

Parametric LTL on Markov Chains

Souymodip Chakraborty and Joost-Pieter Katoen^{*}

RWTH Aachen University, Ahornstraße 55, D-52074 Aachen, Germany

Abstract. This paper is concerned with the verification of finite Markov chains against parametrized LTL (pLTL) formulas. In pLTL, the until-modality is equipped with a bound that contains variables; e.g., $\Diamond_{\leq x} \varphi$ asserts that φ holds within x time steps, where x is a variable on natural numbers. The central problem studied in this paper is to determine the set of parameter valuations $V_{\prec p}(\varphi)$ for which the probability to satisfy pLTL-formula φ in a Markov chain meets a given threshold $\prec p$, where \prec is a comparison on reals and p a probability. As for pLTL determining the emptiness of $V_{>0}(\varphi)$ is undecidable, we consider several logic fragments. We consider parametric reachability properties, a sub-logic of pLTL restricted to next and $\Diamond_{\leq x}$, parametric Büchi properties and finally, a maximal subclass of pLTL for which emptiness of $V_{>0}(\varphi)$ is decidable.

1 Introduction

Verifying a finite Markov chain (MC, for short) M against an LTL-formula φ amounts to determining the probability that M satisfies φ , i.e., the likelihood of the set of infinite paths of M satisfying φ . Vardi [1] considered the qualitative version of this problem, that is, does M almost surely satisfy φ , or with positive probability. Together with Wolper, he showed that the qualitative LTL model-checking problem for MCs is PSPACE-complete. The quantitative verification problem – what is the probability of satisfying φ ? – has been treated by Courcoubetis and Yannakakis [2]. An alternative algorithm that has a time complexity which is polynomial in the size of the MC and exponential in $|\varphi|$ is by Couvreur *et al.* [3]. Recently, practical improvements have been obtained by Chatterjee *et al.* for verifying the LTL(F,G)-fragment on MCs using generalized deterministic Rabin automata [4].

This paper considers the verification of MCs against *parametric* LTL formulas. In parametric LTL [5] (pLTL, for short), temporal operators can be subscripted by a variable ranging over the natural numbers. The formula $\Diamond_{\leq x} a$ means that in at most x steps a occurs, and $\Box \Diamond_{\leq y} a$ means that at every index a occurs within y steps. Note that x and y are variables whose value is not fixed in advance. The central question is now to determine the values of x and y such that the probability of a given MC satisfying the pLTL-formula φ meets a certain

* Currently on sabbatical leave at the University of Oxford, United Kingdom.

threshold p . This is referred to as the valuation set $V_{\prec p}(\varphi)$ for comparison operator \prec . This problem has both a qualitative (threshold > 0 and $= 1$) and a quantitative variant ($0 < p < 1$).

The main results of this paper are as follows. Just as for the setting with Kripke structures [5], it is shown that checking the emptiness of $V_{>0}(\varphi)$ in general is undecidable. We therefore resort to fragments of pLTL. We show that determining $V_{\geq p}(\Diamond_{\leq x} a)$ can be done by searching in a range defined by the *precision* of the input, whereas polynomial time graph algorithms suffice for its qualitative variant. The same applies to formulas of the form $\Box \Diamond_{\leq x} a$. We provide necessary and sufficient criteria for checking the emptiness of $V_{>0}(\varphi)$ (and $V_{=1}(\varphi)$) for the fragments pLTL(F,X) and pLTL $_{\Diamond}$, and prove that checking these criteria are NP-complete and PSPACE-complete, respectively. We also define a representation of these sets and provide algorithms to construct them.

Related work. The verification of parametric probabilistic models in which certain transition probabilities are given as parameters (or functions thereof) has recently received considerable attention. Most of these works are focused on parameter synthesis: for which parameter instances does a given (LTL or PCTL) formula hold? To mention a few, Han *et al.* [6] considered this problem for timed reachability in continuous-time MCs, Hahn *et al.* [7] and Pugelli *et al.* [8] for Markov decision processes (MDPs), and Benedikt *et al.* [9] for ω -regular properties of interval MCs. Hahn *et al.* [10] provide an algorithm for computing the rational function expressing the probability of reaching a given set of states in a parametric (reward) MDP based on exploiting regular expressions as initially proposed by Daws [11]. Other related work includes the synthesis of loop invariants for parametric probabilistic programs [12]. To the best of our knowledge, verifying parametric properties on MCs has not been considered so far. The closest related works are on combining two-variable FO with LTL for MDPs by Benedikt *et al.* [13] and the computation of quantiles by Ummels and Baier [14].

Organization of the paper. Section 2 presents pLTL and MCs and a first undecidability result. Section 3 considers parametric reachability. Section 4 treats the fragment pLTL(F,X) and Section 5 parametric Büchi properties. Section 6 treats the bounded always-free fragment of pLTL. Section 7 concludes the paper. Full version of the paper can be found in the archive.

2 Preliminaries

Parametric LTL. Parametric LTL extends propositional LTL with bounded temporal modalities, for which the bound is either a constant or a variable. Let Var be a finite set of variables ranged over by x, y , and AP be a finite set of propositions ranged over by a and b . Let $c \in \mathbb{N}$. Parametric LTL formulas adhere to the following syntax:

$$\varphi ::= a \mid \neg \varphi \mid \varphi \wedge \varphi \mid \bigcirc \varphi \mid \varphi \mathsf{U} \varphi \mid \Diamond_{\leq x} \varphi \mid \Diamond_{\leq c} \varphi$$

where $\prec \in \{=, \leqslant, <, >, \geqslant\}$. A pLTL structure is a triple (w, i, \mathbf{v}) where $w \in \Sigma^\omega$ with $\Sigma = 2^{AP}$ is an infinite word over sets of propositions, $i \in \mathbb{N}$ is an index, and $\mathbf{v} : \text{Var} \rightarrow \mathbb{N}$ is a variable valuation. Analogously, we consider a valuation \mathbf{v} as a vector in \mathbb{N}^d , where d for pLTL formula φ is the number of variables occurring in φ . E.g. for $d = 1$, the valuation is just a number v . We compare valuations \mathbf{v} and \mathbf{v}' as $\mathbf{v} \leqslant \mathbf{v}'$ iff $\mathbf{v}(x) \leqslant \mathbf{v}'(x)$ for all x . Let $w[i]$ denote the i -th element of w . The satisfaction relation \models is defined by structural induction over φ as follows:

$$\begin{aligned} (w, i, \mathbf{v}) \models a &\quad \text{iff } a \in w[i] \\ (w, i, \mathbf{v}) \models \neg \varphi &\quad \text{iff } (w, i, \mathbf{v}) \not\models \varphi \\ (w, i, \mathbf{v}) \models \varphi_1 \wedge \varphi_2 &\quad \text{iff } (w, i, \mathbf{v}) \models \varphi_1 \text{ and } (w, i, \mathbf{v}) \models \varphi_2 \\ (w, i, \mathbf{v}) \models \Diamond_{\prec_x} \varphi &\quad \text{iff } (w, j, \mathbf{v}) \models \varphi \text{ for some } j \prec \mathbf{v}(x)+i. \end{aligned}$$

