, we fix a finitely satisfiable $\varphi \in \mathcal{L}_1$. For every $X \in \mathit{Com}_\varphi$ and every state s in some Markov chain such that $s \models X$, let $\mathit{Cl}(X,s)$ (the *closure* of X in s) be the least set K of PCTL formulae satisfying the following conditions:

- $X \subseteq K$,
- if $\psi_1 \vee \psi_2 \in K$ and $s \models \psi_1$, then $\psi_1 \in K$,
- if $\psi_1 \vee \psi_2 \in K$ and $s \models \psi_2$, then $\psi_2 \in K$,
- if $\psi_1 \wedge \psi_2 \in K$, then $\psi_1, \psi_2 \in K$,
- if $P(\mathrm{G}\,\psi) \rhd r \in K$, then $\psi \in K$,
- if $P(\mathrm{F}\,\psi) \rhd r \in K$ and $s \models \psi$, then $\psi \in K$.

Observe that $\mathit{Cl}(X,s) \subseteq \mathcal{L}_1$ and $s \models \mathit{Cl}(X,s)$. Furthermore, we define a φ-commitment $X[s] \in \mathit{Com}_\varphi$ as follows:

- $X[s](a) = X(a)$ for every $a \in AP(\varphi)$;
- $X[s](\mathrm{G}\,\psi) = \mathbb{P}_s(\{\pi \in \mathit{Run}(s) \mid \pi \models \mathrm{G}\,\psi\})$ for all $\mathrm{G}\,\psi \in \mathit{Cl}(X,s)$;
- $X[s](\mathrm{F}\,\psi) = \mathbb{P}_s(\{\pi \in \mathit{Run}(s) \mid \pi \models \mathrm{F}\,\psi\})$ for all $\mathrm{F}\,\psi \in \mathit{Cl}(X,s)$ such that $s \not\models \psi$;
- $X[s](\Psi) = 0$ for all other $\Psi \in psub(\varphi)$.

Intuitively, the φ-commitment $X[s]$ may only strengthen the requirements specified by X, i.e., $X[s]$ implies X. At first glance, this does not seem to be the case because $X[s]$ seemingly "disregards" subformulae of the form $\mathrm{F}\,\psi$ where $s \models \psi$ by setting $X[s](\mathrm{F}\,\psi) = 0$. However, $X[s]$ also sets (possibly new) requirements implying the satisfaction of ψ. Thus, $X[s]$ enforces the satisfaction of $\mathrm{F}_{=1}\,\psi$.

Example 2. *Let* $\varphi \equiv \mathrm{F}_{\geq 0.2}(\mathrm{G}_{\geq 0.5}\, a)$*, and let* X *be a* φ*-commitment such that* $X(a) = 1$*,* $X(\mathrm{F}(\mathrm{G}_{\geq 0.5}\, a)) = 0.7$*, and* $X(\mathrm{G}\, a) = 0$*. Consider a state* s *such that* $s \models \mathrm{G}_{=0.8}\, a$*. Then* $X[s](a) = 1$*,* $X[s](\mathrm{F}(\mathrm{G}_{\geq 0.5}\, a)) = 0$*, and* $X[s](\mathrm{G}\, a) = 0.8$*.*

Note that $X[s]$ *implies* X*, and the requirement of satisfying* $\mathrm{F}(\mathrm{G}_{\geq 0.5}\, a)$ *with probability at least* 0.7 *is "replaced" with a new requirement of satisfying* $\mathrm{G}\, a$ *with probability at least* 0.8*. Also observe that* $g(X[s]) < g(X)$*.*

In general, we have that $g(X[s]) \leq g(X)$, but the inequality is not necessarily strict. As we shall see, the decrease of g is achieved by considering appropriate successors of s and "updating" $X[s]$ in these successors. Now we formalize and prove the above observations.

Lemma 5. $X[s]$ *implies* X*.*

Proof. Since $X \subseteq \mathit{Cl}(X, s)$, it suffices to show that $X[s]$ implies ψ for every $\psi \in \mathit{Cl}(X, s)$. Recall that $s \models \mathit{Cl}(X, s)$. We proceed by induction on the structure of ψ.

- $\psi \equiv a$ or $\psi \equiv \neg a$. It suffices to realize that $a \in \mathit{Cl}(X, s)$ iff $X(a) = 1$ iff $X[s](a) = 1$, and $\neg a \in \mathit{Cl}(X, s)$ iff $X(a) = 0$ iff $X[s](a) = 0$.
- $\psi \equiv \psi_1 \wedge \psi_2$ or $\psi \equiv \psi_1 \vee \psi_2$. Immediately by induction hypothesis.
- $\psi \equiv P(\mathrm{G}\, \psi) \rhd r$. Since $s \models P(\mathrm{G}\, \psi) \rhd r$, we have that $X[s](\mathrm{G}\, \psi) \rhd r$ by definition of $X[s]$. Hence, $X[s]$ implies $P(\mathrm{G}\, \psi) \rhd r$.
- $\psi \equiv P(\mathrm{F}\, \psi) \rhd r$. Then $s \models P(\mathrm{F}\, \psi) \rhd r$, and we distinguish two possibilities.
  - $s \models \Psi$. Then $\Psi \in \mathit{Cl}(X, s)$, and hence $X[s]$ implies Ψ by induction hypothesis. This means that $X[s]$ implies $P(\mathrm{F}\, \psi) = 1$, and hence also $P(\mathrm{F}\, \psi) \rhd r$.
  - $s \not\models \Psi$. Then $X[s](\mathrm{F}\, \psi) \rhd r$, and hence $X[s]$ implies $P(\mathrm{F}\, \psi) \rhd r$.

Lemma 6. $g(X[s]) \leq g(X)$*. Furthermore, if there is a formula* $\mathrm{F}\, \psi$ *such that* $X(\mathrm{F}\, \psi) > 0$ *and* $s \models \psi$*, then* $g(X[s]) < g(X)$*.*

Proof. Recall

$$g(X) = \sum_{\substack{\mathrm{F}\,\psi \in psub(\varphi) \\ X(\mathrm{F}\,\psi) > 0}} \langle \mathrm{F}\,\psi \rangle$$

Let $\mathrm{F}\,\psi \in psub(\varphi)$ such that $X(\mathrm{F}\,\psi) > 0$. If $X \not\models \psi$, then the summand for $\mathrm{F}\,\psi$ in $g(X[s])$ does not change. Otherwise, the summand for $\mathrm{F}\,\psi$ in $g(X[s])$ is replaced with (zero or more) summands for subformulae $\mathrm{F}\,\varrho_1, \ldots, \mathrm{F}\,\varrho_k$ where $k \geq 0$ and $\mathrm{F}\,\varrho_i \in psub(\mathrm{F}\,\varphi)$ for every $i \leq k$ (see the definition of $X[s]$ above). Since $\langle \mathrm{F}\,\psi \rangle > \sum_{i=1}^{k} \langle \mathrm{F}\,\varrho_i \rangle$, we are done. □

Note that in Lemma 6, we rely on the syntactic restrictions imposed by the definition of $\mathcal{L}_1$. Since $\mathcal{L}_1$ prohibits the use of F within subformulae of the form $\mathrm{G}\,\psi$, no "new" subformulae of the form $\mathrm{F}\,\varrho$ in $X[s]$ can be generated by the closure of $\mathrm{G}\,\psi$ in s. Without this restriction, Lemma 6 may not hold, as illustrated in the following example:

