

Parameter Synthesis for Parametric Probabilistic Dynamical Systems and Prefix-Independent Specifications

Christel Baier 

Technische Universität Dresden, Germany

Simon Jantsch 

Technische Universität Dresden, Germany

Engel Lefauchaux 

University of Lorraine, CNRS, Inria, LORIA, Nancy, France

David Purser 

University of Warsaw, Poland

James Worrell 

Department of Computer Science,
University of Oxford, UK

Florian Funke 

Technische Universität Dresden, Germany

Toghrul Karimov 

Max Planck Institute for Software Systems,
Saarland Informatics Campus, Germany

Joël Ouaknine 

Max Planck Institute for Software Systems,
Saarland Informatics Campus, Germany

Markus A. Whiteland 

University of Liège, Belgium

Abstract

We consider the model-checking problem for parametric probabilistic dynamical systems, formalised as Markov chains with parametric transition functions, analysed under the distribution-transformer semantics (in which a Markov chain induces a sequence of distributions over states).

We examine the problem of synthesising the set of parameter valuations of a parametric Markov chain such that the orbits of induced state distributions satisfy a prefix-independent ω -regular property.

Our main result establishes that in all non-degenerate instances, the feasible set of parameters is (up to a null set) semialgebraic, and can moreover be computed (in polynomial time assuming that the ambient dimension, corresponding to the number of states of the Markov chain, is fixed).

2012 ACM Subject Classification Theory of computation \rightarrow Logic and verification

Keywords and phrases Model checking, parametric Markov chains, distribution transformer semantics

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2022.29

Funding This work was funded by DFG grant 389792660 as part of TRR 248 – CPEC (see perspicuous-computing.science).

Joël Ouaknine: Also affiliated with Keble College, Oxford as emmy.network Fellow.

1 Introduction

The algorithmic analysis of Markov chains, in particular by means of model checking, is a central topic in probabilistic verification [7]. It is in fact fairly common to consider *parametric* Markov chains (PMCs), in which probabilities are given not as explicit numbers but rather as functions of certain parameters. One is then interested in the set of parameters giving rise to a Markov chain that meets a certain specification.

Markov chains are typically analysed under one of two standard semantics: the *path semantics* considers the set of all possible control-state trajectories, weighted by relevant probabilities, whereas the *distribution-transformer semantics* views the Markov chain as a single sequence of distributions over control states; this sequence is the orbit of the initial



© Christel Baier, Florian Funke, Simon Jantsch, Toghrul Karimov, Engel Lefauchaux, Joël Ouaknine, David Purser and Markus A. Whiteland;

licensed under Creative Commons License CC-BY 4.0

33rd International Conference on Concurrency Theory (CONCUR 2022).

Editors: Bartek Klin, Slawomir Lasota, and Anca Muscholl; Article No. 29; pp. 29:1–29:17

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

distribution under the repeated application of the underlying stochastic linear transformation. In this paper we focus exclusively on the second modelling paradigm, viewing Markov chains as special instances of linear dynamical systems (LDS). We consider parametric Markov chains, in which probabilities are given by rational functions over a set of parameters. Such parameters might account for uncertainties in the environment or in the exact values of the probabilities at hand, etc. Given a particular specification, we are interested in computing the set of all parameter valuations such that the resulting concrete Markov chain meets the specification. More precisely: *Given a parametric Markov chain \mathcal{M} over set of parameters X , i.e., a (parametric) matrix M and initial distribution π (see Definition 3), as well as a specification φ , compute the set of parameter instantiations $p \in \mathbb{R}^X$ for which the orbit $(\pi \cdot M[p]^n)_{n \geq 0}$ of \mathcal{M} satisfies φ .*

Our properties are specified with respect to the *characteristic word* of a Markov chain, which describes the orbit of the Markov chain relative to a set of targets. Given a Markov chain (Q, M, π) and a partition of the space $[0, 1]^Q = T_1 \cup \dots \cup T_k$, the characteristic word is the infinite word $w \in \{1, \dots, k\}^\omega$ such that $w_i = j$ if and only if $\pi \cdot M^i \in T_j$. For a parametric Markov chain, each admissible valuation of the parameters $p \in \mathbb{R}^X$ induces a concrete characteristic word $w[p]$. In this case we call $w : \mathbb{R}^X \rightarrow \{1, \dots, k\}^\omega$ the *parametric characteristic word* of the parametric Markov chain.

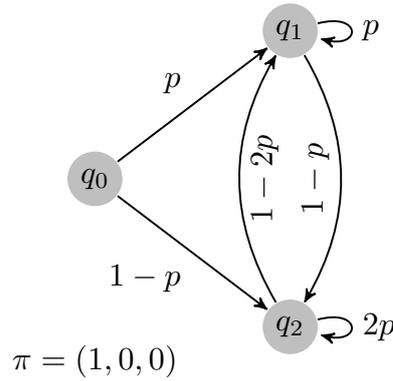
The model-checking problem asks, given a specification φ over $\{1, \dots, k\}$ (typically specified in LTL, MSO, or simply as an automaton), whether the characteristic word w satisfies φ , denoted by $w \models \varphi$. In the parametric setting, we are interested in the set of parameters $D_\varphi = \{p \in \mathbb{R}^X \mid w[p] \models \varphi\}$.

We consider ω -regular *prefix-independent* specifications, i.e., (intuitively speaking) properties that are invariant under finitely many changes to w (see Section 2.4 for the formal definition). The Ultimate Positivity Problem [22] is an example of a prefix-independent property: it asks whether the orbit is eventually trapped inside a certain target (chosen to be the region where a particular quantity is always positive). Other examples include repeatedly revisiting a target, since such a property only depends on any infinite suffix of w , regardless of the initial prefix. On the other hand, reachability is *not* a prefix-independent property. Note that LTL properties starting with “eventually always” or “always eventually” define prefix-independent properties (see [4]), although we do not limit ourselves to LTL properties.

Consider a parametric Markov chain with parametric characteristic word w and a prefix-independent specification φ . The set $D_\varphi = \{p \in \mathbb{R}^X \mid w[p] \models \varphi\}$ of feasible parameters can be a highly complex object. Nevertheless, one of our main results is that, assuming the specification is *non-degenerate* (a fairly mild technical condition), D_φ differs from a semialgebraic set by a null set (a set of Lebesgue measure zero), and moreover we can compute this semialgebraic set (in polynomial time assuming that the ambient dimension, corresponding to the number of states of the Markov chain, is fixed). More precisely, we show how to synthesise a semialgebraic set D' contained in the full set of feasible parameters such that $D_0 = D_\varphi \setminus D'$ is a null set.

Before going into the details of our construction, we show that the restriction to prefix-independent specifications is indeed necessary. Dropping it may lead to situations in which the set of feasible parameters is not semialgebraic, even up to a null set, as the following example shows.

► **Example 1 (A prefix-dependent property).** Consider the parametric Markov chain depicted in Figure 1a, with the single parameter p . Let us denote by (l, r, s) a distribution over Q , with l, r, s denoting the probability in states q_L, q_R, q_S respectively. Consider the following partition of $\text{Dist}(Q)$ (the set of probability distributions over Q) into the sets B, L, R, O ,



■ **Figure 2** Example Parametric Markov chain with parameter p .

In Appendix A we give a second example, which shows that reachability properties also do not have semialgebraic feasibility sets (up to a null set).