For the sake of brevity, we have omitted the semantics of the standard LTL modalities. As usual, $\varphi_1 R \varphi_2 \equiv \neg(\neg \varphi_1 U \neg \varphi_2)$, $\Diamond \varphi \equiv \text{true} U \varphi$ and $\Box \varphi \equiv \neg \Diamond \neg \varphi$. The language of φ is defined by $\mathcal{L}(\varphi) = \{(w, \mathbf{v}) \mid (w, 0, \mathbf{v}) \models \varphi\}$. Alur et al. [5] have shown that other modalities such as $U_{\leqslant x}$, $\Diamond_{>x}$, $\Box_{>x}$, $U_{>x}$, $R_{\leqslant x}$ and $R_{>x}$, can all be encoded in our syntax. For instance, the following equivalences hold:

$$\begin{aligned} \Diamond_{>x} \varphi &\equiv \Box_{\leqslant x} \Diamond \bigcirc \varphi, & \Box_{>x} \varphi &\equiv \Diamond_{\leqslant x} \Box \bigcirc \varphi, \\ \varphi U_{\leqslant x} \psi &\equiv (\varphi U \psi) \wedge \Diamond_{\leqslant x} \psi, & \varphi U_{>x} \psi &\equiv \Box_{\leqslant x} (\varphi \wedge \bigcirc (\varphi U \psi)) \end{aligned} \quad (1)$$

In the remainder of this paper, we focus on bounded always and eventualities where all bounds are upper bounds. We abbreviate $\Diamond_{\leqslant x}$ by \Diamond_x and do similar for the other modalities. For valuation \mathbf{v} and pLTL-formula φ , let $\mathbf{v}(\varphi)$ denote the LTL formula obtained from φ by replacing variable x by its valuation $\mathbf{v}(x)$; e.g., $\mathbf{v}(\Diamond_x \varphi)$ equals $\Diamond_{\mathbf{v}(x)} \mathbf{v}(\varphi)$.

Markov chains. A discrete-time Markov chain M is a quadruple (S, \mathbf{P}, s_0, L) where S is a finite set of states with $m = |S|$, $\mathbf{P} : S \times S \rightarrow [0, 1]$ is a stochastic matrix, $s_0 \in S$ an initial state, and $L : S \rightarrow 2^{AP}$ a state-labeling function. $\mathbf{P}(u, v)$ denotes the one-step probability of moving from state u to v . A trajectory (or path) of a Markov chain (MC, for short) M is a sequence $\{s_i\}_{i \geqslant 0}$ such that $\mathbf{P}(s_i, s_{i+1}) > 0$ for all $i \geqslant 0$. A trajectory $\pi = s_0 s_1 s_2 \dots$ induces the trace $\text{trace}(\pi) = L(s_0) L(s_1) L(s_2) \dots$. Let $\text{Paths}(M)$ denote the set of paths of MC M . A path π satisfies the pLTL-formula φ under the valuation \mathbf{v} , denoted $\pi \models \mathbf{v}(\varphi)$, whenever $(\text{trace}(\pi), 0, \mathbf{v}) \models \varphi$ (or equivalently, $(\text{trace}(\pi), \mathbf{v}) \in \mathcal{L}(\varphi)$). A finite path (or path fragment) satisfies a formula under a valuation if any infinite extension of it also satisfies the formula. Let \Pr be the probability measure on sets of paths, defined by a standard cylinder construction [1]. The probability of satisfying φ by M under valuation \mathbf{v} is given by $\Pr\{\pi \in \text{Paths}(M) \mid \pi \models \mathbf{v}(\varphi)\}$, generally abbreviated as $\Pr(M \models \mathbf{v}(\varphi))$.

Valuation set. The central problem addressed in this paper is to determine the valuation set of a pLTL formula φ . Let M be an MC, $p \in [0, 1]$ a probability bound, and $\prec \in \{=, \leqslant, <, >, \geqslant\}$. Then we are interested in determining:

$$V_{\prec p}(\varphi) = \{\mathbf{v} \mid \Pr(M \models \mathbf{v}(\varphi)) \prec p\},$$

i.e., the set of valuations under which the probability of satisfying φ meets the bound $\prec p$. In particular, we will focus on the decidability and complexity of the emptiness problem for $V_{\prec p}(\varphi)$, i.e., the decision problem whether $V_{\prec p}(\varphi) = \emptyset$ or not, on algorithms (if any) determining the set $V_{\prec p}(\varphi)$, and on the size of the minimal representation of $V_{\prec p}(\varphi)$. In the qualitative setting, the bound $\prec p$ is either > 0 , or $= 1$.

Proposition 1. *For $\varphi \in pLTL$, the problem if $V_{>0}(\varphi) = \emptyset$ is undecidable.*

Proof. The proof is based on [5, Th. 4.1], see the archived version for details. ■

It follows that deciding whether $V_{=1}(\varphi) = \emptyset$ is undecidable, as $V_{>0}(\varphi) = \emptyset$ iff $V_{=1}(\neg\varphi) \neq \emptyset$. As a combination of $\Diamond_{\leq x}$ and $\Box_{\leq x}$ modalities can encode $\mathbf{U}_{=x}$, e.g.,

$$\neg a \wedge \bigcirc (\neg a \mathbf{U}_{=x} a) \equiv \bigcirc (\neg a \mathbf{U}_{\leq x} a) \wedge (\neg a \mathbf{U}_{>x} a),$$

we will restrict ourselves to fragments of pLTL where each formula is in negative normal form and the only parametrized operator is $\Diamond_{\leq x} \varphi$. We refer to this fragment as $pLTL_{\Diamond}$:

$$\varphi ::= a \mid \neg a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \bigcirc \varphi \mid \varphi \mathbf{U} \varphi \mid \varphi \mathbf{R} \varphi \mid \Box \varphi \mid \Diamond_{\leq x} \varphi \mid \Diamond_{\leq c} \varphi \mid \Box_{\leq c} \varphi. \quad (2)$$

We show it is a sub-logic of pLTL for which the emptiness problem for $V_{>0}(\varphi)$ is decidable. The logic has a favourable *monotonicity* property, i.e.,

Remark 1. For every $pLTL_{\Diamond}$ -formula φ , infinite word w and valuations \mathbf{v}, \mathbf{v}' , $\mathbf{v} \leq \mathbf{v}'$ implies $(w, \mathbf{v}) \models \varphi \implies (w, \mathbf{v}') \models \varphi$.

Here $(w, \mathbf{v}) \models \varphi$ is shorthand for $(w, 0, \mathbf{v}) \models \varphi$. We start off with briefly considering (only) parametric eventualities and then consider the sub-logic $pLTL(F, X)$ restricted to next and \Diamond_x . Later on, we also consider parametric Büchi formulas, and finally, $pLTL_{\Diamond}$.

3 Parametric Reachability

In this section, we consider pLTL-formulas of the form $\Diamond_x a$ for proposition a , or equivalently, $\Diamond_x T$ for the set of target states $T = \{s \in S \mid a \in L(s)\}$. We consider bounds of the form $\geq p$ with $0 < p < 1$. The valuation set of interest is thus $V_{\geq p}(\Diamond_x a)$. Let μ_i be the probability of reaching T within i steps; the sequence $\{\mu_i\}$ is ascending. There can be two cases: (a) the sequence reaches a constant value in m steps (m being the size of Markov chain) or (b) the sequence monotonically increases and converges to μ_{∞} . This makes the emptiness problem for $V_{\geq p}(\Diamond_x a)$ decidable. In the first case, we check $\mu_m \geq p$. In the second case, emptiness is decidable in time polynomial in m , by determining $\mu_{\infty} = \Pr(\Diamond a)$ which can be done by solving a system of linear equations with at most m variables. Then, $V_{\geq p}(\Diamond_x a) \neq \emptyset$ iff $p < \mu_{\infty}$.

Assume in the sequel that T is non-empty. Let $\min V_{\geq p}(\Diamond_x a) = n_0$. The valuation set can thus be represented by n_0 (this gives a minimal representation of the set). Membership queries, i.e., does $n \in V_{\geq p}(\Diamond_x a)$, then simply boil down to checking whether $n_0 \leq n$, which can be done in constant time (modulo the size of n_0). The only catch is that n_0 can be very large if p is close to μ_∞ . A simple example elucidates this fact.