Example 3. *Let $\varphi \equiv \mathrm{G}_{\geq 0.4}\,\mathrm{F}_{\geq 0.8}\,a$. Note that $\varphi \notin \mathcal{L}_1$. Consider a φ-commitment X such that $X(\mathrm{G}\,\mathrm{F}_{\geq 0.8}\,a) = 0.5$ and $X(\mathrm{F}\,a) = 0$. Then for every state s such that $s \models X$ we have that $X[s](\mathrm{F}\,a) \geq 0.8$. Observe that $g(X) = 0$ and $g(X[s]) = 1$, and hence Lemma 6 does not hold for φ. Intuitively, the problem is that φ keeps "regenerating" the requirement $\mathrm{F}_{\geq 0.8}\,a$, and hence the progress in satisfying φ cannot be measured by the progress in satisfying its* F*-subformulae.*

Now we show that the two conditions of Definition 7 are satisfied for the functions c, h and g defined above. To verify the first condition, consider a state s of a finite Markov chain such that $s \models \varphi$. Then $\mathcal{X}_s \in Com_\varphi$, $\mathcal{X}_s$ implies X, and $g(\mathcal{X}_s) \leq h(\varphi)$ as required. Now let $X \in Com_\varphi$ such that $\mathcal{H}(X) > 0$, and let $s \models X$ where s is a state in a finite Markov chain $M = (S, P, v)$. Consider the set B of all $t \in S$ such that t belongs to some BSCC of M or $t \models \psi$ for some $\mathrm{F}_{\triangleright r}\,\psi \in X[s]$. Furthermore, for every $t \in B$, let

- $\nu(t)$ be the probability of all runs initiated in s visiting the state t so that all states preceding the first visit to t are not contained in B;
- Y_t be the φ-commitment obtained by "updating" $X[s]$ in t, i.e., $Y_t(a) = \mathcal{X}_t(a)$ for all $a \in AP(\varphi)$, and for every path formula Φ we have that $Y_t(\Phi)$ is equal either to 0 or $\mathcal{X}_t(\Phi)$, depending on whether $X[s](\Phi) = 0$ or $X[s](\Phi) > 0$, respectively.

Observe that $t \models Y_t$ and $g(Y_t) \leq g(X)$. If t belongs to a BSCC of M, then $\mathcal{H}(Y_t) = \mathcal{H}(Y_t[t]) = 0$. Otherwise, $g(Y_t[t]) < g(Y_t)$ by Lemma 6, hence $g(Y_t[t]) < g(X)$.

Let T be the set of all $t \in B$ such that $\nu(t) > 0$. Since the probability of all runs initiated is s that eventually visit a BSCC of M is equal to one, we obtain $\sum_{t \in T} \nu(t) = 1$. Now consider Markov chain $M' = (S', P', v')$ with a SCC $C = \{s'\}$ such that $v'(s') = v(s)$, $Desc(C) = \{t' \mid t \in T\}$, and $P'(s', t') = \nu(t)$ for every $t \in T$. Then $|C| = 1 = c(\varphi)$, and a C-assignment $\mathcal{A}$ such that $\mathcal{A}(t') = Y_t[t]$ is g-progressive for X. It remains to verify that $\mathcal{A}$ is safe for X. To see this, recall that $X[s]$ implies X and realize the following:

- For every formula $\mathrm{F}\,\psi$ such that $X[s](\mathrm{F}\,\psi) > 0$ we have that $X[s](\mathrm{F}\,\psi) = \sum_{t \in T} \nu(t) \cdot Y_t[t](\mathrm{F}\,\psi)$.
- For every formula $\mathrm{G}\,\psi$ such that $X[s](\mathrm{G}\,\psi) > 0$ we have that $X[s](\mathrm{G}\,\psi) \leq \sum_{t \in T} \nu(t) \cdot Y_t[t](\mathrm{G}\,\psi)$. Note that the inequality can be strict because the formula ψ can be invalid in some states visited along a path from s to a state of T.

Hence, if every t' is replaced with a state satisfying $Y_t[t]$, then $s' \models X[s]$ and hence also $s' \models X$ as required.

4. Loop Progressive Fragments

In this section, we introduce a special class of effectively progressive PCTL fragments called *loop-progressive fragments*, where the "progress SCC" C of Definition 7 is a simple loop with one exit state. To some extent, the presented notions and observations generalize the ones presented in Section 3.1.

As a running example, we use the formula

$$\varphi \quad \equiv \quad G_{=1}\big(F_{\geq 0.5}(a \wedge F_{\geq 0.2}\neg a) \vee a\big) \;\wedge\; F_{=1}\, G_{=1}\, a \;\wedge\; \neg a$$

and its model of Fig. 1.

Definition 9 (the sets $C(\xi, s)$ and $E(\xi, s)$). *For a PCTL formula ξ and a state s in a Markov chain M such that $s \models \xi$, we define the set $C(\psi, s)$ as the least set K satisfying the following conditions:*

- *$\xi \in K$;*

- *if* $\psi_1 \vee \psi_2 \in K$ *and* $s \models \psi_1$*, then* $\psi_1 \in K$*;*
- *if* $\psi_1 \vee \psi_2 \in K$ *and* $s \models \psi_2$*, then* $\psi_2 \in K$*;*
- *if* $\psi_1 \wedge \psi_2 \in K$*, then* $\psi_1, \psi_2 \in K$*;*
- *if* $\mathrm{F}_{\triangleright r}\, \psi \in K$ *and* $s \models \psi$*, then* $\psi \in K$*.*

The set $E(\xi, s)$ *is defined in the same way as* $C(\xi, s)$*, except the last condition is omitted.*

Recall that every φ-commitment X can be interpreted as a PCTL formula (see Remark 1), which means that $C(X, s)$ and $E(X, s)$ are defined.

Observe that $C(\xi, s)$ contains some but not necessarily *all* subformulae of ξ that are valid in s. In particular, there is no rule saying that if $\mathrm{G}_{\triangleright r}\, \psi \in K$, then $\psi \in K$. As we shall see, the subformulae within the scope of the $\mathrm{G}_{\triangleright r}$ operator are treated in a special way. Also note that $C(\xi, s)$ and $E(\xi, s)$ are defined differently from the set $Cl(\xi, s)$ used in Section 3.1.

In our next definition, we introduce the set Com_φ consisting of "relevant" φ-commitments used to progressively simplify the original φ in the sense of Definition 7.

Definition 10 (φ-commitment $\mathcal{Y}_{s,\xi}$). *Let* $\mathcal{L}$ *be a PCTL fragment and* $\varphi \in \mathcal{L}$*. For every* $\xi \in \mathcal{L}$ *and every state* s *in some finite-state Markov chain such that* $psub(\xi) \subseteq psub(\varphi)$ *and* $s \models \xi$*, we define the* φ*-commitment* $\mathcal{Y}_{s,\xi}$ *where*

- *for every* $a \in AP(\varphi)$*, we have that* $\mathcal{Y}_{s,\xi}(a) = 1$ *if* $s \models a$*, and* $\mathcal{Y}_{s,\xi}(a) = 0$ *otherwise,*
- $\mathcal{Y}_{s,\xi}(\Phi) = \mathbb{P}_s(\{\pi \in Run(s) \mid \pi \models \Phi\})$ *for all* $\Phi \in psub(\varphi)$ *such that* $Cl(\xi, s)$ *contains a formula of the form* $P(\Phi) \triangleright r$*,*
- $\mathcal{Y}_{s,\xi}(\Phi) = 0$ *for the other* $\Phi \in psub(\varphi)$*.*

The set Com_φ *consists of all* $\mathcal{Y}_{s,\xi}$*.*

Note that Com_φ is eligible, because every element of Com_φ is finitely satisfiable and $\mathcal{Y}_{s,\varphi}$ implies φ for every s such that $s \models \varphi$.