Let us consider another example, highlighting how we can use the limit distribution of an aperiodic Markov chain to decide prefix-independent properties.

► **Example 2 (Ultimate Positivity).** Consider the parametric Markov chain, depicted in Figure 2, with a single parameter p . The system represented by the diagram is a Markov chain for $p \in [0, 0.5]$ and has constant structure for all for $p \in (0, 0.5)$ (that is, each edge either exists for all p in the interval, or for none).

Consider the property that the probability distribution in states q_1 is eventually above 0.4 and q_2 is eventually above 0.55. We are interested in the set of parameters $D = \{p \in [0, 0.5] \mid \exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^n \geq (0, 0.4, 0.55)\}$.

The limit distribution of the Markov chain is $(0, \frac{1-2p}{2-3p}, \frac{1-p}{2-3p})$. Hence D , up to a null set, corresponds to the interval $(0, 0.5) \cap \{p \mid \frac{1-2p}{2-3p} > 0.4\} \cap \{p \mid \frac{1-p}{2-3p} > 0.55\} = (\frac{2}{13}, \frac{1}{4})$. Moreover all parameters in the interval $(\frac{2}{13}, \frac{1}{4})$ satisfy the property.

1.1 Related work

The papers [18, 19] introduce the logic iLTL to specify LTL-definable properties of the orbit of a Markov chain, where atomic propositions correspond to half-spaces. The authors devise a model-checking procedure which assumes that the Markov chain is aperiodic and diagonalisable, and that the unique limit distribution, which exists due to the aperiodicity condition, does not lie on the boundary of any of the half-spaces used to define the property. The paper [19] also presents case studies in the areas of software reliability and medicine. Our work extends these previous works in three directions: we consider parametric Markov chains, we allow the Markov chains to be periodic, and we allow semialgebraic sets as atomic propositions. Due to new difficulties that arise in the parametric setting we do not cover full LTL, rather we handle arbitrary prefix-independent ω -regular properties.

Agrawal et al. [1] consider the model-checking problem of Markov chains under the distribution-transformer semantics, where the target sets are specified as intervals on each component. They also remark that full ω -regular model checking will not be possible in general, and instead they consider whether an approximation of the trajectory satisfies a property.

In [17], a related problem for Markov decision processes (MDPs) is studied. The orbit of an MDP is not fixed but depends on the scheduler, and this additional feature often leads to

undecidability. Several restrictions on schedulers and specifications are studied [17] under which decidability can be achieved. [10] considers a restriction on the MDP.

Markov chains under the distribution-transformer semantics are a special case of LDS. Presumably one is interested in expressive specifications, e.g., those that are specifiable in LTL or MSO. Unfortunately, even simple reachability queries for LDS are known to be extremely challenging [11], and the attendant hardness propagates to Markovian dynamical systems as well [2].

Baier et al. [5] showed that parametric *point-to-point reachability*, which asks whether there *exist* parameter choices under which a given state distribution is reachable, is decidable only for a single parameter, and Skolem-hard for two or more parameters. The problem is well-known to be decidable in polynomial time for LDS [16] (and thus for non-parametric Markov chains). We circumvent this limitation, allowing us to consider an arbitrary number of parameters, by synthesising, up to a null set, the set of parameter choices for which an arbitrary prefix-independent property holds (rather than reachability of a single point target).

Model checking prefix-independent properties on diagonalisable LDS is decidable [3] (see also [22] for Ultimate Positivity specifically). However, in general, the decidability status of the Ultimate Positivity Problem is a major open question—in fact, decidability of Ultimate Positivity for LDS of dimension 6 would entail major breakthroughs in number theory as it would solve certain longstanding open problems in Diophantine approximation of transcendental numbers that are widely believed to be hard [21].

Typically, only very few border cases are particularly difficult, and thus in the parametric setting such border cases amount to a null set which we can exclude. This is the case for all but degenerate instances in which *all* of the parameter valuations lead to such hard border cases. It is therefore necessary to impose a technical restriction on the expressible targets in order to exclude these degenerate instances.

The problem of model checking parametric Markov chains with respect to the standard trace semantics has been considered extensively [12, 20, 14, 6, 13]. In this setting one can express properties such as “the set of traces reaching a certain state has probability above λ ”, which can be described using standard logics such as PCTL [15, 7]. This semantics does not allow specifying properties such as “the probability of being in state s_1 is eventually larger than the probability of being in state s_2 ”, which can be expressed by the properties we consider.

2 Preliminaries

2.1 Parametric Markov chains

Given a set of variables X , we denote the field of rational functions over X with base field \mathbb{Q} by $\mathbb{Q}(X)$. We denote the set of all probability distributions over Q by $\text{Dist}(Q)$.

- **Definition 3.** A parametric Markov chain (PMC) is a tuple $\mathcal{M} = (Q, X, M, \pi)$, where
- Q is a finite set of states;
 - X is a finite set of variables, here typically called parameters;
 - $M \in \mathbb{Q}(X)^{Q \times Q}$ is the parametrised transition matrix;
 - $\pi \in \text{Dist}(Q)$ is an initial distribution.

Given a concrete instantiation $p \in \mathbb{R}^X$ of the parameters X , we denote by $M[p] \in \mathbb{R}^{Q \times Q}$ the matrix $M[p]_{s,t} = M_{s,t}(p)$, provided that $M_{s,t}(p)$ is defined for every $s, t \in Q$. That is, $M[p]$ is the concrete transition function obtained by replacing in M every occurrence of a parameter $v \in X$ by the value assigned to v in p . We call $p \in \mathbb{R}^X$ *admissible* if $M[p]$ is a

probabilistic transition function, i.e., $0 \leq M[p]_{s,t} \leq 1$ for all $s, t \in Q$, and $\sum_{t \in Q} M[p]_{s,t} = 1$ for all $s \in Q$. The Markov chain induced by the parameter value p will be denoted by $\mathcal{M}[p] = (Q, M[p], \pi)$. Finally, we remark that parametrised initial distributions can be encoded in our framework by adding a single state to the Markov chain that is visited only once at the beginning. The probabilities associated with the outgoing edges of the new start state are then used to simulate the parametrised initial probabilities.

2.2 The topological structure of a PMC

Throughout the paper we will use structural arguments about the *underlying graph* (or *topological structure*) of a Markov chain (Q, M, π) , which is defined as $(Q, \{(s, t) \mid M_{s,t} > 0\})$. For a parametric Markov chain (Q, X, M, π) we consider the *main structure* $(Q, \{(s, t) \mid \exists p. M[p]_{s,t} \neq 0\})$. That is, we only keep the entries of M that are not identically zero. We will show that the main structure matches the structure of $M[p]$ almost everywhere (that is, everywhere except possibly on a set with null measure). This means that w.l.o.g. we can assume that the given PMC has a constant topological structure. We begin by recalling a well-known fact which is immediate from the observation that a non-zero polynomial is non-zero almost everywhere (see, e.g., [9]).