Example 1. Consider the MC M with $S = \{s_0, t\}$, $L(t) = \{a\}$, $L(s_0) = \emptyset$, $\mathbf{P}(s_0, s_0) = \frac{1}{2} = \mathbf{P}(s_0, t)$ and $\mathbf{P}(t, t) = 1$. Then $\Pr(M \models \Diamond_n a) = 1 - (\frac{1}{2})^n$. It follows that $\min V_{\geq p}(\Diamond_x a)$ goes to infinity when p approaches one.

The following bound on n_0 can nonetheless be provided. This bound allows for obtaining the minimum value n_0 by a binary search.

Proposition 2. *For MC M , $\min V_{\geq p}(\Diamond_x a) \leq \log_\gamma(1 - (1 - \gamma)^{\frac{p}{b}})$, where $0 < \gamma < 1$ and $b > 0$.*

Proof. Collapse all a -states into a single state t and make it absorbing (i.e., replace all outgoing transitions by a self-loop with probability one). Let t be the only *bottom strongly connected component* (BSCC) of M (other BSCCs can be safely ignored). Let $\{1, \dots, m\}$ be the states of the modified MC M , with the initial state s_0 and the target state t represented by 1 and m , respectively. Let \mathbf{Q} be the $(m-1) \times (m-1)$ transition matrix of the modified MC without the state t . That is, $\mathbf{Q}(i, j) = \mathbf{P}(i, j)$ iff $j \neq m$ where \mathbf{P} is the transition probability matrix of M . We have the following observation:

1. Let the coefficient of ergodicity $\tau(\mathbf{Q})$ of \mathbf{Q} defined as

$$\tau(\mathbf{Q}) = 1 - \min_{i,j} \left(\sum_k \min\{\mathbf{Q}(i, k), \mathbf{Q}(j, k)\} \right).$$

As \mathbf{Q} is sub-stochastic and no row of \mathbf{Q} is zero, it follows $0 < \tau(\mathbf{Q}) < 1$.

2. Let vector $\mathbf{r}^T = (r_1, \dots, r_{m-1})$ with $r_i = \mathbf{P}(i, m)$, r_{\max} be the maximum element in \mathbf{r} and \mathbf{i}^T be $(1, 0, \dots, 0)$. The probability of reaching the state m from the state 1 in at most $n+1$ steps is the probability of being in some state $i < m$ within n steps and taking the next transition to m :

$$\mu_{n+1} = \sum_{j=0}^{n+1} \mathbf{i}^T \mathbf{Q}^j \mathbf{r} \leq \sum_{j=0}^{n+1} \tau(\mathbf{Q})^j r_{\max}.$$

Let $\tau(\mathbf{Q}) = \gamma$ and $r_{\max} = b$. The integer n_0 is the smallest integer such that $\mu_{n_0} \geq p$, which implies that $b \cdot \frac{1-\gamma^{n_0}}{1-\gamma} \geq p$. This yields $n_0 \leq \log_\gamma(1 - (1 - \gamma)^{\frac{p}{b}})$. ■

As in the non-parametric setting, it follows that (for finite MCs) the valuation sets $V_{>0}(\Diamond_x a)$ and $V_{=1}(\Diamond_x a)$ can be determined by a graph analysis, i.e. no inspection of the transition probabilities is necessary for qualitative parametric reachability properties.

Proposition 3. *The problem $V_{>0}(\Diamond_x a) = \emptyset$ is NL-complete.*

Proof. The problem is the same as reachability in directed graphs. ■

Proposition 4. *The sets $V_{>0}(\Diamond_x a)$ and $V_{=1}(\Diamond_x a)$ can be determined in polynomial time by a graph analysis of MC M .*

Proof. Collapse all the a -states into a target state t and make t absorbing. If $V_{>0}(\Diamond_x a)$ is non-empty, it suffices to determine $\min V_{>0}(\Diamond_x a)$ which equals the length of a shortest path from s_0 to t . To determine whether $V_{=1}(\Diamond_x a)$ is empty or not, we proceed as follows. If a cycle without t is reachable from s_0 , then no finite n exists for which the probability of reaching t within n steps equals one. Thus, $V_{=1}(\Diamond_x a) = \emptyset$. If this is not the case, then the graph of M is a DAG (apart from the self-loop at t), and $\min V_{=1}(\Diamond_x a)$ equals the length of a longest path from s_0 to t . ■

4 The Fragment pLTL(F,X)

This section considers the fragment pLTL(F,X) which is defined by:

$$\varphi ::= a \mid \neg a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \bigcirc \varphi \mid \Diamond \varphi \mid \Diamond_{\leq x} \varphi \mid \Diamond_{\leq c} \varphi$$

Our first result is a necessary and sufficient condition for the emptiness of $V_{>0}(\varphi)$.

Theorem 1. *For $\varphi \in pLTL(F,X)$ and MC M with m states, $V_{>0}(\varphi) \neq \emptyset$ iff $\bar{v} \in V_{>0}(\varphi)$ with $\bar{v}(x) = m \cdot |\varphi|$.*

Proof. Let φ be a pLTL(F,X)-formula and assume $V_{>0}(\varphi) \neq \emptyset$. By monotonicity, it suffices to prove that $v \in V_{>0}(\varphi)$ with $v \not\leq \bar{v}$ implies $\bar{v} \in V_{>0}(\varphi)$. The proof proceeds in a number of steps. (1) We show that it suffices to consider formulas without disjunction. (2) We show that if path fragment $\pi[0..l] \models \bar{\varphi}$, (where LTL(F,X)-formula $\bar{\varphi}$ is obtained from φ by omitting all parameters from φ) then $\pi[0..l] \models v_l(\varphi)$ with $v_l(x) = l$ for every x . (3) We construct a deterministic Büchi automaton (DBA) $A_{\bar{\varphi}}$ for $\bar{\varphi}$ such that its initial and final state are at most $|\bar{\varphi}|$ transitions apart. (4) We show that reachability of a final state in the product of MC M and DBA $A_{\bar{\varphi}}$ implies the existence of a finite path in M of length at most $m \cdot |\varphi|$ satisfying $\bar{\varphi}$.

- As disjunction distributes over $\wedge, \bigcirc, \Diamond$, and \Diamond_x , each formula can be written in disjunctive normal form. Let $\varphi \equiv \varphi_1 \vee \dots \vee \varphi_k$, where each φ_i is disjunction-free. Evidently, $|\varphi_i| \leq |\varphi|$. Assume $v \in V_{>0}(\varphi)$. Then, $v \in V_{>0}(\varphi_i)$ for some $0 < i \leq k$. Assuming the theorem holds for φ_i (this will be proven below), $\bar{v}_i \in V_{>0}(\varphi_i)$ with $\bar{v}_i(x) = |\varphi_i| \cdot m$. Since $\bar{v} \geq \bar{v}_i$, it follows by monotonicity that $\bar{v} \in V_{>0}(\varphi_i)$, and hence, $\bar{v} \in V_{>0}(\varphi)$. It thus suffices in the remainder of the proof to consider disjunction-free formulas.
- For pLTL(F,X)-formula φ , let $\bar{\varphi}$ be the LTL(F,X)-formula obtained from φ by replacing all occurrences of \Diamond_x by \Diamond , e.g., for $\varphi = \Diamond_x(a \wedge \Diamond_y b)$, $\bar{\varphi} = \Diamond(a \wedge \Diamond b)$. We claim that $\pi[0..l] \models \bar{\varphi}$ implies $\pi[0..l] \models v_l(\varphi)$ with $v_l(x) = l$ for all x . This is proven by induction on the structure of φ . The base cases