Example 4. *For the formula φ and the state s of our running example, we obtain*

$$C(\varphi, s) \quad = \quad \{\varphi, \quad \mathrm{G}_{=1}\big(\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\,\neg a) \vee a\big), \quad \mathrm{F}_{=1}\ \mathrm{G}_{=1}\ a, \quad \neg a\}$$

Observe that although $\mathrm{F}_{=1}\ \mathrm{G}_{=1}\ a \in C(\varphi, s)$, the formula $\mathrm{G}_{=1}\ a$ is not included into $C(\varphi, s)$ because $s \not\models \mathrm{G}_{=1}\ a$. Furthermore,

- $\mathcal{Y}_{s,\varphi}\big(\,\mathrm{G}\,\big(\,\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\,\neg a) \vee a\big)\big) = 1,$
- $\mathcal{Y}_{s,\varphi}(\mathrm{F}\ \mathrm{G}_{=1}\ a) = 1,$

and all other components of $\mathcal{Y}_{s,\varphi}$ (incl. $\mathcal{Y}_{s,\varphi}(a)$) are zero.

Now we introduce several notions needed to define a loop-progressive PCTL fragment.

Consider a φ-commitment $\mathcal{Y}_{s,\xi} \in \mathit{Com}_\varphi$ where s is a state of some finite Markov chain. Intuitively, we aim to simplify $\mathcal{Y}_{s,\xi}$ not only by constructing appropriate successors of s, but also by connecting a suitable "predecessor" loop $\mathscr{L}$ with states $\ell_0, \ldots, \ell_n$ and one exit edge leading to s (see Fig. 3).

We start by formalizing the notion of a progress loop.

Definition 11 (progress loop). *Let φ be a finitely satisfiable PCTL formula and $\mathcal{Y}_{s,\xi} \in \mathit{Com}_\varphi$. A* progress loop *for $\mathcal{Y}_{s,\xi}$ is a finite sequence $\mathscr{L} = L_0, \ldots, L_n$ of subsets of $sub(\mathcal{Y}_{s,\xi})$ satisfying the following conditions:*

(1) $\mathcal{Y}_{s,\xi} \subseteq L_i$ for some $i \in \{0, \ldots, n\}$;

(2) $L_0, \ldots, L_n$ are pairwise different (this induces an upper bound on n);

(3) for every $i \in \{0, \ldots, n\}$, we have that

- *if $a \in L_i$, then $\neg a \notin L_i$;*
- *if $\psi_1 \wedge \psi_2 \in L_i$, then $\psi_1, \psi_2 \in L_i$;*
- *if $\psi_1 \vee \psi_2 \in L_i$, then $\psi_1 \in L_i$ or $\psi_2 \in L_i$;*
- *if $\mathrm{G}_{\rhd r}\,\psi \in L_i$, then $\psi \in L_j$ for every $j \in \{0, \ldots, n\}$.*

Definition 11 by itself does not guarantee any progress in satisfying $\mathcal{Y}_{s,\xi}$. This is achieved by additional conditions specified later. A better intuitive understanding of the actual purpose of a progress loop can be developed after introducing the necessary notions (see Remark 3).

Example 5. *Consider the φ-commitment $\mathcal{Y}_{s,\varphi}$ of our running example. Then L_0, L_1, where*

$$\begin{array}{rcl} L_0 & = & \{\varphi,\ \mathrm{G}_{=1}\left(\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a) \vee a\right),\ \mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a) \vee a, \\ & & \mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a),\ \mathrm{F}_{=1}\mathrm{G}_{=1}a,\ \neg a\} \\ L_1 & = & \{\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a) \vee a,\ \mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a), a \wedge \mathrm{F}_{\geq 0.2}\neg a,\ a,\ \mathrm{F}_{\geq 0.2}\neg a\} \end{array}$$

is a progress loop for $\mathcal{Y}_{s,\varphi}$. Note that L_0 is the least *set containing φ satisfying the closure properties of Definition 11. The set L_1 then inevitably contains the formula $\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a) \vee a$, but the conditions of Definition 11 would be satisfied even if we used*

$$L_1' \quad = \quad \{\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a) \vee a, a\}$$

instead of L_1. However, we will impose additional restrictions on progress loops preventing using L_1'.

For a given progress loop $\mathscr{L}$, we use $\Delta(\mathscr{L})$ to denote the set of all formulae occurring in $\mathscr{L}$ whose satisfaction is not guaranteed by $\mathscr{L}$ itself. More precisely, the set $\Delta(\mathscr{L})$ is defined as follows:

Definition 12 (the set $\Delta(\mathscr{L})$). *Let $\mathscr{L} = L_0, \ldots, L_n$ be a progress loop for $\mathcal{Y}_{s,\xi} \in Com_\varphi$. The set $\Delta(\mathscr{L})$ is the union of all $\psi \in L_0 \cup \cdots \cup L_n$ such that one of the following conditions holds:*

- $\psi \equiv \mathrm{G}_{\rhd r}\,\varrho$;
- $\psi \equiv \mathrm{F}_{\rhd r}\,\varrho$ *and* $\varrho \notin L_0 \cup \cdots \cup L_n$;
- $\psi \equiv \mathrm{F}_{=1}\,\varrho$ *and* $\mathrm{F}_{=1}\,\varrho \in L_i$ *for some i such that* $\varrho \notin L_i \cup \cdots \cup L_n$.

Example 6. *Consider the progress loop $\mathscr{L} = L_0, L_1$ for $\mathcal{Y}_{s,\varphi}$ of Example 5. Then*

$$\Delta(\mathscr{L}) \quad = \quad \{\mathrm{G}_{=1}\left(\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a) \vee a\right),\ \mathrm{F}_{=1}\mathrm{G}_{=1}a\}\,.$$

Note that $\mathcal{Y}_{s,\Delta(\mathscr{L})} = \mathcal{Y}_{s,\varphi}$. For the progress loop $\mathscr{L}' = L_0, L_1'$ of Example 5, we obtain

$$\Delta(\mathscr{L}') \quad = \quad \Delta(\mathscr{L}) \cup \{\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\neg a)\}$$

Now we define loop progressive PCTL fragments.