► **Lemma 4.** *Any non-zero rational function $f \in \mathbb{Q}(X)$ is almost everywhere defined and non-zero.*

► **Lemma 5** (Constant topological structure). *Let $D \subseteq \mathbb{R}^X$ be the set of parameters defined as $D = \{p \mid p \text{ is admissible and } M[p] \text{ has the main structure}\}$. Then $\mathbb{R}^X \setminus D$ has null measure.*

Proof. Observe that

$$\mathbb{R}^X \setminus D = \bigcup_{s,t \in Q} \{p \mid M_{s,t} \text{ is not well-defined at } p\} \cup \bigcup_{s,t \in Q} \{p \mid \text{the structure of } M[p] \text{ differs from the main structure at } (s, t)\}$$

which is a finite union of sets of measure zero. Hence $\mathbb{R}^X \setminus D$ also has a null measure. ◀

Henceforth we define D to be the set described above.

We recall some basic structural notions about Markov chains. These descriptions also apply to the main structure of a parametric Markov chain. We say a collection of states $\mathcal{C} \subseteq Q$ is *strongly connected* if there is a path from any state to another in the restriction of the underlying graph of \mathcal{M} to \mathcal{C} . We only refer to (maximally) *strongly connected components* (SCCs), that is, SCCs for which there does not exist $s \in Q$ such that $\mathcal{C} \cup \{s\}$ is also strongly connected. A singleton state with no self loops and no other path to itself is considered its own SCC. Given an SCC \mathcal{C} , its *period* is the greatest common divisor of the lengths of cycles in \mathcal{C} . A SCC is called *aperiodic* if its period is 1, and otherwise it is called *periodic*. We say that a SCC \mathcal{C} is a *bottom SCC*, or *recurrent*, if for all $s \in \mathcal{C}$, $M_{s,s'} = 0$ for all $s' \in Q \setminus \mathcal{C}$. That is, no probability is lost from \mathcal{C} . If \mathcal{C} is not recurrent, it is called *transient*.

A Markov chain is *aperiodic* if all of its SCCs are aperiodic (and otherwise periodic) and *recurrent* if all its SCCs are recurrent. If the Markov chain consists of only one recurrent SCC then the Markov chain is said to be irreducible. As these properties are structural, depending only on the matrix M of $\mathcal{M} = (Q, X, M, \pi)$, we may say M is aperiodic or irreducible.

2.3 Semialgebraic targets

A set $T \subseteq \mathbb{R}^d$ is semialgebraic if it is a finite Boolean combination of sets specified by a polynomial inequality. That is, T can be obtained from sets of the form $\{x \in \mathbb{R}^d \mid f(x) \bowtie 0\}$ for some $\bowtie \in \{\geq, \leq, >, <, =\}$ using finitely many union and intersection operations. In fact, without loss of generality we can assume the sets to be of the form $\{x \in \mathbb{R}^d \mid f(x) \triangleright 0\}$ for $\triangleright \in \{\geq, >\}$. Written in disjunctive normal form, with \wedge corresponding to \cap and \vee corresponding to \cup , we can write T as $\bigcup_{i=1}^k \bigcap_{j=1}^{l_i} \{x \in \mathbb{R}^d \mid f_{ij}(x) \triangleright_{ij} 0\}$. Note that many restricted classes of target sets, such as singleton points and Boolean combinations of linear inequalities (e.g., polyhedra, halfspaces, and cones) are all examples of semialgebraic sets.

We will be considering the semialgebraic targets T_1, \dots, T_k within the universe of $U = \text{Dist}(Q) \subseteq \mathbb{R}^Q$, which will be endowed with the subspace topology with respect to the usual Euclidean topology on \mathbb{R}^Q . In this topology, a vector (i.e., a probability distribution) x is in the interior T° of a target T if and only if there exists $\epsilon > 0$ such that $B_\epsilon(x) \cap U \subseteq T$, where $B_\epsilon(x)$ is the ϵ -ball around x in \mathbb{R}^Q . We will be particularly interested in points on the boundary of T . The boundary of T , denoted ∂T , is the set of all limit points of T in U that are not in the interior of T . That is, $\partial T = \overline{T} \setminus T^\circ$, where \overline{T} is the closure of T in U .

We denote by $\text{vol}(D)$ the Lebesgue measure of a measurable set $D \subseteq \mathbb{R}^X$. Recalling that a vector v lies on the boundary of T if $v \in \partial T$, we say that a parametrised vector $(v[p])_{p \in D}$ is contained within the boundary of T if $v[p] \in \partial T$ **for all** $p \in D$. Given a parametrised vector, we will often be interested in the quantity $\text{vol}(\{p \in D \mid v[p] \in \partial T\})$.

2.4 Prefix-independent model checking

Let $\{T_1, \dots, T_k\}$ be a partition of the ambient space $\text{Dist}(Q)$ and $\Sigma = \{1, \dots, k\}$. Recall that properties over the predicates T_1, \dots, T_k are modelled by the subsets of Σ^ω . An ω -regular property P is *prefix-independent* if for every infinite word w and every finite word u acting as a prefix, $w \in P \iff uw \in P$. For such a property P it holds that for every $w, w' \in \Sigma^\omega$ that can be obtained from one another through finitely many insertions and deletions, $w \in P \iff w' \in P$.²

Given a property φ over Σ , we say a Markov chain \mathcal{M} satisfies φ , denoted $\mathcal{M} \models \varphi$, when the characteristic word of the Markov chain with respect to the targets T_1, \dots, T_k satisfies φ . In this paper we assume the property to be given as an ω -automaton (e.g., a non-deterministic Büchi automaton) over Σ . Then, one can check whether a given ultimately periodic word is accepted by such an automaton. This is done by checking non-emptiness on the automaton built by the product construction on the given automaton and an automaton for the ultimately periodic word. Properties given in other specification languages such as LTL or MSO can be handled by first creating an equivalent non-deterministic Büchi automaton, provided that the input property is prefix-independent.

2.5 Problems: synthesising parameters

First, we consider the set of parameters such that the sequence of distributions of the resulting Markov chain is ultimately trapped inside one of the target sets (“the positive set”). This is the parametric analogue of the well-known Ultimate Positivity Problem [22] (with the

² To see this, consider the common suffix v such that $w = uv$ and $w' = u'v$ and then observe that $w \in P \iff v \in P \iff u'v \in P$.

halfspace generalised to arbitrary semialgebraic set). Formally, we consider the following problem:

► **Problem 6** (Ultimate Positivity on PMCs). *Given a PMC $\mathcal{M} = (Q, X, M, \pi)$ and a semialgebraic set $T \subseteq \text{Dist}(Q)$, and letting $D \subseteq \mathbb{R}^X$ be the set of admissible parameter instantiations that give rise to the main structure, synthesise the set of feasible parameters $\{p \in D \mid \exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^n \in T\}$.*

Since the set of parameters could give rise to concrete instances which are hard, we do not synthesise the full set of feasible parameters exactly, but rather compute a semialgebraic subset that differs from the full set by a null set. In particular all of the parameter valuations in the set we compute give rise to an ultimately positive instance. If the computed semialgebraic set is non-empty, one can be sure that there does exist a parameter choice satisfying the property, and that such a parameter valuation can be computed. However, if the set is empty, then one cannot be sure that there does not exist a choice; but in this case one would know that even if there is such a parameter choice, there are ‘not too many choices’.

In Theorem 9 of Section 3 we compute this set for aperiodic recurrent finite-state parametric Markov chains, before generalising the result to periodic Markov chains in Theorem 13 in Section 4.

Being ultimately trapped inside a semialgebraic set is a prefix-independent property. Next, we generalise the problem to any prefix-independent property.