- a and $\neg a$ are obvious. For the induction step, conjunctions, $\bigcirc \varphi$ and $\Diamond \varphi$ are straightforward. It remains to consider $\Diamond_x \varphi$. Assume $\pi[0..l] \models \Diamond \bar{\varphi}$. Thus, for some $i \leq l$, $\pi[i..l] \models \bar{\varphi}$. By induction hypothesis, $\pi[i..l] \models \mathbf{v}_i(\varphi)$ with $\mathbf{v}_i(y) = l-i$ for each variable y in φ . Thus, $\pi[0..l] \models \mathbf{v}_l(\Diamond_x \varphi)$ with $\mathbf{v}_l(x) = l$ and for all y in φ , $\mathbf{v}_l(y) = l$.
3. We provide a DBA $A_{\bar{\varphi}} = \langle Q, \Sigma, \delta, q_0, F \rangle$ with $\Sigma = 2^{AP}$ for each LTL(F,X)-formula $\bar{\varphi}$ using the construction from [15]. We first treat $\bar{\varphi} = a$ and $\bar{\varphi} = \Diamond a$. As every LTL(F,X)-formula can be obtained from $\Diamond(a \wedge \varphi)$, $\varphi_1 \wedge \varphi_2$ and $\bigcirc \varphi$, we then treat these inductive cases. (Negations are treated similarly.) For $\bar{\varphi} = a$, $A_a = \langle \{q_0, q_1\}, \Sigma, \delta, q_0, \{q_1\} \rangle$ with $\delta(q_0, a) = q_1$ and $\delta(q_1, \text{true}) = q_1$. For $\bar{\varphi} = \Diamond a$, the DBA $A_{\Diamond a} = \langle \{q_0, q_1\}, \Sigma, \delta, q_0, \{q_1\} \rangle$, where $\delta(q_0, a) = q_1$, $\delta(q_0, \neg a) = q_0$ and $\delta(q_1, \text{true}) = q_1$. This completes the base cases. For the three inductive cases, the DBA is constructed as follows.
 - (a) Let $A_{\bar{\varphi}} = \langle Q, \Sigma, \delta, q_0, F \rangle$. $A_{\Diamond(a \wedge \bar{\varphi})} = \langle Q \cup \{q'_0\}, \Sigma, \delta', q'_0, F \rangle$ where q'_0 is fresh, $\delta'(q, \cdot) = \delta(q, \cdot)$ if $q \in Q$, $\delta'(q'_0, a) = \delta(q_0, a)$, and $\delta'(q'_0, \neg a) = q'_0$.
 - (b) For $\bar{\varphi}_1 \wedge \bar{\varphi}_2$, the DBA is a standard synchronous product of the DBA for $\bar{\varphi}_1$ and $\bar{\varphi}_2$.
 - (c) Let $A_{\bar{\varphi}} = \langle Q, \Sigma, \delta, q_0, F \rangle$. $A_{\bigcirc \bar{\varphi}} = \langle Q \cup \{q'_0\}, \Sigma, \delta', q'_0, F \rangle$ where q'_0 is fresh, $\delta'(q'_0, a) = q_0$ for all $a \in \Sigma$ and $\delta'(q, a) = \delta(q, a)$ for every $q \in Q$.
- A few remarks are in order. The resulting DBA have a single final state. In addition, the DBA enjoy the property that the reflexive and transitive closure of the transition relation is a partial order [15]. Formally, $q \preceq q'$ iff $q' \in \delta^*(q, w)$ for some $w \in \Sigma^\omega$. The diameter of $A_{\bar{\varphi}}$ is the length of a longest simple path from the initial to the final state. This implies that the diameter of $A_{\Diamond(a \wedge \bar{\varphi})}$ and $A_{\bigcirc \bar{\varphi}}$ is $n+1$ where n is this diameter of $A_{\bar{\varphi}}$, and the diameter of $A_{\bar{\varphi}_1 \wedge \bar{\varphi}_2}$ is $n_1 + n_2$ where n_i is the diameter of $A_{\bar{\varphi}_i}$, $i \in \{1, 2\}$.
4. Let $\varphi \equiv \varphi_1 \vee \dots \vee \varphi_k$, where each φ_i is disjunction-free, with DBA $A_{\bar{\varphi}_i}$. Evidently, $V_{>0}(\varphi) \neq \emptyset$ iff $V_{>0}(\varphi_i) \neq \emptyset$ for some disjunct φ_i . Consider the product of MC M and DBA $A_{\bar{\varphi}_i}$, denoted $M \otimes A_{\bar{\varphi}_i}$; see, e.g., [16, Def. 10.50]. By construction, $M \otimes A_{\bar{\varphi}_i}$ is partially ordered and has diameter at most $m \cdot |\varphi_i|$. We have that $\Pr(M \models \bar{\varphi}_i) > 0$ iff an accepting state in $M \otimes A_{\bar{\varphi}_i}$ is reachable. Thus, there exists a finite path $\pi[0..m \cdot |\varphi_i|]$ in M with $\pi[0..m \cdot |\varphi_i|] \models \bar{\varphi}$, or, $\pi[0..m \cdot |\varphi|] \models \bar{v}(\varphi)$. This concludes the proof.

$M \otimes A_{\bar{\varphi}_i}$ can also be used to show that, if we have a valuation \mathbf{v} such that $\mathbf{v}(x) > m \cdot |\varphi|$ and for all other variables $y \neq x$, $\mathbf{v}(y) \leq m \cdot |\varphi|$ and $\mathbf{v} \in V_{>0}(\varphi)$ then $\mathbf{v}' \in V_{>0}(\varphi)$, where $\mathbf{v}'(x) = m \cdot |\varphi|$ and for $y \neq x$, $\mathbf{v}'(y) = \mathbf{v}(y)$. The argument proceeds as induction on $\bar{\varphi}_i$. ■

The above Theorem 1 leads to the following proposition.

Proposition 5. *For $\varphi \in pLTL(F,X)$, deciding if $V_{>0}(\varphi) = \emptyset$ is NP-complete.*

For almost sure properties, a similar approach as for $V_{>0}(\varphi)$ suffices.

Theorem 2. *For $\varphi \in pLTL(F,X)$ and MC M with m states, $V_{=1}(\varphi) \neq \emptyset$ iff $\bar{v} \in V_{=1}(\bar{\varphi})$ with $\bar{v}(x) = m \cdot |\varphi|$.*

Proof. Consider the direction from left to right. The argument goes along similar lines as the proof of Theorem 1. We build the DBA $A_{\bar{\varphi}}$ for $\bar{\varphi}$ and take the cross product with Markov chain M . There are $m \cdot |\varphi|$ state in the cross product. If $\Pr(M \models \bar{v}(\varphi)) < 1$ then there is some cycle in the cross product that does not contain the final state. Thus, $V_{=1}(\varphi)$ is empty. ■

Theorem 1 suggests that $\min V_{>0}(\varphi)$ lies in the hyper-cube $H = \{0, \dots, N\}^d$, where $N = m \cdot |\varphi|$. A possible way to find $\min V_{>0}(\varphi)$ is to apply the bisection method in d -dimensions. We recursively choose a middle point of the cube, say $\mathbf{v} \in H$ —in the first iteration $\mathbf{v}(x) = N/2$ — and divide H in 2^d equally sized hypercubes. If $\mathbf{v} \in V_{>0}(\varphi)$, then the hypercube whose points exceed \mathbf{v} is discarded, else the cube whose points are below \mathbf{v} is discarded. The asymptotic time-complexity of this procedure is given by the recurrence relation:

$$T(k) = (2^d - 1) \cdot T(k \cdot 2^{-d}) + F \quad (3)$$

where k is the number of points in the hypercube and F is the complexity of checking $\mathbf{v} \in V_{>0}(\varphi)$ where $|\mathbf{v}| \leq N$. Section 6 presents an algorithm working in $\mathcal{O}(m \cdot N^d \cdot 2^{|\varphi|})$ for a somewhat more expressive logic. From (3), this yields a complexity of $\mathcal{O}(m \cdot N^d \cdot 2^{|\varphi|} \cdot \log N)$. The size of a set of minimal points can be exponential in the number of variables, as shown below.