Definition 13 (loop progressive PCTL fragment). *A PCTL fragment $\mathcal{L}$ is* loop progressive *if for every finitely satisfiable $\varphi \in \mathcal{L}$ and every $\mathcal{Y}_{s,\xi} \in Com_\varphi$ such that $\mathcal{H}(\mathcal{Y}_{s,\xi}) > 0$ there exists a progress loop $\mathscr{L} = L_0, \ldots, L_n$ such that*

(1) $\Delta(\mathscr{L}) \in \mathcal{L}$;

(2) $s \models \Delta(\mathscr{L})$;

(3) $s \not\models \varrho$ *for every formula of the form* $\mathrm{F}_{\triangleright r}\, \varrho$ *such that* $\mathrm{F}_{\triangleright r}\, \varrho \in \Delta(\mathscr{L})$;

(4) $deg_s(\Delta(\mathscr{L})) \subset deg_s(\mathcal{Y}_{s,\xi})$ *or* $cf_s(\Delta(\mathscr{L})) \subseteq cf_s(\mathcal{Y}_{s,\xi})$.
Here, the sets $deg_s(X)$ and $cf_s(X)$, where X is a set of PCTL fomulae, are defined as follows:

- *$deg_s(X)$ consists of all formulae $\mathrm{G}\, \varrho$ such that $sub(X)$ contains a formula of the form $\mathrm{G}_{\triangleright r}\, \varrho$ and $s \not\models \mathrm{G}_{=1}\, \varrho$;*
- *$cf_s(X)$ consists of all formulae $\mathrm{F}\, \varrho$ such that X contains a formula of the form $\mathrm{F}_{\triangleright r}\, \varrho$, $s \not\models \varrho$, and there is a finite path from s to a state t where $t \models \varrho$ and $deg_t(X) = deg_s(X)$.*

Remark 2. *Intuitively, the sets $deg_s(X)$ and $cf_s(X)$ represent the "complexity of* G *and* F *requirements imposed by X on s". More specifically,*

- *$deg_s(X)$ contains $\mathrm{G}\, \varrho$ formulae that are "mentioned" in X as subformulae but are* not *satisfied in s with probability one.*

- *$cf_s(X)$ contains some of the formulae $\mathrm{F}\, \varrho$ where $\mathrm{F}_{\triangleright r}\, \varrho$ occurs in X. Intuitively, we do* not *need to include such a formula into $cf_s(X)$ if either $s \models \varrho$, or every state t reachable from s such that $t \models \varrho$ satisfies $deg_t(X) \subset deg_s(X)$. In the first case, the $\mathrm{F}_{\triangleright r}\, \varrho$ requirement is satisfied immediately in s. In the second case, the future satisfaction of ϱ* inevitably *causes a decrease in the complexity of* G *requirements measured by deg.*

Let us illustrate the technical conditions of Definition 13 on our running example.

Example 7. *Consider the φ-commitment $\mathcal{Y}_{s,\varphi}$. Recall that $\mathcal{Y}_{s,\varphi}$ can be interpreted as a set of formulae*

$$\mathcal{Y}_{s,\varphi} \;=\; \{\neg a, \quad \mathrm{G}_{=1}\left(\mathrm{F}_{\geq 0.5}(a \wedge \mathrm{F}_{\geq 0.2}\, \neg a) \vee a\right), \quad \mathrm{F}_{=1}\, \mathrm{G}_{=1}\, a\}$$

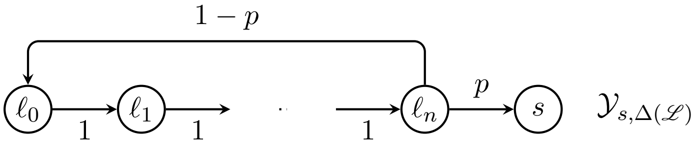

Figure 3: A graph for a progress loop $L_0, \ldots, L_n$.

We have that

$$\begin{aligned} deg_s(\mathcal{Y}_{s,\varphi}) &= \{\mathrm{G}\, a\} \\ cf_s(\mathcal{Y}_{s,\varphi}) &= \emptyset \end{aligned}$$

In particular, note that the formula $\mathrm{F}\,\mathrm{G}_{=1}\, a$ *does* not *belong to* $cf_s(\mathcal{Y}_{s,\varphi})$, *because the only state* t *reachable from* s *such that* $t \models \mathrm{G}_{=1}\, a$ *also satisfies* $deg_t(\mathcal{Y}_{s,\varphi}) = \emptyset$, *hence* $deg_t(\mathcal{Y}_{s,\varphi}) \subset deg_s(\mathcal{Y}_{s,\varphi})$ *(see Fig. 1).*

Furthermore, for the progress loops $\mathscr{L}, \mathscr{L}'$ *of Example 5, we obtain*

$$\begin{aligned} deg_s(\Delta(\mathscr{L})) &= \{\mathrm{G}\, a\} \\ cf_s(\Delta(\mathscr{L})) &= \emptyset \\ deg_s(\Delta(\mathscr{L}')) &= \{\mathrm{G}\, a\} \\ cf_s(\Delta(\mathscr{L}')) &= \{\mathrm{F}(a \wedge \mathrm{F}_{\geq 0.2}\, \neg a)\} \end{aligned}$$

Hence, the conditions of Definition 13 are satisfied for the progress loop $\mathscr{L}$ *but not for the progress loop* $\mathscr{L}'$.

Remark 3. *Example 7 also reveals the actual purpose of progress loops. States of a progress loop* $\mathscr{L}$ *are used to ensure that satisfaction of formulae of the form* $\mathrm{F}_{\triangleright r}\, \varrho$ *generated by* G*-formulae occurring in* $\mathcal{Y}_{s,\xi}$ *happens on the progress loop itself. In other words,* $\mathscr{L}$ *needs to be chosen so that* $\Delta(\mathscr{L})$ *does not contain any "new" formulae of the form* $\mathrm{F}_{\triangleright r}\, \varrho$ *that are not contained in* $\mathcal{Y}_{s,\xi}$. *The only exceptions are when*

- $deg_s(\Delta(\mathscr{L})) \subset deg_s(\mathcal{Y}_{s,\xi})$, *which is a sufficiently strong progress indicator by itself (this condition may hold even though* $\mathcal{Y}_{s,\xi}$ *is included in some state of* $\mathscr{L}$, *see Definition 12);*
- *every state* t *reachable from* s *such that* $t \models \varrho$ *satisfies* $deg_t(\Delta(\mathscr{L})) \subset deg_s(\Delta(\mathscr{L}))$. *Then,* $\mathrm{F}_{\triangleright r}\, \varrho$ *does not contribute to the "complexity of* F*-requirements imposed by* $\Delta(\mathscr{L})$ *on* s*" in the sense of Remark 2.*

A proof of the next theorem shows how to construct the descendants of s in a progress loop for $\mathcal{Y}_{s,\xi} \in \mathit{Com}_\varphi$ together with a safe and progressive assignment.

Theorem 7. *If a PCTL fragment $\mathcal{L}$ is loop progressive, then $\mathcal{L}$ is effectively progressive.*

Proof. Let $\mathcal{L}$ be a loop progressive PCTL fragment. For every finitely satisfiable $\varphi \in \mathcal{L}$, we put $c(\varphi) = 2^{|psub(\varphi)|}$. The functions h and g required in Definition 7 are defined later.

For the rest of this proof, we fix a finitely satisfiable $\varphi \in \mathcal{L}$. Let $\mathcal{Y}_{s,\xi} \in \mathit{Com}_\varphi$, where s is a state of a finite Markov chain N and $\mathcal{H}(\mathcal{Y}_{s,\xi}) > 0$. Let $\mathscr{L} = L_0, \ldots, L_n$ be a progress loop for $\mathcal{Y}_{s,\xi}$ satisfying the conditions of Definition 13. We construct a finite Markov chain M by extending N with fresh states $\ell_0, \ldots, \ell_n$ in the way shown in Fig. 3. The states $\ell_0, \ldots, \ell_n$ correspond to $L_0, \ldots, L_n$ and form the only top SCC C of M where $\mathit{Desc}(C) = \{s\}$. The valuation v of N is extended to M so that $a \in v(\ell_i)$ iff $a \in L_i$. The probability $p > 0$ is chosen so that $1 - p$ is strictly larger than the maximal $r \neq 1$ appearing in formulae of the form $\mathrm{F}_{\triangleright r}\, \varrho \in L_0 \cup \cdots \cup L_n$.