► **Problem 7** (Prefix-independent model checking on PMCs). *Given*

- a PMC $\mathcal{M} = (Q, X, M, \pi)$,
- semialgebraic sets T_1, \dots, T_k which form a partition of $\text{Dist}(Q)$, and
- a prefix-independent property φ over T_1, \dots, T_k ,

and letting $D \subseteq \mathbb{R}^X$ be the set of admissible parameter instantiations that give rise to the main structure, synthesise the set of the feasible parameters, i.e., those satisfying φ : $\{p \in D \mid \mathcal{M}[p] \models \varphi\}$.

In Theorem 15 of Section 5 we compute a semialgebraic subset of the feasible parameters differing up to a null set for the prefix-independent model checking problem.

3 Synthesising satisfying parameters for Ultimate Positivity in aperiodic and irreducible PMCs

Let $\mathcal{M} = (Q, X, M, \pi)$ be an aperiodic and irreducible Markov chain. It is well-known that any such Markov chain has a unique stationary distribution and that this distribution is also the unique limit distribution. In our case this means that for every choice of parameters $p \in D$ there is a unique probability distribution $\mu[p] \in \text{Dist}(Q)$ such that $\mu[p] \cdot M[p] = \mu[p]$. The fact that \mathcal{M} is irreducible implies that $\mu[p]$ will be strictly positive in each entry. The following lemma assures that this distribution is also a rational function in X and can be effectively computed.

► **Lemma 8.** *Given an aperiodic and irreducible PMC $\mathcal{M} = (Q, X, M, \pi)$, let $D \subseteq \mathbb{R}^X$ be the set of admissible parameters leading to the main structure of \mathcal{M} . There exists a unique parametric limit distribution $\mu \in \mathbb{Q}(X)^Q$ such that $\lim_{n \rightarrow \infty} \pi \cdot M[p]^n = \mu[p]$ for all $p \in D$. Furthermore, μ can be effectively computed.*

Proof. The stationary distribution $\mu : D \rightarrow \text{Dist}(Q)$ is the unique solution of the linear equation system $\mu[p] \cdot M[p] = \mu[p]$ in probability distributions. Hence μ can be computed by

performing Gaussian elimination on the system $\mu[p] \cdot M[p] = \mu[p]$ followed by a normalisation step. This shows that every entry $\mu[p]_s$ of $\mu[p]$ is a rational function in p . ◀

We now establish the main theorem of our paper, showing how to compute a semialgebraic set of parameters which, up to a null set, equals the set of admissible parameters satisfying Ultimate Positivity. Our approach relies on the assumption that the volume of limit distributions lying on the boundary of a target T is null, that is, $\text{vol}(\{p \in D \mid \mu[p] \in \partial T\}) = 0$. We say that an instance of Problem 6 is *degenerate* if $\text{vol}(\{p \in D \mid \mu[p] \in \partial T\}) > 0$. If one considers only half-spaces as target sets, our requirement of non-degeneracy corresponds exactly to the third condition of [19, Theorem 1], which tackles the corresponding model-checking question for non-parametric Markov chains.

To see why such an assumption is strictly needed we show that, without this assumption, Problem 6 is as hard as ultimate positivity. That is, one would need to answer potentially intractable instances of the Ultimate Positivity Problem. Consider the following scenario. There is a single parameter p and all of its instantiations lead to the same non-parametric Markov chain \mathcal{M} . For any Markov chain \mathcal{M} , such a parametric Markov chain $\mathcal{M}[p]$ can easily be constructed. Recall that the Ultimate Positivity Problem for stochastic matrices asks whether there exists N such that for all $n \geq N$, it holds that $(\pi \cdot M^n)_s \geq 1/2$, for a given stochastic matrix M , an initial vector π and a state s [2]. Ultimate Positivity Problem for stochastic matrices can be easily expressed as an Ultimate Positivity Problem for PMCs (Problem 6). Then, the answer to the non-parametric Ultimate Positivity instance is yes if and only if the measure of parameters satisfying the formula is 1 (and in case the answer is no, the computed semialgebraic set will be empty, having measure zero). The decidability status of the Ultimate Positivity Problem is a major open question. However, it is solvable if the limit distribution of M in state s is not zero. Our non-degeneracy assumption essentially excludes the currently intractable cases of this problem.

► **Theorem 9.** *Consider a non-degenerate instance of Problem 6, in which the following are given:*

- *an aperiodic and irreducible PMC $\mathcal{M} = (Q, X, M, \pi)$, for which $D \subseteq \mathbb{R}^X$ is the semialgebraic set of admissible parameter values that give rise to the main structure,*
- *the parametric limit distribution $\mu \in \mathbb{Q}(X)^Q$ such that $\lim_{n \rightarrow \infty} \pi \cdot M[p]^n = \mu[p]$ for all $p \in D$, and*
- *a semialgebraic set $T = \bigcup_{i=1}^k \bigcap_{j=1}^{l_i} \{x \in \text{Dist}(Q) \mid f_{ij}(x) \triangleright_{ij} 0\}$, for which $\text{vol}(\{p \in D \mid \mu[p] \in \partial T\}) = 0$.*

Then a semialgebraic set D'_T , contained in $D_T = \{p \in D \mid \exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^n \in T\}$ but differing from D_T only by a null set, can be effectively computed.

Proof. Since for all $p \in D$, we have $\lim_{n \rightarrow \infty} \pi M[p]^n = \mu[p]$ then for all $p \in D$ such that $\mu[p] \in T^\circ$ it holds that there exists N such that for all $n \geq N$, $\pi M[p]^n \in T$. Clearly, if $p \in D$ is such that $\mu[p] \notin \bar{T}$, then the sequence of distributions of $\mathcal{M}[p]$ is eventually outside of T . It remains to consider the case where $\mu[p] \in \partial T$. Since by our assumption $\text{vol}(\{p \in D \mid \mu[p] \in \partial T\}) = 0$ it holds that D_T differs from $D'_T = \{p \in D \mid \mu[p] \in T^\circ\}$ by only a null set. Therefore it suffices to show how to compute a representation for D'_T .

The set D'_T is a semialgebraic set, for which an implicit representation in the first order theory of the reals can be found in polynomial time. To see this, observe that $D'_T = \{p \in \mathbb{R}^X \mid p \in D \wedge \exists y. y = \mu[p] \wedge y \in T^\circ\}$, with also D semialgebraic. The set T° is itself a semialgebraic set, which can easily be seen by specification in the theory of the reals as $\{x \in U \mid \exists \epsilon > 0. \forall z \in U, |z - x| \leq \epsilon \implies z \in T\}$ (recall from Section 2.3 that $U = \text{Dist}(Q)$ is the universe of probability distributions over Q). Finally, $z \in T$ can be

expressed in the theory of the reals by asserting that $\bigvee_{i=1}^k \bigwedge_j^{l_i} f_{ij}(x) \triangleright_{ij} 0$ where $f_{i,j} \triangleright_{ij} 0$ are the polynomial inequalities defining T . This concludes the proof in case one is satisfied with the set D'_T represented in the first order theory of the reals. In case an explicit form is required, quantifier elimination can be used to compute D'_T as boolean combination of polynomial inequalities, i.e., of the form $\{x \in \mathbb{R}^X \mid \bigwedge_i \bigvee_j g_{ij}(x) \triangleright_{ij} 0\}$ [23, Theorem 1.2]. ◀

► **Remark 10.** Since the Lebesgue measure is complete (i.e., every subset of a null set is measurable), it follows from the second part of Theorem 9 that the set D_T is Lebesgue measurable and hence Problem 6 is well-defined.