Proposition 6. $|\min V_{>0}(\varphi)| \leq (N \cdot d)^{d-1}$.




Fig. 1. MC and $\min V_{>0}(\varphi)$ for pLTL(F,X)-formula $\varphi = \Diamond_{x_1} r \wedge \Diamond_{x_2} b \wedge \Diamond_{x_3} g$

Example 2. There exist MCs for which $|\min V_{>0}(\varphi)|$ grows exponentially in d , the number of parameters in φ , whereas the number m of states in the MC grows linearly in d . For instance, consider the MC M in Fig. 1 and $\varphi = \Diamond_{x_1} r \wedge \Diamond_{x_2} b \wedge \Diamond_{x_3} g$, i.e., $d=3$. We have $|\min V_{>0}(\varphi)| = 4^2$ as indicated in the table.

We conclude this section by briefly considering the membership query: does $\mathbf{v} \in V_{>0}(\varphi)$ for pLTL(F,X)-formula φ with d parameters? Checking membership

of a valuation $\mathbf{v} \in V_{>0}(\varphi)$ boils down to deciding whether there exists a $\mathbf{v}' \in \min V_{>0}(\varphi)$ such that $\mathbf{v} \geq \mathbf{v}'$. A representation of $\min V_{>0}(\varphi)$ facilitating an efficient membership test can be obtained by putting all elements in this set in lexicographical order. This involves sorting over all d coordinates. A membership query then amounts to a recursive binary search over d dimensions. This yields:

Proposition 7. *For pLTL(F, X)-formula φ , $\mathbf{v} \in V_{>0}(\varphi)$? takes $\mathcal{O}(d \cdot \log N \cdot d)$ time, provided a representation of $\min V_{>0}(\varphi)$ is given.*

5 Qualitative Parametric Büchi

In this section, we consider pLTL-formulas of the form $\varphi = \square \Diamond_x a$, for proposition a . We are interested in $V_{>0}(\varphi)$, i.e., does the set of infinite paths visiting a -states that are maximally x apart infinitely often, have a positive measure? Let MC $M = (S, \mathbf{P}, s_0, L)$. A *bottom strongly-connected component* (BSCC) $B \subseteq S$ of M is a set of mutually reachable states with no edge leaving B . For BSCC B , let $n_{a,B} = \max\{|\pi| \mid \forall i \leq |\pi|, \pi[i] \in B \wedge a \notin L(\pi[i])\}$.

Proposition 8. *Let B be a BSCC and $s \in B$. Then, $\forall n \in \mathbb{N}, n > n_{a,B} \Leftrightarrow \Pr(s \models \square \Diamond_n a) = 1$ and $n \leq n_{a,B} \Leftrightarrow \Pr(s \models \square \Diamond_n a) = 0$.*

Proof. If $n > n_{a,B}$, then each path π from any state $s \in B$ will have at least one a -state in finite path fragment $\pi[i, \dots, i+n]$ for all i . Hence, $\Pr(s \models \square \Diamond_n a) = 1$. If $n \leq n_{a,B}$, then there exists a finite path fragment ρ of B , such that, for all $i \leq n$, $a \notin L(\rho[i])$. Consider an infinite path π starting from any arbitrary $s \in B$. As $s \in B$, π will almost surely infinitely often visit the initial state of ρ . Therefore, by [16, Th.10.25], π will almost surely visit every finite path fragment starting in that state, in particular ρ . Path π thus almost surely refutes $\square \Diamond_n a$, i.e. $\Pr(s \models \square \Diamond_n a) = 0$. ■

For any BSCC B and $\square \Diamond_x a$, $n_{a,B} < \infty$ iff every cycle in B has at least one a -state. Hence, $n_{a,B}$ can be obtained by analysing the digraph of B (in $\mathcal{O}(m^2)$, the number of edges). BSCC B is called *accepting* for $\square \Diamond_x a$ if $n_{a,B} < \infty$ and B is reachable from the initial state s_0 . Note that this may differ from being an accepting BSCC for $\square \Diamond a$. Evidently, $V_{>0}(\square \Diamond_x a) \neq \emptyset$ iff $n_{a,B} < \infty$. This result can be extended to *generalized Büchi* formula $\varphi = \square \Diamond_{x_1} a_1 \wedge \dots \wedge \square \Diamond_{x_d} a_d$, by checking $n_{a_i,B} < \infty$ for each a_i .

As a next problem, we determine $\min V_{>0}(\square \Diamond_x a)$. For the sake of simplicity, let MS M have a single accepting BSCC B . For states s and t in MC M , let $d(s, t)$ be the distance from s to t in the graph of M . (Recall, the distance between state s and t is the length of the shortest path from s to t .) For BSCC B , let $d_{a,B}(s) = \min_{t \in B, a \in L(t)} d(s, t)$, i.e., the minimal distance from s to an a -state in B . Let the proposition a_B hold in state s iff $s \in B$ and $a \in L(s)$. Let $G_a = (V, E)$ be the digraph defined as follows: V contains all a -states of M and the initial state s_0 and $(s, s') \in E$ iff there is path from s to s' in M . Let c be a cost function defined on a finite path $s_0 \dots s_n$ in graph G_a as: $c(s_0 \dots s_n) = \max_i d(s_i, s_{i+1})$, (d is defined on the graph of M). Using these auxiliary notions we obtain the following characterization for $\min V_{>0}(\square \Diamond_x a)$:

Theorem 3. $\min V_{>0}(\Box\Diamond_x a) = n_0$ where $n_0 = \max \left(n_{a,B}, \min_{\pi=s_0 \dots s_n, s_n \models a_B} c(\pi) \right)$ if $n_{a,B} < d_{a,B}(s_0)$ and $n_0 = n_{a,B}$ otherwise.

Proof. We show for $n \geq n_0$, $\Pr(\Box\Diamond_n a) > 0$, and for $n < n_0$, $\Pr(\Box\Diamond_n a) = 0$. Distinguish:

1. $n_{a,B} \geq d_{a,B}(s_0)$. Then, from s_0 an a -state in B can be reached within $n_{a,B}$ steps, i.e., $\Pr(s_0 \models \Diamond_{n_{a,B}} a_B) > 0$. For this a_B -state, s , say, by Proposition 8 it follows $\Pr(s \models \Box\Diamond_{n_{a,B}} a) = 1$. Together this yields $\Pr(s_0 \models \Box\Diamond_n a) > 0$ for each $n \geq n_{a,B} = n_0$. For $n < n_0 = n_{a,B}$, it follows by Proposition 8 that $\Pr(s \models \Box\Diamond_n a) = 0$ for every a_B -state s . Thus, $\Pr(s_0 \models \Box\Diamond_n a) = 0$.
2. $n_{a,B} < d_{a,B}(s_0)$. As B is accepting, $d_{a,B}(s_0) \neq \infty$. Consider a simple path π from s_0 to an a -state in B . Let $c(\pi)$ be the maximal distance between two consecutive a -states along this path. Then it follows $\Pr(s_0 \models \Box\Diamond_k a) > 0$ where $k = \max(c(\pi), n_{a,B})$. By taking the minimum c_{\min} over all simple paths between s_0 and B , it follows $\Pr(s_0 \models \Box\Diamond_n a) > 0$ for each $n \geq n_0 = \max(n_{a,B}, c_{\min})$ with $c_{\min} = \min_{\pi \in \text{Paths}(s_0, \Diamond a_B)} c(\pi)$. For $n < n_0$, distinguish between $n_0 = n_{a,B}$ and $n_0 = c_{\min}$. In the former case, it follows (as in the first case) by Proposition 8 that $\Pr(s_0 \models \Box\Diamond_n a) = 0$ for all $n \geq n_0$. Consider now $n_0 = c_{\min} \geq n_{a,B}$. Let $n < n_0$. By contra-position. Assume $\Pr(s_0 \models \Box\Diamond_n a) > 0$. Let $\pi = s_0 \dots s_{1,a} \dots s_{2,a} \dots \dots s_{k,a}$ be a finite path fragment in M where $s_{i,a} \models a$ and $s_{k,a}$ is the first a -state along π which belongs to B . Then, by definition of the digraph G_a , the sequence $\pi = s_0 s_{1,a} s_{2,a} \dots s_{k,a}$ is a path in G_a satisfying $c(s_{i,a}, s_{i+1,a}) \leq n$ for all $0 \leq i < k$. But then $c_{\min} \leq n$. Contradiction. ■

If MC M has more than one accepting BSCC, say $\{B_1, \dots, B_k\}$ with $k > 1$, then $n_0 = \min_i n_{0,B_i}$, where n_{0,B_i} for $0 < i \leq k$ is obtained as in Theorem 3.