Recall that $\mathcal{Y}_{s,\xi} \subseteq L_i$ for some $0 \leq i \leq n$. First, we show that a C-assigment $\mathcal{A}$ such that $\mathcal{A}(s) = \mathcal{Y}_{s,\Delta(\mathscr{L})}$ is safe for $\mathcal{Y}_{s,\xi}$. This is achieved by proving a *stronger* claim saying that if $\theta \in L_i$ where $0 \leq i \leq n$, then $\ell_i \models \theta$. We proceed by induction on the structure of θ.

- $\theta \equiv a$ or $\theta \equiv \neg a$. If $\theta \in L_i$, then $\ell_i \models \theta$ by the definition of v.
- $\theta \equiv \psi_1 \wedge \psi_2$. If $\psi_1 \wedge \psi_2 \in L_i$, then $\psi_1, \psi_2 \in L_i$, and hence $\ell_i \models \psi_1 \wedge \psi_2$ by induction hypothesis.
- $\theta \equiv \psi_1 \vee \psi_2$. Similarly as above.
- $\theta \equiv \mathrm{G}_{\triangleright r}\, \psi$. If $\mathrm{G}_{\triangleright r}\, \psi \in L_i$, then for every $0 \leq j \leq n$ we have that $\psi \in L_j$ and hence $\ell_j \models \psi$ by induction hypothesis. Furthermore, $\mathrm{G}_{\triangleright r}\, \psi \in \Delta(\mathscr{L})$, hence $s \models \mathrm{G}_{\triangleright r}\, \psi$ because $s \models \Delta(\mathscr{L})$. This implies $\ell_i \models \mathrm{G}_{\triangleright r}\, \psi$.
- $\theta \equiv \mathrm{F}_{\triangleright r}\, \psi$. Then there are three cases:
  - $r < 1$ and $\psi \in L_0 \cup \cdots \cup L_n$. If $\mathrm{F}_{\triangleright r}\, \psi \in L_i$, then $\ell_i \models \mathrm{F}_{\triangleright r}\, \psi$ by our choice of p and induction hypothesis.

- $r < 1$ and $\psi \notin L_0 \cup \cdots \cup L_n$. If $\mathrm{F}_{\triangleright r}\,\psi \in L_i$, then $\mathrm{F}_{\triangleright r}\,\psi \in \Delta(\mathscr{L})$, hence $s \models \mathrm{F}_{\triangleright r}\,\psi$, and thus also $\ell_i \models \mathrm{F}_{\triangleright r}\,\psi$ because the probability of all runs initiated in ℓ_i that visit the state s is equal to one.
- $r = 1$ and for every i such that $\mathrm{F}_{=1}\,\psi \in L_i$ we have that $\psi \in L_i \cup \cdots \cup L_n$. Then $\ell_i \models \mathrm{F}_{=1}\,\psi$ by induction hypothesis.
- $r = 1$ and $\mathrm{F}_{=1}\,\psi \in L_i$ for some i such that $\psi \notin L_i \cup \cdots \cup L_n$. Then $\mathrm{F}_{=1}\,\psi \in \Delta(\mathscr{L})$, hence $s \models \mathrm{F}_{=1}\,\psi$, and we obtain $\ell_i \models \mathrm{F}_{=1}\,\psi$.

Since the φ-commitment $\mathcal{Y}_{s,\Delta(\mathscr{L})}$ is not necessarily different from $\mathcal{Y}_{s,\xi}$ (see Example 6), it is generally not "simpler" than $\mathcal{Y}_{s,\xi}$. Now we show that there exist a finite set T of states of N and a probability distribution ν over T such that

$$\mathcal{Y}_{s,\Delta(\mathscr{L})} \quad \leq \quad \sum_{t \in T} \nu(t) \cdot \mathcal{Y}_{t,\Delta_t(\mathscr{L})}\,. \tag{4}$$

Here, $\Delta_t(\mathscr{L})$ denotes the set $\Delta(\mathscr{L})$ "updated" to t, i.e., $\Delta_t(\mathscr{L})$ is obtained from $\Delta(\mathscr{L})$ by changing every $P(\Phi) \triangleright r \in \Delta(\mathscr{L})$ into $P(\Phi) \geq r'$, where $r' = \mathbb{P}_t(\{\pi \in \mathit{Run}(t) \mid \pi \models \Phi\})$. If $r' = 0$, the formula $P(\Phi) \geq r'$ is *not* included into $\Delta_t(\mathscr{L})$. Observe that $t \models \Delta_t(\mathscr{L})$.

Furthermore, we prove that $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) < g(\mathcal{Y}_{s,\xi})$ for every $t \in T$ such that $\mathcal{H}(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})) \geq 1$ where $g : \mathit{Com}_\varphi \to \mathbb{N}$ is the complexity measure defined below. Thus, the second condition of Definition 7 is established, because we can then replace s with T by setting $\mathit{Desc}(C) = T$ and $P(\ell_n, t) = p \cdot \nu(t)$ for every $t \in T$. Clearly, the C-assignment defined by $\mathcal{A}(t) = \mathcal{Y}_{t,\Delta_t(\mathscr{L})}$ is both g-progressive and safe for $\mathcal{Y}_{s,\xi}$.

Let B be the set of all states t of N such that t either belongs to a BSCC of N or there exists a formula $\mathrm{F}_{\triangleright r}\,\psi \in \Delta(\mathscr{L})$ such that $t \models \psi$. For every $t \in B$, let $\nu(t)$ be the probability of all runs initiated in s visiting the state t so that all states preceding the first visit to t are not contained in B. Let T be the set of all $t \in B$ such that $\nu(t) > 0$. Since the probability of all runs initiated is s that eventually visit a BSCC of N is equal to one, we obtain $\sum_{t \in T} \nu(t) = 1$. Now consider the vector

$$X \quad = \quad \sum_{t \in T} \nu(t) \cdot \mathcal{Y}_{t,\Delta_t(\mathscr{L})}$$

Clearly, for every $\mathrm{F}_{\triangleright r}\,\psi \in \Delta(\mathscr{L})$ we have that $\mathcal{Y}_{s,\Delta(\mathscr{L})}(\mathrm{F}\,\psi) = X(\mathrm{F}\,\psi)$. Furthermore, for every $\mathrm{G}_{\triangleright r}\,\psi \in \Delta(\mathscr{L})$ we obtain $\mathcal{Y}_{s,\Delta(\mathscr{L})}(\mathrm{G}\,\psi) \leq X(\mathrm{G}\,\psi)$. This inequality can be strict because ψ can become invalid along a path from s

before visiting a state of B (the runs initiated by such a path do not satisfy $\mathrm{G}\,\psi$). This proves Inequality (4).