► **Remark 11.** Observe that it is decidable whether the instance is degenerate. This amounts to asking whether $\text{vol}(\{p \in D \mid \mu[p] \in \partial T\}) = 0$, which is the case if and only if the interior of the set is empty. The set $D_{\partial T} = \{p \in D \mid \mu[p] \in \partial T\}$ is semialgebraic, thus the interior $D_{\partial T}^\circ$ is also semialgebraic, for which one can test emptiness.

Secondly, we note that degenerate instances are somewhat unlikely. Recall the limit distribution is a rational function. Should the limit distribution coincide with the boundary of T for a positive volume of points then the function must essentially correspond with one of the polynomials defining a boundary of T . If the difference is zero with positive measure then by Lemma 4 the difference is the zero function, and we conclude that they must be the same function. This would seem to indicate that the target had been constructed adversarially with a priori knowledge of the limit distribution.

3.1 Complexity

Together Lemma 8 (which shows how to compute μ using Gaussian elimination) and Theorem 9 produce, up to a null set, the set of parameters of \mathcal{M} satisfying ultimate positivity for a target T in the case that \mathcal{M} is an irreducible and aperiodic PMC. We now consider the complexity of this reduction.

In general, the number of terms of a rational functions one gets from applying Gaussian elimination over the field of rational functions may become exponential. However, for a fixed number of parameters the parametrised stationary distribution μ (from Lemma 8) can be computed in polynomial time using *fraction-free* Gaussian elimination [6] (thus the representation of μ needs at most polynomial space).

It is then straightforward to see that the implicit representation, given as a sentence in the first order theory of the reals, can be found in polynomial time. We consider the complexity of computing the explicit representation in the following lemma and observe that this is polynomial time for fixed Markov chains \mathcal{M} .

► **Lemma 12.** *The explicit representation of D'_T can be found in time $p(x)^{O(|X||Q|^2)}$, where p is a polynomial in the size of the inputs $\mathcal{M} = (Q, X, M, \pi)$, μ , and T (represented by x).*

Proof. Consider an implicit description of a semialgebraic set given by a sentence in the theory of the reals of the form $\{y \in \mathbb{R}^\ell \mid Q_1 x_1 \in \mathbb{R}^{n_1} \dots Q_\omega x_\omega \in \mathbb{R}^{n_\omega} P(B_1(y, x), \dots, B_m(y, x))\}$, where $Q_i \neq Q_{i+1}$ are quantifiers in $\{\exists, \forall\}$, P is a Boolean formula in m variables and the B_i 's are polynomial inequalities of degree at most d in variables from x_1, \dots, x_ω and integer coefficients of bit-size at most L . Define $K_1 = \ell \prod_{k=1}^\omega n_k$ and $K_2 = \ell + \sum_{k=1}^\omega n_k$. By Theorem 1.2 of [23] the explicit description can be found from the implicit description using $L^{O(1)}(md)^{2^{O(\omega)}K_1}$ many arithmetic operations and $(md)^{K_2}$ evaluations of the Boolean formula P .

The formula described in Theorem 9 implicitly uses inequalities on rational functions, with rational coefficients. Observe that this can be converted to polynomial inequalities

with integer coefficients, e.g., $f(x)/g(x) > 0 \iff g(x) \neq 0 \wedge f(x)g(x) > 0$. Then, rational coefficients can be removed by multiplying through by the lcm of denominators.

In the proof of Theorem 9, the implicit representation of D'_T is constructed in polynomial time by suitably describing the set in the first order theory of the reals, in time polynomial in the sizes of \mathcal{M} , μ and T . Let $q(x)$ be such a polynomial. We observe that the resulting representation has bounded quantifier alternation. In particular, composing the descriptions of D_T and T° into a single formula, the description has free parameters $p \in D$ (thus $\ell = |X|$) and quantification of the form $\exists y \in \text{Dist}(Q), \epsilon > 0, \forall z \in \text{Dist}(Q)$ followed by a Boolean combination of polynomial inequalities. Hence, there are two blocks of quantifiers ($\omega = 2$), of size $n_1 = |Q| + 1$ and $n_2 = |Q|$. The degree d of the polynomials and the number of such polynomial inequalities m are polynomial in the same parameters to describe T and μ . Let $u(x)$ be such a polynomial.

Since $|P|$ is at most $q(x)$, the Boolean formula can be evaluated in linear time, i.e., $q(x)$. Thus the conversion to explicit representation using the procedure of Renegar thus takes $O(u(x)^{O(|X||Q|^2)}q(x))$ many operations. But $O(u(x)^{O(|X||Q|^2)}q(x)) = O(p(x)^{O(|X||Q|^2)})$ for some larger polynomial $p(x)$ (assuming $|X| \neq 0$ and $|Q| \neq 0$), which concludes the proof. \blacktriangleleft

Recall that when $|X|$ is fixed then μ can be computed in time polynomial in \mathcal{M} , then the size of x is itself polynomial. Further, when the size of $|Q|$ is fixed, then the procedure is polynomial time in the size of \mathcal{M} and T and polynomial in T when \mathcal{M} is fixed (the parametric Markov chain may be considered fixed when the chain is given but the problem needs to be considered for several possible targets).

4 The limit distribution of periodic Markov chains with transient states

We have observed that we can compute the parametric stationary distribution for aperiodic and irreducible PMCs. Next, we show how to drop both of these restrictions by handling periodicity and transient states.

4.1 Managing periodicity

We observe that we can assume that the matrix is aperiodic for all parameters by considering subsequences. Recall that we can assume that the topological structure of $\mathcal{M} = (Q, X, M, \pi)$ is constant. When a Markov chain is periodic with period H we have that M^H is aperiodic. We consider H many parametric Markov chains $\mathcal{M}^{(h)} = (Q, X, M^H, \pi \cdot M[p]^h)$ for each $h \in \{0, \dots, H-1\}$, each leading to the parametric orbit $(\pi \cdot M[p]^h (M[p]^H)^n)_n$. Each Markov chain $\mathcal{M}^{(h)}$ has the same aperiodic update matrix $M[p]^H$ but a different starting point $\pi \cdot M[p]^h$. For reachability questions, we can simply analyse each subsequence independently, although we must suitably interleave the results if considering more general properties.

4.2 Managing transient states

Secondly, we consider transient states, that is, the states outside of a bottom strongly connected component. Let us assume $\mathcal{M} = (Q, X, M, \pi)$ is an aperiodic Markov chain. We know from the standard literature that the limit probability for any transient state is zero. However, we must decide how much of the total weight which started in a transient state ultimately reaches each of the bottom strongly connected components and weight the respective stationary distributions accordingly.