Proposition 9. The sets $V_{>0}(\Box\Diamond_x a)$ and $V_{=1}(\Box\Diamond_x a)$ can be determined in polynomial time by a graph analysis of MC M .

Determining $\min V_{\geq p}(\Box\Diamond_x a)$ for arbitrary p reduces to reachability of *accepting* BSCCs. In a similar way as for parametric reachability (cf. Section 3), this can be done searching. For generalized Büchi formula $\varphi = \Box\Diamond_{x_1} a_1 \wedge \dots \wedge \Box\Diamond_{x_d} a_d$ and BSCC B , $n_{a,B}$ is at most m . Thus, $\min V_{>0}(\varphi) \in \{0, \dots, m \cdot d\}^d$ and can be found by the bisection method, similar to the procedure described in Section 4.

6 The Fragment pLTL $_{\Diamond}$

This section is concerned with the logical fragment pLTL $_{\Diamond}$, as defined in (2):

$$\varphi ::= a \mid \neg a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \bigcirc \varphi \mid \varphi \mathsf{U} \varphi \mid \varphi \mathsf{R} \varphi \mid \Box \varphi \mid \Diamond_{\leq x} \varphi.^1$$

We will focus on the emptiness problem: is $V_{>0}(\varphi) = \emptyset$. The decision problem whether $V_{=1}(\varphi)$ is very similar. Similar as for pLTL(F,X), we obtain necessary and sufficient criteria for both cases. The proofs for these criteria depend on an algorithm that checks whether $v \in V_{>0}(\varphi)$. This algorithm is presented first.

¹ The modalities $\Diamond_{\leq c}$ and $\Box_{\leq c}$ can be removed with only quadratic blow up.

Automata constructions. Let φ be a pLTL $_{\Diamond}$ -formula, and \mathbf{v} a variable valuation. W.l.o.g. we assume that each variable occurs once in φ . We will extend the classical automaton-based approach for LTL by constructing a nondeterministic Büchi automaton for φ that is amenable to treat the variables occurring in φ . To that end, inspired by [17], we proceed in a number of steps:

1. Construct an automaton G_φ for φ , independent from the valuation \mathbf{v} , with two types of acceptance sets, one for treating until and release-modalities (as standard for LTL [18]), and one for treating the parameter constraints.
2. Establish how for a given valuation \mathbf{v} , a Büchi automaton $B_\varphi(\mathbf{v})$ can be obtained from G_φ such that for infinite word w , $(w, \mathbf{v}) \in \mathcal{L}(\varphi)$ iff w is an accepting run of $B_\varphi(\mathbf{v})$.
3. Exploit the technique advocated by Couvreur *et al.* [3] to verify MC M versus $B_\varphi(\mathbf{v})$.

We start with constructing G_φ . Like for the LTL-approach, the first step is to consider consistent sets of sub-formulas of φ . Let $cl(\varphi)$ be the set of all sub-formulas of φ . Set $H \subseteq cl(\varphi)$ is *consistent*, when:

- $a \in H$ iff $\neg a \notin H$,
- $\varphi_1 \wedge \varphi_2 \in H$ iff $\varphi_1 \in H$ and $\varphi_2 \in H$,
- $\varphi_1 \vee \varphi_2 \in H$ iff $\varphi_1 \in H$ or $\varphi_2 \in H$,
- $\varphi_2 \in H$ implies $\varphi_1 \cup \varphi_2 \in H$,
- $\varphi_1, \varphi_2 \in H$ implies $\varphi_1 R \varphi_2 \in H$,
- $\varphi_1 \in H$ implies $\Diamond_x \varphi_1 \in H$.

We are now in a position to define G_φ , an automaton with two acceptance sets. For $\varphi \in \text{pLTL}_\Diamond$, let $G_\varphi = (Q, 2^{AP}, Q_0, \delta, Acc_B, Acc_P)$ where

- Q is the set of all consistent sub-sets of $cl(\varphi)$ and $Q_0 = \{ H \in Q \mid \varphi \in H \}$.
- $(H, a, H') \in \delta$, where $a \in 2^{AP}$ whenever:
 - $H \cap AP = \{ a \}$,
 - $\bigcirc \varphi_1 \in H \iff \varphi_1 \in H'$,
 - $\varphi_1 \cup \varphi_2 \in H \iff \varphi_2 \in H$ or $(\varphi_1 \in H \text{ and } \varphi_1 \cup \varphi_2 \in H')$,
 - $\varphi_1 R \varphi_2 \in H \iff \varphi_2 \in H$ and $(\varphi_1 \in H \text{ or } \varphi_1 R \varphi_2 \in H')$,
 - $\Diamond_x \varphi_1 \in H \iff \varphi_1 \in H$ or $\Diamond_x \varphi_1 \in H'$,
- (generalized) Büchi acceptance Acc_B and parametric acceptance Acc_P :
 - $Acc_B = \{ F_{\varphi'} \mid \varphi' \in cl(\varphi) \wedge (\varphi' = \varphi_1 \cup \varphi_2 \vee \varphi' = \varphi_1 R \varphi_2) \}$ where
 - * $F_{\varphi'} = \{ H \mid \varphi' \in H \Rightarrow \varphi_2 \in H \}$ if $\varphi' = \varphi_1 \cup \varphi_2$, and
 - * $F_{\varphi'} = \{ H \mid \varphi_2 \in H \Rightarrow \varphi' \in H \}$ if $\varphi' = \varphi_1 R \varphi_2$,
 - $Acc_P = \{ F_{x_i} \mid \Diamond_{x_i} \varphi_i \in cl(\varphi) \}$ with $F_{x_i} = \{ H \mid \Diamond_{x_i} \varphi_i \in H \Rightarrow \varphi_i \in H \}$.

A run $\rho \in Q^\omega$ of G_φ is *accepting* under valuation \mathbf{v} if it visits each set in Acc_B infinitely often and each $F_{x_i} \in Acc_P$ in every infix of length $\mathbf{v}(x_i)$. $\mathcal{L}(G_\varphi)$ contains all pairs (w, \mathbf{v}) such that there is an accepting run of w under the valuation \mathbf{v} . G_φ is *unambiguous* if $q \xrightarrow{a} q'$ and $q \xrightarrow{a} q''$ implies $\mathcal{L}(q') \cap \mathcal{L}(q'') = \emptyset$, where $\mathcal{L}(q)$ is the language starting from the state q .