Now we define the function $g : \mathit{Com}_\varphi \to \mathbb{N}$. Recall the functions $\mathit{deg}_s(X)$ and $\mathit{cf}_s(X)$ introduced in Definition 13, and for every $\mathcal{Y}_{u,\varrho} \in \mathit{Com}_\varphi$, we put

$$g(\mathcal{Y}_{u,\varrho}) \quad = \quad |\mathit{deg}_u(\mathcal{Y}_{u,\varrho})| \cdot \Big(1 + \sum_{\Phi \in \mathit{psub}(\mathcal{Y}_{u,\varrho})} \langle \Phi \rangle \Big) + \sum_{\Phi \in \mathit{cf}_u(\mathcal{Y}_{u,\varrho})} \langle \Phi \rangle$$

Here, $\langle \Phi \rangle$ is defined for every path formula Φ of the form $\mathrm{F}\,\eta$ or $\mathrm{G}\,\eta$ inductively by $\langle \Phi \rangle = 1 + \sum_{\Psi \in \mathit{psub}(\eta)} \langle \Psi \rangle$ (the empty sum denotes 0). Observe that if $g(\mathcal{Y}_{u,\varrho}) = 0$, then $s \models \mathrm{G}_{=1}\,\eta$ for every $\mathrm{G}\,\eta \in \mathit{psub}(\mathcal{Y}_{u,\varrho})$, and if $\mathcal{Y}_{u,\varrho}$ contains a formula of the form $\mathrm{F}_{\triangleright r}\,\eta$, then $s \models \eta$. Hence, $\mathcal{Y}_{u,\varrho}$ has a strongly connected model, i.e., $\mathcal{H}(\mathcal{Y}_{u,\varrho}) = 0$.

It remains to show that $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) < g(\mathcal{Y}_{s,\xi})$ for every $t \in T$ such that t does *not* belong to a BSCC of N (if t belongs to a BSCC of N, then $\mathcal{H}(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) = 0$). Realize the following:

(a) $\mathit{deg}_t(\Delta_t(\mathscr{L})) \subseteq \mathit{deg}_s(\Delta(\mathscr{L})) \subseteq \mathit{deg}_s(\mathcal{Y}_{s,\xi})$. This follows directly from the definition of *deg* (also recall that the probability of all runs initiated in s visiting t is positive).

(b) $\sum_{\Phi \in \mathit{psub}(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})} \langle \Phi \rangle \quad \leq \quad \sum_{\Phi \in \mathit{psub}(\mathcal{Y}_{s,\xi})} \langle \Phi \rangle$

According to item (4) of Definition 13, we can distinguish two possibilities:

1. $\mathit{deg}_s(\Delta(\mathscr{L})) \subset \mathit{deg}_s(\mathcal{Y}_{s,\xi})$. Then, we obtain $\mathit{deg}_t(\Delta_t(\mathscr{L})) \subset \mathit{deg}_s(\mathcal{Y}_{s,\xi})$ by applying (a). Also observe $\mathit{deg}_t(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) = \mathit{deg}_t(\Delta_t(\mathscr{L}))$. Hence, we can use (b) to conclude that the difference between the first summand of $g(\mathcal{Y}_{s,\xi})$ and the first summand of $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})$ is at least

$$1 + \sum_{\Phi \in \mathit{psub}(\mathcal{Y}_{s,\xi})} \langle \Phi \rangle$$

Since $\mathit{cf}_t(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) \subseteq \mathit{psub}(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})$, by applying (b) we also obtain that the difference between the second summand of $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})$ and the second summand of $g(\mathcal{Y}_{s,\xi})$ is at most

$$\sum_{\Phi \in \mathit{psub}(\mathcal{Y}_{s,\xi})} \langle \Phi \rangle$$

Hence, $g(\mathcal{Y}_{s,\xi}) - g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) \geq 1$, which proves $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) < g(\mathcal{Y}_{s,\xi})$.

2. $cf_s(\Delta(\mathscr{L})) \subseteq cf_s(\mathcal{Y}_{s,\xi})$. Since $deg_t(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) = deg_t(\Delta_t(\mathscr{L})) \subseteq deg_s(\mathcal{Y}_{s,\xi})$ (see (a)), the difference between the first summand of $g(\mathcal{Y}_{s,\xi})$ and the first summand of $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})$ is *non-negative*. Furthermore, $cf_t(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) \subset cf_s(\Delta(\mathscr{L}))$ because $t \models \psi$ for some $\mathrm{F}_{\rhd r}\,\psi \in \Delta(\mathscr{L})$. This means that the second summand of $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})})$ is strictly smaller than the second summand of $g(\mathcal{Y}_{s,\xi})$, hence $g(\mathcal{Y}_{t,\Delta_t(\mathscr{L})}) < g(\mathcal{Y}_{s,\xi})$. □

Since $g(X) \leq 3 \cdot \|\varphi\|^3$ for every φ-commitment X, we put $h(\varphi) = 3 \cdot \|\varphi\|^3$ to satisfy the first condition of Definition 7. □

An immediate consequence of Theorem 7, Theorem 4, and Proposition 1 is the following:

Theorem 8. *Let $\mathcal{L}$ be a loop progressive PCTL fragment. Then the finite satisfiability problem for $\mathcal{L}$ is in* **2-EXPSPACE**.

Theorem 8 can be applied to various PCTL fragments by demonstrating their loop progressivity, and can be interpreted as a "unifying principle" behind these concrete decidability results. To illustrate this, we give examples of loop progressive fragments in Section 4.1.

4.1. Examples of Loop Progressive Fragments

In this section, we give examples of loop progressive fragments. Let us note that the fragment $\mathcal{L}_1$ of Section 3.1 is also loop progressive. For a finitely satisfiable formula of $\mathcal{L}_1$, a progress loop can always be chosen as a singleton, and hence is not too interesting. In the following abstract syntax equations, the constraint $\rhd r$ has the same meaning as in Definition 1, and $\rhd w$ stands for an arbitrary constraint except for '=1'.

Fragment $\mathcal{L}_2$

$$\begin{array}{lcl}
\varphi & ::= & a \mid \neg a \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \mathrm{F}_{\rhd r}\,\varphi \mid \mathrm{G}_{=1}\,\theta \\
\theta & ::= & a \mid \neg a \mid \theta_1 \wedge \theta_2 \mid \theta_1 \vee \theta_2 \mid \mathrm{F}_{\rhd w}\,\theta
\end{array}$$

Fragment $\mathcal{L}_3$

$$\begin{array}{lcl}
\varphi & ::= & a \mid \neg a \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \mathrm{F}_{\rhd r}\,\varphi \mid \mathrm{G}_{=1}\,\theta \mid \mathrm{G}_{=1}\,\varrho \\
\theta & ::= & a \mid \neg a \mid \theta_1 \wedge \theta_2 \mid \theta_1 \vee \theta_2 \mid \mathrm{F}_{\rhd w}\,\theta \\
\varrho & ::= & \varrho_1 \wedge \varrho_2 \mid \varrho_1 \vee \varrho_2 \mid \mathrm{F}_{\rhd w}\,\theta \mid \mathrm{G}_{=1}\,\theta \mid \mathrm{G}_{=1}\,\varrho
\end{array}$$

Fragment $\mathcal{L}_4$

$$\begin{array}{lcl}
\varphi & ::= & a \ \mid\ \neg a \ \mid\ \varphi_1 \wedge \varphi_2 \ \mid\ \varphi_1 \vee \varphi_2 \ \mid\ \mathrm{F}_{\triangleright r}\, \varphi \ \mid\ \mathrm{G}_{=1}\, \theta \\
\theta & ::= & a \ \mid\ \neg a \ \mid\ \theta_1 \wedge \theta_2 \ \mid\ \theta_1 \vee \theta_2 \ \mid\ \mathrm{F}_{>0}\, \theta \ \mid\ \mathrm{G}_{=1}\, \theta
\end{array}$$

Observe that $\mathcal{L}_2$ and $\mathcal{L}_3$ contain the formula φ defined by (1) in Section 1. The above fragments are chosen so that they are not covered by the results of [6] and illustrate various properties of Definition 13. All of them contain formulae requiring non-bottom SCCs with more than one state.