We are interested in the absorption probability of each bottom SCC. We consider a new (parametric) Markov chain (Q', X, N, π) where each BSCC \mathcal{C} is reduced to a single, aperiodic, absorbing state $q_{\mathcal{C}}$. Let $B = \{b_{\mathcal{C}} \mid \mathcal{C} \text{ is a bottom SCC.}\}$, $F = \{q \in Q \mid q \text{ is transient}\}$ and $Q' = F \cup B$, the set of transient states and the new representative bottom states. Let $N \in \mathbb{Q}(X)^{Q' \times Q'}$, be defined such that $N_{q,q'} = M_{q,q'}$ if q, q' are in F , $N_{q,b_{\mathcal{C}}} = \sum_{q' \in \mathcal{C}} M_{q,q'}$, $N_{b_{\mathcal{C}},b_{\mathcal{C}}} = 1$ and $N_{b_{\mathcal{C}},q'} = 0$ if $q' \neq b_{\mathcal{C}}$.

We compute absorbing probabilities $a \in \mathbb{Q}(X)^{Q' \times B}$, where $a_{q,b_{\mathcal{C}}}$ is the probability of reaching bottom state $b_{\mathcal{C}}$ starting in state q . Note that this is parametric in variables X . To compute a , we solve the linear equation system, where for each $b_{\mathcal{C}} \in B$ we require that $a_{b_{\mathcal{C}},b_{\mathcal{C}}} = 1$ (every bottom SCC is absorbing), $a_{q,b_{\mathcal{C}}} = 0$ if q cannot reach $b_{\mathcal{C}}$, and $a_{q,b_{\mathcal{C}}} = \sum_{q' \in F} N_{q,q'} a_{q',b_{\mathcal{C}}}$ if q can reach $b_{\mathcal{C}}$.

We can also compute the limit distribution $\mu^{\mathcal{C}}$ for each bottom strongly connected component \mathcal{C} in isolation, this is the stationary distribution as computed in Lemma 8. We can then reweight these stationary distributions according to the probability which reaches each bottom SCC using the absorbing probabilities. For a bottom strongly connected component \mathcal{C} , the limit distribution of state $s \in \mathcal{C}$, is $\ell[p]_s = (\sum_{q \in Q'} \pi_q a[p]_{q,b_{\mathcal{C}}} + \sum_{q \in \mathcal{C}} \pi_q) \mu^{\mathcal{C}}[p]_s$, when the initial distribution is π . For states not in any bottom strongly connected component, $s \in F$, we have $\ell[p]_s = 0$. Note that ℓ is also a rational function, since it is simply the product and sums of functions found by Gaussian elimination.

4.3 Managing periodic Markov chains with transient states

We now induce a limit distribution for each of the H aperiodic Markov chains $(\mathcal{M}^{(h)})_{h=0}^{H-1}$ found in Section 4.1. For each such chain, the matrix is $M[p]^H$, and we assume the stationary distributions $\mu^{\mathcal{C}}$ for each bottom SCC \mathcal{C} (this does not depend on h). However, we must consider the limit distribution for each of the H starting points we consider. That is the initial distribution is $\pi \cdot M[p]^h$ for each $h \in \{0, \dots, H-1\}$. Thus for each subsequence, distinguished by h , we can compute a unique limit distribution, where

$$\ell^{(h)}[p]_s = \left(\sum_{q \in Q'} (\pi \cdot M[p]^h)_q a[p]_{q,b_{\mathcal{C}}} + \sum_{q \in \mathcal{C}} (\pi \cdot M[p]^h)_q \right) \mu^{\mathcal{C}}[p]_s \quad \text{for } s \in \mathcal{C}, \text{ and}$$

$$\ell^{(h)}[p]_s = 0 \quad \text{for } s \text{ transient,}$$

such that $\lim_{n \rightarrow \infty} (M[p]^H)^n (\pi \cdot M[p]^h) = \ell^{(h)}[p]$ for all $p \in D$.

4.4 Ultimate Positivity in the general case

Using the limit distribution established in this section, we now complete the proof of ultimate Positivity for periodic Markov chains with transient states.

► **Theorem 13.** *Consider a non-degenerate instance of Problem 6, in which the following are given:*

- $\mathcal{M} = (Q, X, M, \pi)$ is a PMC with period H , for which $D \subseteq \mathbb{R}^X$ is the semialgebraic set of admissible parameter values that give rise to the main structure,
- H parametric limit distributions $\ell^{(h)} \in \mathbb{Q}(X)$ for $h \in \{0, \dots, H-1\}$ such that $\lim_{n \rightarrow \infty} (M[p]^H)^n (\pi \cdot M[p]^h) = \ell^{(h)}[p]$ for all $p \in D$.
- a semialgebraic set $T = \bigcup_{i=1}^k \bigcap_{j=1}^{l_i} \{x \in \text{Dist}(Q) \mid f_{ij}(x) \triangleright_{ij} 0\}$, such that, for all limit distributions $(\ell^{(h)})_{h=0}^{H-1}$, we have $\text{vol}(\{p \in D \mid \ell^{(h)}[p] \in \delta T\}) = 0$.

Then a semialgebraic set D'_T , contained in $D_T = \{p \in D \mid \exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^n \in T\}$ but differing from D_T only by a null set, can be effectively computed.

Proof. Ultimate Positivity requires that $\exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^n \in T$. Since we have split the orbit into H subsequences, we require that *all* of these subsequences eventually enter and remain inside T . That is there exists $N \in \mathbb{N}$ such that for all $h \in \{0, \dots, H\}$ and all $n \geq N$ we have $\pi \cdot M[p]^h (M[p]^H)^n \in T$. To compute a set D'_T as required, we use Theorem 9 on the limit distribution $\ell^{(h)}$. This gives us a set D'_h contained in $\{p \in D \mid \exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^h (M[p]^H)^n \in T\}$, but of the same measure as D_h , for each h . Then, the set $D'_T = \bigcap_{h \in \{0, \dots, H\}} D'_h$ satisfies the requirements which concludes the proof. \blacktriangleleft

► **Remark 14.** In the case \mathcal{M} is aperiodic but not irreducible then the complexity result of Section 3.1 also applies. Note that in general the period H may be exponential, thus the periodic case requires consideration of exponentially many subsequences, for which the matrix of the system is M^H , which could be much larger than M .

5 Synthesising satisfying parameters for prefix-independent model checking

Finally we show how to compute a set with volume equivalent to the parameters which induce a Markov chain satisfying a prefix-independent property.

Consider semialgebraic targets T_1, \dots, T_k which partition $\text{Dist}(Q)$. That is $T_i \cap T_j = \emptyset$ for $i \neq j$ and $\text{Dist}(Q) = T_1 \cup \dots \cup T_k$.

We generalise the notion of *degenerate* instances to multiple targets. We say that an instance is *non-degenerate* if the the volume of any of the limit distributions lying on the boundary of any of the targets is zero. That is, $\text{vol}(\{p \in D \mid \ell^{(h)}[p] \in \partial T_i\}) = 0$ for each T_i and $\ell^{(h)}$. This allows us to be sure that every subsequence is eventually inside one of the targets, for all but a null-set of parameters.

► **Theorem 15.** Consider a non-degenerate instance of Problem 7, in which the following are given:

- $\mathcal{M} = (Q, X, M, \pi)$ is a PMC, with period H , for which $D \subseteq \mathbb{R}^X$ is the semialgebraic set of admissible parameter values that give rise to the main structure,
- H parametric limit distributions $\ell^{(h)} \in \mathbb{Q}(X)$ for $h \in \{0, \dots, H-1\}$ such that $\lim_{n \rightarrow \infty} (M[p]^H)^n (\pi \cdot M[p]^h) = \ell^{(h)}[p]$ for all $p \in D$,
- T_1, \dots, T_k are semialgebraic targets partitioning $\text{Dist}(Q)$ such that, for all limit distributions $(\ell^{(h)})_{h=0}^{H-1}$, and all targets T_i , we have $\text{vol}(\{p \in D \mid \ell^{(h)}[p] \in \partial T_i\}) = 0$, and
- φ is a prefix-independent ω -regular property over T_1, \dots, T_k .