Proposition 10 ([17]). *For $\varphi \in \text{pLTL}_\Diamond$, the automaton G_φ is unambiguous and $\mathcal{L}(G_\varphi) = \mathcal{L}(\varphi)$.*

The automaton G_φ can be constructed in $\mathcal{O}(2^{|\varphi|})$. Apart from the parametric acceptance condition, G_φ behaves as a generalized Büchi automaton (GNBA) with accepting set $Acc_B = \{F_1, \dots, F_k\}$. In order to obtain a non-deterministic automaton, we first apply a similar transformation as for GNBA to NBA [16]. We convert G_φ to $U_\varphi = (Q', 2^{AP}, Q'_0, \delta', Acc'_B, Acc'_P)$ where $Q' = Q \times \{1, \dots, k\}$, $Q'_0 = Q_0 \times \{1\}$. If $(q, a, q') \in \delta$, then $((q, i), a, (q', i')) \in \delta'$ with $i=i'$ if $q \notin F_i$ else $i' = (i \bmod k) + 1$. $Acc_B = F_1 \times \{1\}$ and $Acc'_P = \{F'_{x_i} \mid F_{x_i} \in Acc_P\}$, where $F'_{x_i} = F_{x_i} \times \{1, \dots, k\}$. Note that the construction preserves unambiguity and the size of U_φ is in $\mathcal{O}(|\varphi| \cdot 2^{|\varphi|})$.

For a given valuation \mathbf{v} , U_φ can be converted into an NBA $B_\varphi(\mathbf{v})$. This is done as follows. Let $U_\varphi = (Q', 2^{AP}, Q'_0, \delta', Acc'_B, Acc'_P)$ and \mathbf{v} a valuation of φ with d parameters. Then $B_\varphi(\mathbf{v}) = (Q'', 2^{AP}, Q''_0, \delta'', Acc)$ with:

- $Q'' \subseteq Q' \times \{0, \dots, \mathbf{v}(x_1)\} \times \dots \times \{0, \dots, \mathbf{v}(x_d)\}$,
- $((q, \mathbf{n}), a, (q', \mathbf{n}')) \in \delta''$ if $(q, a, q') \in \delta'$ and for all x_i :
 - if $q' \in F'_{x_i}$ and $\mathbf{n}(x_i) < \mathbf{v}(x_i)$ then $\mathbf{n}'(x_i) = 0$,
 - if $q' \notin F'_{x_i}$ and $\mathbf{n}(x_i) < \mathbf{v}(x_i)$ then $\mathbf{n}'(x_i) = \mathbf{n}(x_i) + 1$.
- $Q''_0 = Q'_0 \times 0^d$ and $Acc = Acc'_B \times \{0, \dots, \mathbf{v}(x_1)\} \times \dots \times \{0, \dots, \mathbf{v}(x_d)\}$.

It follows that $B_\varphi(\mathbf{v})$ is unambiguous for any valuation \mathbf{v} . Furthermore, every run of $B_\varphi(\mathbf{v})$ is either finite or satisfies the parametric acceptance condition for valuation \mathbf{v} . Thus we have:

Proposition 11. *An infinite word $w \in \mathcal{L}(B_\varphi(\mathbf{v}))$ if and only if $(w, \mathbf{v}) \in \mathcal{L}(\varphi)$.*

The size of $B_\varphi(\mathbf{v})$ is in $\mathcal{O}(c_v \cdot |\varphi| \cdot 2^{|\varphi|})$ where $c_v = \prod_{x_i} (\mathbf{v}(x_i) + 1)$.

As a next step, we exploit the fact that $B_\varphi(\mathbf{v})$ is unambiguous, and apply the technique by Couvreur *et al.* [3] for verifying MC M against $B_\varphi(\mathbf{v})$. Let $M \otimes B_\varphi(\mathbf{v})$ be the synchronous product of M and $B_\varphi(\mathbf{v})$ [16], Π_1 the projection to M and Π_2 the projection to $B_\varphi(\mathbf{v})$. Let $\mathcal{L}(s, q) = \{\pi \in Paths(s) \mid trace(\pi) \in \mathcal{L}(q)\}$ and $Pr(s, q) = Pr(\mathcal{L}(s, q))$. Let $Pr(M \otimes B_\varphi(\mathbf{v})) = \sum_{q_0 \in Q_0} Pr(s_0, q_0)$. As $B_\varphi(\mathbf{v})$ is unambiguous, we have for any (s, q) :

$$Pr(s, q) = \sum_{(t, q') \in \delta(s, q)} \mathbf{P}(s, t) \cdot Pr(t, q'),$$

where δ is the transition relation of $M \otimes B_\varphi(\mathbf{v})$ and $\mathbf{P}(s, t)$ is the one-step transition probability from s to t in MC M . A (maximal) strongly connected component (SCC, for short) $C \subseteq S$ is *complete* if for any $s \in \Pi_1(C)$:

$$Paths(s) = \bigcup_{(s, q) \in C} \mathcal{L}_C(s, q)$$

where $\mathcal{L}_C(s, q)$ restricts runs to C (runs only visits states from C). The SCC C is *accepting* if $Acc \cap \Pi_2(C) \neq \emptyset$ (where Acc is the set of accepting states in $B_\varphi(\mathbf{v})$).

Proposition 12 ([3]). *Let C be a complete and accepting SCC in $M \otimes B_\varphi(\mathbf{v})$. Then for all $s \in \Pi_1(C)$:*

$$\Pr\left(\bigcup_{(s,q) \in C} \mathcal{L}_C(s, q)\right) = 1.$$

Moreover, since $B_\varphi(\mathbf{v})$ is unambiguous, $\Pr(M \otimes B_\varphi(\mathbf{v})) > 0$ implies there exists a reachable, complete and accepting SCC.

Finding complete and accepting SCC in $M \otimes B_\varphi(\mathbf{v})$ is done by standard graph analysis. Altogether, $\mathbf{v} \in V_{>0}(\varphi)$ is decided in $\mathcal{O}(m \cdot c_v \cdot |\varphi| \cdot 2^{|\varphi|})$. The space complexity is polynomial in the size of the input (including the valuation), as $M \otimes B_\varphi(\mathbf{v})$ can be stored in $\mathcal{O}(\log m + |\varphi| + \log c_v)$ bits. In the sequel, we exploit these results to obtain a necessary and sufficient criterion for the emptiness of $V_{>0}(\varphi)$ for φ in $p\text{LTL}_\diamond$.

Theorem 4. *For $\varphi \in p\text{LTL}_\diamond$, $V_{>0}(\varphi) \neq \emptyset$ iff $\bar{\mathbf{v}} \in V_{>0}(\varphi)$ s.t. $\bar{\mathbf{v}}(x) = m \cdot |\varphi| \cdot 2^{|\varphi|}$.*

Proof. Consider the direction from left to right. The only non-trivial case is when there exists a valuation $\mathbf{v} \not\leq \bar{\mathbf{v}}$ such that $\mathbf{v} \in V_{>0}(\varphi)$ implies $\bar{\mathbf{v}} \in V_{>0}(\varphi)$. In the model checking algorithm described above, we first construct G_φ , and then U_φ with a single Büchi accepting set Acc'_B and d parametric accepting sets F'_{x_i} , one for each variable x_i in φ . For the sake of clarity, assume $d = 1$, i.e., we consider valuation v . The explanation extends to the general case in a straightforward manner. For valuation v , consider $M \otimes B_\varphi(v)$. We show that, for $r < v$, $\Pr(M \otimes B_\varphi(v)) > 0$ implies $\Pr(M \otimes B_\varphi(r)) > 0$, where $r = m \cdot |U_\varphi|$, which is in $\mathcal{O}(m \cdot |\varphi| \cdot 2^{|\varphi|})$.

Note that every cycle in $M \otimes B_\varphi(r)$ contains a state (s, q, i) with $i = 0$. Moreover, the graph of $M \otimes B_\varphi(r)$ is a sub-graph of $M \otimes B_\varphi(v)$. We now prove that, if a (maximal) SCC C of $M \otimes B_\varphi(r)$ is not complete (or accepting) then any SCC C' of $M \otimes B_\varphi(v)$ containing C is also not complete (or accepting, respectively).