Proposition 9. *Fragment $\mathcal{L}_2$ is loop progressive.*

PROOF. Let $\varphi \in \mathcal{L}_2$ and $\mathcal{Y}_{s,\xi} \in \mathit{Com}_\varphi$ where $\mathcal{H}(\mathcal{Y}_{s,\xi}) \geq 1$. We need to show the existence of a progress loop $\mathscr{L} = L_0, \ldots, L_n$. We construct the loop inductively so that every L_i is associated with some state t_i reachable from s where $t_i \models L_i$.

The set L_0 is the least set K satisfying the following conditions:

- $\mathcal{Y}_{s,\xi} \subseteq K$;
- if $\psi \in K$, then $C(\psi, s) \subseteq K$;
- if $\mathrm{G}_{\triangleright r}\, \psi \in K$, then $\psi \in K$.

We put $t_0 = s$ (observe that $s \models L_0$). Furthermore, let N be the set of all ϱ such that $\mathrm{G}_{=1}\, \varrho \in L_0$.

Suppose that $L_0, \ldots, L_n$ are the sets constructed so far where $t_i \models L_i$ for every $i \in \{0, \ldots, n\}$. Now we distinguish two possibilities:

- If for every formula of the form $\mathrm{F}_{\triangleright r}\, \varrho \in L_0 \cup \ldots \cup L_n$ where $\mathrm{F}_{\triangleright r}\, \varrho \notin \mathcal{Y}_{s,\xi}$ there exists $i \in \{0, \ldots, n\}$ such that $\varrho \in L_i$, then the construction terminates.
- Otherwise, let $\mathrm{F}_{\triangleright r}\, \varrho \in L_i$ be a formula such that $\mathrm{F}_{\triangleright r}\, \varrho \notin \mathcal{Y}_{s,\xi}$ and $\varrho \notin L_0 \cup \ldots \cup L_n$. It follows from the definition of $\mathcal{L}_2$ that $r \neq 1$. Furthermore, $t_i \not\models \varrho$ (this is guaranteed by the closure rules defining L_0 and L_{k+1}, see below). Since $t_i \models \mathrm{F}_{\triangleright r}\, \varrho$, there exists a state t reachable from t_i (and hence also from s) such that $t \models \varrho$. Moreover, $t \models N$. Now, we construct L_{n+1} as the least set K satisfying the following conditions:

- $\varrho \in K$;
- $N \subseteq K$;
- if $\psi \in K$, then $C(\psi, t) \subseteq K$.

Observe that if $\mathrm{G}_{=1}\,\psi \in K$, then $\mathrm{G}_{=1}\,\psi \in N$ because ϱ does not contain any subformula of the form $\mathrm{G}_{\rhd r}\,\psi$ (this follows directly from the definition of $\mathcal{L}_2$). Furthermore, $t \models L_{n+1}$.

Note that if $\mathrm{F}_{\rhd r}\,\psi \in \Delta(\mathscr{L})$, then this formula belongs also to $\mathcal{Y}_{s,\xi}$. Now it is easy to verify that the constructed $\mathscr{L} = L_0, \ldots, L_n$ is a progress loop. $\square$

The argument for $\mathcal{L}_3$ is similar, but a progress loop needs to be constructed somewhat differently. Let $\mathcal{Y}_{s,\xi} \in \mathit{Com}_\varphi$ where $\mathcal{H}(\mathcal{Y}_{s,\xi}) \geq 1$. For every state t reachable from s, let L_t be the least set K satisfying the following conditions, where θ and ϱ range over the sets of formulae defined by the corresponding abstract syntax equations in the definition of $\mathcal{L}_3$:

(i) if there is $\mathrm{F}_{\rhd w}\,\theta \in \mathit{sub}(\mathcal{Y}_{s,\xi})$ such that $t \models \theta$, then $E(\theta, t) \subseteq K$;

(ii) if there is $\mathrm{G}_{=1}\,\theta \in \mathit{sub}(\mathcal{Y}_{s,\xi})$ such that $t \models \theta$, then $E(\theta, t) \subseteq K$;

(iii) if there is $\mathrm{G}_{=1}\,\varrho \in \mathit{sub}(\mathcal{Y}_{s,\xi})$ such that $s \models \varrho$, then $E(\varrho, s) \subseteq K$;

(iv) if $t = s$, then $C(\mathcal{Y}_{s,\xi}, s) \subseteq K$.

Note that $L_t \subseteq \mathit{sub}(\mathcal{Y}_{s,\xi})$. The syntactic restrictions of $\mathcal{L}_3$ imply the following:

- The closure rule *(iii)* cannot add new formulae of the form a or $\neg a$ into L_t.
- The closure rules *(i)* and *(ii)* cannot add new formulae of the form $G_{\rhd r}\lambda$ into L_t.

Let $\{L_0, \ldots, L_n\}$ be the set of all L_t such that t is reachable from s (we assume $L_i \neq L_j$ for $i \neq j$). We put $\mathscr{L} = L_0, \ldots, L_n$. The conditions of Definition 13 are now easy to verify.

A progress loop for $\mathcal{L}_4$ is constructed by altering the definition of L_t above and considering only a suitable subset of states reachable from s.

5. Conclusions

We have designed a general method for proving the decidability of the finite satisfiability problem for PCTL fragments. As an application, we derived a meta-result about loop progressive fragments with four concrete instances $\mathcal{L}_1, \ldots, \mathcal{L}_4$. A natural continuation of our work is considering fragments requiring a more complex shape of a progress-achieving SCC. Natural candidates are loops with several exit states, and SCCs with arbitrary topology but one exit state. Here, increasing the probability of satisfying $\mathrm{F}\,\varphi$ subformulae can be "traded" for decreasing the probability of satisfying $\mathrm{G}\,\varphi$ subformulae, and understanding this phenomenon is another important step towards solving the finite satisfiability problem for the whole PCTL.

Acknowledgements

This work is supported by the Czech Science Foundation, grant No. 21-24711S.

Appendix A. Encoding PCTL bounded satisfiability in existential theory of the reals

In this section, we sketch a (non-deterministic) polynomial space algorithm deciding bounded PCTL satisfiability. Let φ be a PCTL formula and $n \in \mathbb{N}$ a bound on the size of the model. Without restrictions[3], we assume that φ is constructed according to the abstract syntax equation

$$\varphi \ ::= \ a \mid \neg a \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \mathrm{F}_{\bowtie r}\,\varphi$$

where $\bowtie\ \in \{\geq, >, \leq, <\}$. We disregard the trivial probability constraints '$\geq$0', '$>$1', '$<$0', and '$\leq$ 1'.

The algorithm starts by guessing a finite directed graph $(V, \rightarrow)$, where $V = \{v_1, \ldots, v_m\}$ and $m \leq n$. Furthermore, for every subformula $\psi \in \mathit{sub}(\varphi)$, the algorithm guesses a subset $V(\psi) \subseteq V$ so that

- $V(a) = V \setminus V(\neg a)$ for every atomic proposition a such that $\neg a \in \mathit{sub}(\varphi)$;
- $V(\xi_1 \wedge \xi_2) = V(\xi_1) \cap V(\xi_2)$ for every $\xi_1 \wedge \xi_2 \in \mathit{sub}(\varphi)$;
- $V(\xi_1 \vee \xi_2) = V(\xi_1) \cup V(\xi_2)$ for every $\xi_1 \vee \xi_2 \in \mathit{sub}(\varphi)$;

[3] Observe that every occurrence of $\mathrm{G}_{\rhd r}\,\varphi$ can be replaced with $\mathrm{F}_{\lhd r}\,\neg\varphi$.