Then, a semialgebraic set D'_φ , contained in $D_\varphi = \{p \in D \mid \mathcal{M}[p] \models \varphi\}$, but differing from D_φ only by a null set, can be effectively computed.

Proof. We know that any aperiodic Markov chain M will eventually converge to its limit distribution ℓ , that is, for any ϵ for sufficiently large n we have $|\pi M^n - \ell| < \epsilon$. So if the limit distribution is not on the boundary of a target, eventually the Markov chain stays inside the target or outside the target.

Hence, for all but a null set of $p \in D$ and each $h \in \{0, \dots, H-1\}$, we have that the orbit $\pi M[p]^h (M[p]^H)^n$ enters, and stays in, exactly one of T_1, \dots, T_k from some point on. Given p , we can determine this final target by checking in which set T_1, \dots, T_k the point $\ell^{(h)}[p]$ lies. Then, since $\pi M[p]^h (M[p]^H)^n$ is stationary from some point on, we have that every

H th character of the characteristic word of $\mathcal{M}[p]$ w.r.t. T_1, \dots, T_k is fixed, and therefore the characteristic word is eventually periodic.

We consider each of the possible k^H periodic words describing the limit behaviour. That is, we consider w^ω for a word $w \in \{1, \dots, k\}^H$, i.e., w repeated infinitely many times. For all but a null set of parameters, the resulting characteristic word must have such a suffix. We can model check each such word, and decide if the word satisfies the property φ by asking whether w is accepted by the automaton representing φ .

We discard the parameter values leading to a periodic suffix which does not satisfy the specification φ . However, for each periodic word w that does satisfy φ , we compute $D'_w \subseteq D$ which, up to a null set, represents the parameters leading to this word. Fix $w \in \{1, \dots, k\}^H$. We compute, up to a null set, the set of parameters for which the periodic word of $\mathcal{M}[p]$ matches w at each position. Using Theorem 9 on limit distribution $\ell^{(h)}$ and target T_{w_h} compute the set $D'_{w,h} \subseteq D$, $D'_{w,h} = \{p \in D \mid \exists N \in \mathbb{N}. \forall n \geq N. \pi \cdot M[p]^h (M[p]^H)^n \in T_{w_h}\}$. To represent the whole word, we take intersection, that is let $D'_w = \bigcap_{h \in \{0, \dots, H-1\}} D'_{w,h}$.

Finally, we compute $D'_\varphi = \bigcup_{w \in \{1, \dots, k\}^H \text{ s.t. } w^\omega \models \varphi} D'_w$, which is contained in D_φ (the set of parameters for which the PMC satisfies the property φ), and differs from D_φ by at most a null set. \blacktriangleleft

A Reachability properties may fail to have semialgebraic feasible sets

We have shown how to compute, for prefix-independent specifications, the feasible parameters up to a null set, whilst carefully circumventing hard instances. Example 1 demonstrates that our approach will not extend in general to properties that depend on the prefix of the characteristic word. However, the property considered ($B \text{ U } L$) appears more complicated than a reachability property. In this section we give a parametric Markov chain for which the set of parameters satisfying a reachability property cannot be represented as a semialgebraic set, even up to a null set.

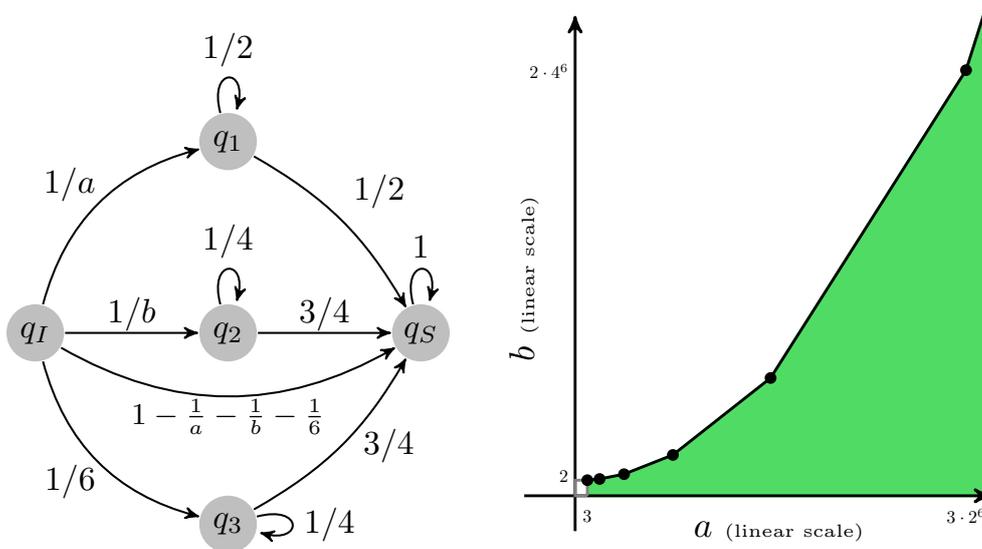
Let \mathcal{M} be the 2-parameter Markov chain with states $(q_I, q_1, q_2, q_3, q_S)$ depicted in Figure 3a. For convenience we restrict the parameters to $D = \{(a, b) \mid a > 3, b > 2\}$. The initial distribution of \mathcal{M} is $(1, 0, 0, 0, 0)$, the limit distribution is $(0, 0, 0, 0, 1)$ and after $n > 0$ steps the probability of being in states q_1, q_2, q_3 is $\frac{1}{a} \frac{1}{2^{n-1}}, \frac{1}{b} \frac{1}{4^{n-1}}, \frac{1}{6} \frac{1}{4^{n-1}}$, respectively. Let $T = \{(u, x, y, z, w) \mid \frac{2}{3z} - \frac{2}{x} + \frac{6z}{y} < 0\}$ be the semialgebraic target³. Observe that the specification is non-degenerate. We will show that the set $D_T = \{(a, b) \in D \mid \exists n. \pi \cdot M^n[(a, b)] \in T\}$ is not semialgebraic, even up to a null set.

By definition of T it holds that $(a, b) \in D_T$ if and only if the linear recurrence sequence $u_n = 4^n - 2^n a + b$ is negative for some n . Hence

$$D_T = \bigcup_{n \in \mathbb{N}} \{(a, b) \mid 4^n - 2^n a + b < 0\} = \bigcup_{n \in \mathbb{N}} \{(a, b) \mid b < 2^n a - 4^n\}.$$

By analysing the family of inequalities above we can show that the set D_T is a polytope with infinite vertices $\{(3 \cdot 2^n, 2 \cdot 4^n) \mid n \in \mathbb{N}\}$, as depicted in Figure 3b. To prove the desired result, assume for contradiction that there exists semialgebraic $S \subseteq D_T$ such that the measure of $D_T \setminus S$ is null. As D_T is open, it follows that $D_T \subseteq \text{Interior}(\text{Closure}(S))$. On the other hand, inspecting the accumulation points of D_T yields the reverse containment, so that $D_T = \text{Interior}(\text{Closure}(S))$, whence D_T itself is semialgebraic. Finally, observe that we can write the set of vertices $V = \{(3 \cdot 2^n, 2 \cdot 4^n) \mid n \in \mathbb{N}\}$ as the set of all points on the boundary

³ One can write $T = \{(u, x, y, z, w) \mid \frac{2}{3}xy - 2yz + 6xz^2 < 0\}$ in order to make the inequality a polynomial.