(a) Suppose C is not complete. Then there exists a finite path $\sigma = s s_1 \dots s_k$ of M , such that from any q , with $(s, q, 0) \in C$, the run $\rho = (s, q, 0)(s_1, q_1, 1) \dots (s_j, q_j, j)$ leads to a deadlock state. This can have two causes: either (s_j, q_j, j) has no successor for any j . Then, C' is not complete. Or, the path ρ terminates in (s_j, q_j, j) where $j = r$. This means, for all $(s', q', j+1) \in \delta(s_j, q_j, j)$ in C' , $q' \notin F_x$. As the length of ρ exceeds r , there are states in the run whose first and second component appear multiple times. Thus, we can find another path σ' (possibly longer than σ) for C' which goes through states where the first and the second component of some of its states are repeated sufficiently many times to have a run $(s, q, 0)(s_1, q_1, 1) \dots (s_j, q_j, v)$ which is a deadlock state. Thus, C' is not complete.

(b) Suppose C' is accepting. Then there exists (s', q', i') with $q' \in \text{Acc}$. Since C' is an SCC and $C \subseteq C'$, there is a path from $(s, q, 0) \in C$ to (s', q', i') . If the length of the path is less than r , then we are done. If $i' > r$, then some (s'', q'') pair in the path must be repeated. Thus, we can find another path of length less

than r to a state (s', q', i) , where $i \leq r$. Therefore, C is accepting. The rest of the proof follows from Proposition 12. ■

For almost sure properties, a similar approach as for $V_{>0}(\varphi)$ suffices.

Theorem 5. *For $\varphi \in pLTL_\diamond$, $V_{=1}(\varphi) \neq \emptyset$ iff $\bar{v} \in V_{=1}(\bar{\varphi})$ with $\bar{v}(x) = m \cdot |\varphi| \cdot 2^{|\varphi|}$.*

Let $N_{\varphi M} = m \cdot |\varphi| \cdot 2^{|\varphi|}$. Note that $c_{\bar{v}}$ equals $(N_{\varphi M})^d$. Thus, we have:

Proposition 13. *For $\varphi \in pLTL_\diamond$, deciding if $V_{>0}(\varphi) = \emptyset$ is PSPACE-complete.*

Proof. Theorem 4 gives an algorithm in PSPACE, as $M \otimes B_\varphi(\bar{v})$ can be stored in $O(\log m + |\varphi| + d \log N_{\varphi M})$ bits. PSPACE hardness follows trivially, as for LTL formula φ and MC M , deciding $\Pr(M \models \varphi) > 0$ (which is known to be a PSPACE complete problem) is the same as checking the emptiness of $V_{>0}(\varphi)$. ■

Just as for pLTL(F,X), we can use the bisection method to find $\min V_{>0}(\varphi)$. The search procedure invokes the model checking algorithm multiple times. We can reuse the space each time we check $\Pr(M \models v(\varphi)) > 0$. Hence, $\min V_{>0}(\varphi)$ can be found in polynomial space. The time complexity of finding $\min V_{>0}(\varphi)$ is $\mathcal{O}(m \cdot (N_{\varphi M})^d \cdot 2^{|\varphi|} \cdot \log N_{\varphi M})$. Membership can also be similarly solved.

Proposition 14. *For $pLTL_\diamond$ -formula φ , $v \in V_{>0}(\varphi)$? takes $\mathcal{O}(d \cdot \log \frac{N_{\varphi M}}{d})$ time, provided a representation of $V_{>0}(\varphi)$ is given.*

7 Concluding Remarks

This paper considered the verification of finite MCs against parametric LTL. We obtained several results on the emptiness problem for qualitative verification problems, including necessary and sufficient conditions as well as some complexity results. Future work consists of devising more efficient algorithms for the quantitative verification problems, and lifting the results to extended temporal logics [19] and stochastic games, possibly exploiting [17].

Acknowledgement. This work was partially supported by the EU FP7 projects MoVeS and Sensation, the EU Marie Curie project MEALS and the Excellence initiative of the German federal government.

References

1. Vardi, M.Y.: Automatic verification of probabilistic concurrent finite-state programs. In: FOCS, pp. 327–338. IEEE Computer Society (1985)
2. Courcoubetis, C., Yannakakis, M.: The complexity of probabilistic verification. J. ACM 42(4), 857–907 (1995)
3. Couvreur, J.M., Saheb, N., Sutre, G.: An optimal automata approach to LTL model checking of probabilistic systems. In: Vardi, M.Y., Voronkov, A. (eds.) LPAR 2003. LNCS, vol. 2850, pp. 361–375. Springer, Heidelberg (2003)

4. Chatterjee, K., Gaiser, A., Kretínský, J.: Automata with generalized Rabin pairs for probabilistic model checking and LTL synthesis. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 559–575. Springer, Heidelberg (2013)
5. Alur, R., Etessami, K., La Torre, S., Peled, D.: Parametric temporal logic for “model measuring”. ACM Trans. Comput. Log. 2(3), 388–407 (2001)
6. Han, T., Katoen, J.P., Mereacre, A.: Approximate parameter synthesis for probabilistic time-bounded reachability. In: IEEE Real-Time Systems Symposium (RTSS), pp. 173–182. IEEE Computer Society (2008)
7. Hahn, E.M., Han, T., Zhang, L.: Synthesis for PCTL in parametric markov decision processes. In: Bobaru, M., Havelund, K., Holzmann, G.J., Joshi, R. (eds.) NFM 2011. LNCS, vol. 6617, pp. 146–161. Springer, Heidelberg (2011)
8. Puggelli, A., Li, W., Sangiovanni-Vincentelli, A.L., Seshia, S.A.: Polynomial-time verification of PCTL properties of mDPs with convex uncertainties. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 527–542. Springer, Heidelberg (2013)
9. Benedikt, M., Lenhardt, R., Worrell, J.: LTL model checking of interval markov chains. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013. LNCS, vol. 7795, pp. 32–46. Springer, Heidelberg (2013)
10. Hahn, E.M., Hermanns, H., Zhang, L.: Probabilistic reachability for parametric Markov models. STTT 13(1), 3–19 (2011)
11. Daws, C.: Symbolic and parametric model checking of discrete-time markov chains. In: Liu, Z., Araki, K. (eds.) ICTAC 2004. LNCS, vol. 3407, pp. 280–294. Springer, Heidelberg (2005)
12. Katoen, J.P., McIver, A., Meinicke, L., Morgan, C.C.: Linear-invariant generation for probabilistic programs: In: Cousot, R., Martel, M. (eds.) SAS 2010. LNCS, vol. 6337, pp. 390–406. Springer, Heidelberg (2010)
13. Benedikt, M., Lenhardt, R., Worrell, J.: Two variable vs. linear temporal logic in model checking and games. Logical Methods in Computer Science 9(2) (2013)
14. Ummels, M., Baier, C.: Computing quantiles in Markov reward models. In: Pfennig, F. (ed.) FOSSACS 2013. LNCS, vol. 7794, pp. 353–368. Springer, Heidelberg (2013)
15. Alur, R., La Torre, S.: Deterministic generators and games for LTL fragments. ACM Trans. Comput. Log. 5(1), 1–25 (2004)
16. Baier, C., Katoen, J.P.: Principles of Model Checking. MIT Press (2008)
17. Zimmermann, M.: Optimal bounds in parametric LTL games. Theor. Comput. Sci. 493, 30–45 (2013)
18. Vardi, M.Y.: An automata-theoretic approach to linear temporal logic. In: Moller, F., Birtwistle, G. (eds.) Logics for Concurrency. LNCS, vol. 1043, pp. 238–266. Springer, Heidelberg (1996)
19. Vardi, M.Y., Wolper, P.: Reasoning about infinite computations. Information and Computation 115, 1–37 (1994)