- $V(\varphi) \neq \emptyset$.

Then, the algorithm constructs the following formula of existential theory of the reals, where $k = |E|$ and $\mathit{Fsub}(\varphi)$ is the set of all subformulae of φ of the form $\mathrm{F}_{\bowtie r}\, \psi$.

$$\exists x_1, \ldots, x_k \quad : \quad \bigwedge_{i=1}^{n} 0 < x_i \leq 1 \quad \wedge \quad \bigwedge_{v \in V} \mathit{Distr}(v) \quad \wedge \quad \bigwedge_{\psi \in \mathit{Fsub}(\varphi)} \mathit{Correct}(V(\psi))$$

The variables $x_1, \ldots, x_n$ represent the (positive) edge probabilities. We write $v_i \xrightarrow{x_t} v_j$ to indicate that x_t represents the probability of $v_i \to v_j$.

The formula $\mathit{Distr}(v)$ says that the sum of the variables associated with the outgoing edges of v is equal to 1, i.e.,

$$\sum_{v \xrightarrow{x_t} v_j} x_t = 1$$

The formula $\mathit{Correct}(V(\mathrm{F}_{\bowtie r}\, \psi))$ says that the set of vertices satisfying the formula $\mathrm{F}_{\bowtie r}\, \psi$ is precisely $V(\mathrm{F}_{\bowtie r}\, \psi)$, assuming that $V(\psi)$ is correct.

$$\begin{aligned} \exists y_1, \ldots, y_n \quad : \quad & \bigwedge_{v_i \in V(\psi)} y_i{=}1 \quad \wedge \bigwedge_{v_i \in \mathit{Out}(V(\psi))} y_i{=}0 \\ \wedge \quad & \bigwedge_{v_i \in \mathit{Other}(V(\psi))} y_i = \sum_{v_i \xrightarrow{x_t} v_j} x_t \cdot y_j \\ \wedge \quad & \bigwedge_{v_i \in V(\mathrm{F}_{\bowtie r}\, \psi)} y_i \bowtie r \quad \wedge \bigwedge_{v_i \notin V(\mathrm{F}_{\bowtie r}\, \psi)} y_i \not\bowtie r \end{aligned}$$

Here, $\mathit{Out}(V(\psi))$ is the set of all vertices $v \in V$ such that there is no path from v to a state of $V(\psi)$ in $(V, \to)$, and $\mathit{Other}(V(\psi)) = V \setminus (V(\psi) \cup \mathit{Out}(V(\psi)))$. Hence, the variable y_i represents the probability of all runs initiated in v_i visiting a vertex in $V(\psi)$.

Observe that the constructed formula belongs to existential theory of the reals and its size is polynomial in the size of φ and n. Our algorithm outputs 'yes' or 'no' depending on whether the formula is valid or not (which is decidable in space polynomial in the size of φ and n [24]). Thus, the existence of a model of φ with at most n states is decided in polynomial space.

References

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

[2] E. Emerson, Temporal and modal logic, Handbook of Theoretical Computer Science B (1991) 995–1072.

[3] T. Brázdil, V. Forejt, J. Křetínský, A. Kučera, The satisfiability problem for probabilistic CTL, in: Proceedings of LICS 2008, IEEE Computer Society Press, 2008, pp. 391–402.

[4] S. Kraus, D. Lehmann, Decision procedures for time and chance (extended abstract), in: Proceedings of 24th Annual Symposium on Foundations of Computer Science, IEEE Computer Society Press, 1983, pp. 202–209.

[5] T. Brázdil, V. Forejt, J. Křetínský, A. Kučera, The satisfiability problem for probabilistic CTL, Technical report FIMU-RS-2008-03, Faculty of Informatics, Masaryk University (2008).

[6] J. Křetínský, A. Rotar, The satisfiability problem for unbounded fragments of probabilistic CTL, in: Proceedings of CONCUR 2018, Vol. 118 of Leibniz International Proceedings in Informatics, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018, pp. 32:1–32:16.

[7] S. Chakraborty, J. Katoen, On the satisfiability of some simple probabilistic logics, in: Proceedings of LICS 2016, 2016, pp. 56–65.

[8] E. Emerson, J. Halpern, Decision procedures and expressiveness in the temporal logic of branching time, in: Proceedings of STOC'82, ACM Press, 1982, pp. 169–180.

[9] B. Banieqbal, H. Barringer, Temporal logic with fixed points, in: Temporal Logic in Specification, Vol. 398 of Lecture Notes in Computer Science, Springer, 1987, pp. 62–74.

[10] M. Fischer, R. Ladner, Propositional dynamic logic of regular programs, Journal of Computer and System Sciences 18 (1979) 194–211.

[11] D. Lehman, S. Shelah, Reasoning with time and chance, Information and Control 53 (1982) 165–198.

[12] S. Hart, M. Sharir, Probabilistic temporal logic for finite and bounded models, in: Proceedings of POPL'84, ACM Press, 1984, pp. 1–13.

[13] N. Bertrand, J. Fearnley, S. Schewe, Bounded satisfiability for PCTL, in: Proceedings of CSL 2012, Vol. 16 of Leibniz International Proceedings in Informatics, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2012, pp. 92–106.

[14] C. Baier, M. Kwiatkowska, Model checking for a probabilistic branching time logic with fairness, Distributed Computing 11 (3) (1998) 125–155.

[15] A. Bianco, L. de Alfaro, Model checking of probabilistic and nondeterministic systems, in: Proceedings of FST&TCS'95, Vol. 1026 of Lecture Notes in Computer Science, Springer, 1995, pp. 499–513.

[16] M. Huth, M. Kwiatkowska, Quantitative analysis and model checking, in: Proceedings of LICS'97, IEEE Computer Society Press, 1997, pp. 111–122.

[17] C. Baier, J.-P. Katoen, Principles of Model Checking, The MIT Press, 2008.

[18] J. Esparza, A. Kučera, R. Mayr, Model-checking probabilistic pushdown automata, Logical Methods in Computer Science 2 (1:2) (2006) 1–31.

[19] T. Brázdil, A. Kučera, O. Stražovský, On the decidability of temporal properties of probabilistic pushdown automata, in: Proceedings of STACS 2005, Vol. 3404 of Lecture Notes in Computer Science, Springer, 2005, pp. 145–157.

[20] K. Etessami, M. Yannakakis, Model checking of recursive probabilistic systems, ACM Transactions on Computational Logic 13 (2012).

[21] T. Brázdil, V. Forejt, A. Kučera, Controller synthesis and verification for Markov decision processes with qualitative branching time objectives, in: Proceedings of ICALP 2008, Part II, Vol. 5126 of Lecture Notes in Computer Science, Springer, 2008, pp. 148–159.

[22] T. Brázdil, V. Brožek, V. Forejt, A. Kučera, Stochastic games with branching-time winning objectives, in: Proceedings of LICS 2006, IEEE Computer Society Press, 2006, pp. 349–358.

[23] P. Billingsley, Probability and Measure, Wiley, 1995.

[24] J. Canny, Some algebraic and geometric computations in PSPACE, in: Proceedings of STOC'88, ACM Press, 1988, pp. 460–467.