(a) Markov chain with parameters $p = (a, b)$. (b) Parameters (green) for which $\pi \cdot M[p]^n$ hits T .

■ **Figure 3** A parametric Markov chain \mathcal{M} and the parameter set satisfying reachability in T .

of D_T that cannot be expressed as a convex combination of two distinct points in $D \setminus D_T$. That is,

$$V = \{x \in \text{Closure}(D_T) \setminus D_T \mid \neg \exists y, z \in D \setminus D_T. y \neq z \wedge \exists \lambda \in (0, 1). x = \lambda y + (1 - \lambda)z\}.$$

This in turn makes V a semialgebraic set. However V is an infinite discrete set, and as such has infinitely many distinct connected components, contradicting a well-known property enjoyed by semialgebraic sets.

References

- 1 Manindra Agrawal, S. Akshay, Blaise Genest, and P. S. Thiagarajan. Approximate verification of the symbolic dynamics of markov chains. *J. ACM*, 62(1):2:1–2:34, 2015. doi:10.1145/2629417.
- 2 S. Akshay, Timos Antonopoulos, Joël Ouaknine, and James Worrell. Reachability problems for Markov chains. *Inf. Process. Lett.*, 115(2):155–158, 2015. doi:10.1016/j.ipl.2014.08.013.
- 3 Shaull Almagor, Toghrul Karimov, Edon Kelmendi, Joël Ouaknine, and James Worrell. Deciding ω -regular properties on linear recurrence sequences. *Proc. ACM Program. Lang.*, 5(POPL):1–24, 2021. doi:10.1145/3434329.
- 4 Tomáš Babiak, Mojmir Křetínský, Vojtěch Řehák, and Jan Strejček. LTL to Büchi Automata Translation: Fast and More Deterministic. In Cormac Flanagan and Barbara König, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science, pages 95–109, Berlin, Heidelberg, 2012. Springer. doi:10.1007/978-3-642-28756-5_8.
- 5 Christel Baier, Florian Funke, Simon Jantsch, Toghrul Karimov, Engel Lefauchaux, Florian Luca, Joël Ouaknine, David Purser, Markus A. Whiteland, and James Worrell. The orbit problem for parametric linear dynamical systems. In Serge Haddad and Daniele Varacca, editors, *32nd International Conference on Concurrency Theory, CONCUR 2021*, volume 203 of *LIPICs*, pages 28:1–28:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.CONCUR.2021.28.

- 6 Christel Baier, Christian Hensel, Lisa Hutschenreiter, Sebastian Junges, Joost-Pieter Katoen, and Joachim Klein. Parametric Markov chains: PCTL complexity and fraction-free Gaussian elimination. *Information and Computation*, 272:104504, jun 2020. doi:10.1016/j.ic.2019.104504.
- 7 Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008.
- 8 Saugata Basu, Richard Pollack, and Marie-Françoise Roy. Semi-algebraic sets. In *Algorithms in Real Algebraic Geometry*, pages 83–99. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. doi:10.1007/3-540-33099-2_4.
- 9 Richard Caron and Tim Traynor. The zero set of a polynomial. Technical report, WSMR report 05-02, 2005.
- 10 Rohit Chadha, Vijay Anand Korthikanti, Mahesh Viswanathan, Gul Agha, and YoungMin Kwon. Model checking mdps with a unique compact invariant set of distributions. In *Eighth International Conference on Quantitative Evaluation of Systems, QEST 2011, Aachen, Germany, 5-8 September, 2011*, pages 121–130. IEEE Computer Society, 2011. doi:10.1109/QEST.2011.22.
- 11 Ventsislav Chonev, Joël Ouaknine, and James Worrell. The polyhedron-hitting problem. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 940–956. SIAM, 2014.
- 12 Conrado Daws. Symbolic and Parametric Model Checking of Discrete-Time Markov Chains. In Zhiming Liu and Keijiro Araki, editors, *Theoretical Aspects of Computing - ICTAC 2004*, Lecture Notes in Computer Science, pages 280–294, Berlin, Heidelberg, 2005. Springer. doi:10.1007/978-3-540-31862-0_21.
- 13 Christian Dehnert, Sebastian Junges, Nils Jansen, Florian Corzilius, Matthias Volk, Harold Buntjes, Joost-Pieter Katoen, and Erika Ábrahám. PROPhESY: A PRObabilistic ParamETER SYnthesis Tool. In Daniel Kroening and Corina S. Păsăreanu, editors, *Computer Aided Verification*, Lecture Notes in Computer Science, pages 214–231, Cham, 2015. Springer International Publishing. doi:10.1007/978-3-319-21690-4_13.
- 14 Ernst Moritz Hahn, Holger Hermanns, and Lijun Zhang. Probabilistic reachability for parametric Markov models. *International Journal on Software Tools for Technology Transfer*, 13(1):3–19, January 2011. doi:10.1007/s10009-010-0146-x.
- 15 Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, September 1994. doi:10.1007/BF01211866.
- 16 Ravindran Kannan and Richard J. Lipton. Polynomial-time algorithm for the orbit problem. *J. ACM*, 33(4):808–821, 1986. doi:10.1145/6490.6496.
- 17 Vijay Anand Korthikanti, Mahesh Viswanathan, Gul Agha, and YoungMin Kwon. Reasoning about MDPs as Transformers of Probability Distributions. In *2010 Seventh International Conference on the Quantitative Evaluation of Systems*, pages 199–208, September 2010. doi:10.1109/QEST.2010.35.
- 18 YoungMin Kwon and Gul Agha. Linear Inequality LTL (iLTL): A Model Checker for Discrete Time Markov Chains. In Jim Davies, Wolfram Schulte, and Mike Barnett, editors, *Formal Methods and Software Engineering*, Lecture Notes in Computer Science, pages 194–208, Berlin, Heidelberg, 2004. Springer. doi:10.1007/978-3-540-30482-1_21.
- 19 YoungMin Kwon and Gul Agha. Verifying the Evolution of Probability Distributions Governed by a DTMC. *IEEE Transactions on Software Engineering*, 37(1):126–141, January 2011. doi:10.1109/TSE.2010.80.
- 20 Ruggero Lanotte, Andrea Maggiolo-Schettini, and Angelo Troina. Parametric probabilistic transition systems for system design and analysis. *Formal Aspects of Computing*, 19(1):93–109, March 2007. doi:10.1007/s00165-006-0015-2.
- 21 Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 366–379. SIAM, 2014. doi:10.1137/1.9781611973402.27.

- 22 Joël Ouaknine and James Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014*, volume 8573 of *Lecture Notes in Computer Science*, pages 330–341. Springer, 2014. doi:10.1007/978-3-662-43951-7_28.
- 23 James Renegar. On the computational complexity and geometry of the first-order theory of the reals, part I: introduction. preliminaries. the geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. *J. Symb. Comput.*, 13(3):255–300, 1992. doi:10.1016/S0747-7171(10)80003-3.