

Parametric Model Checking Continuous-Time Markov Chains

Catalin-Andrei Ilie¹

Department of Computer Science, University of Bucharest, Romania
cilie@fmi.unibuc.ro

James Ben Worrell

Department of Computer Science, University of Oxford, UK
jbw@cs.ox.ac.uk

Abstract

CSL is a well-known temporal logic for specifying properties of real-time stochastic systems, such as continuous-time Markov chains. We introduce PCSL, an extension of CSL that allows using existentially quantified parameters in timing constraints, and investigate its expressiveness and decidability over properties of continuous-time Markov chains. Assuming Schanuel’s Conjecture, we prove the decidability of model checking the one-parameter fragment of PCSL on continuous-time Markov chains. Technically, the central problem we solve (relying on Schanuel’s Conjecture) is to decide positivity of real-valued exponential polynomial functions on bounded intervals. A second contribution is to give a reduction of the Positivity Problem for matrix exponentials to the PCSL model checking problem, suggesting that it will be difficult to give an unconditional proof of the decidability of model checking PCSL.

2012 ACM Subject Classification Theory of computation → Logic and verification; Theory of computation → Verification by model checking; Theory of computation → Random walks and Markov chains

Keywords and phrases Probabilistic Continuous Stochastic Logic, Continuous-time Markov Chains, model checking, Schanuel’s Conjecture, positivity problem

Digital Object Identifier 10.4230/LIPIcs.TIME.2020.7

Funding *James Ben Worrell*: Supported by EPSRC Fellowship EP/N008197/1.

1 Introduction


Continuous-time Markov chains (CTMC) have been intensively investigated for a long time, especially because they are simple stochastic models with a wide range of real life applications, being suitable for modelling properties such as expected failure time for systems or expected time between system events. Given the omnipresence of continuous-time Markov chains, it has been natural to seek a logical formalism to describe their properties. A popular example is Continuous Stochastic Logic (CSL), introduced by Aziz et al in [2]. CSL is a branching-time, temporal logic, that allows expressing quantitative bounds on certain properties of continuous-time Markov chains.

Let us consider the CTMC \mathcal{M} in Figure 1 modelling the state transitions of a simple system. One can express the property that the probability of encountering an error in the continuous time interval $[0, 4]$ is greater than 0.5 in CSL by the following state formula:

$$\varphi := \mathbb{P}_{>0.5}(true \mathbf{U}_{[0,4]} s_3). \quad (1)$$

¹ Most of the research was done as part of Andrei’s dissertation while he was an undergraduate student at the University of Oxford working under the supervision of James Worrell.





■ **Figure 1** A simple CTMC modelling a system which is considered to run properly in states s_1 and s_2 , and to malfunction in state s_3 .

Another natural property that one might want to express is whether there exists a “dangerous” short period in $[0, 4]$, say of length 0.1, in which our example system fails with probability at least 0.3. This could then lead to isolating such periods and taking appropriate action. However, CSL does not allow expressing such properties. This is why we extend CSL to allow existential quantifiers over time bounds, giving rise to the logic Parametric CSL (PCSL), in which we can express the desired property by a state formula:

$$\psi := \exists t \in [0, 3.9] \cdot \mathbb{P}_{>0.3}(\text{true} \mathbf{U}_{[t, t+0.1]} s_3). \quad (2)$$

In general, checking if mathematical models satisfy certain properties is a central part of formal verification. This gives rise to *model checking* problems, in which we want to find procedures to determine if a model verifies properties that are usually expressed formally within a logic. In CSL, the model checking problem consists of deciding if properties expressed by state formulas are true or false in certain states of a CTMC. The main result of [2] is that CSL model checking is decidable. The proof is non-trivial, as it employs results in algebraic and transcendental number theory such as the Lindemann-Weierstrass theorem [10]. There exist state-of-the-art model checking software, such as PRISM [8], which allow verifying properties of systems, including CTMCs, expressed formally by logics like CSL, PCTL. However, in this project we deal theoretically with the fundamental problem regarding PCSL model checking.

We define the model checking problem of PCSL similarly to the one of CSL, with the simple exception that we allow state formulas to be evaluated over initial distributions instead of states. Therefore, we want to decide if a CTMC together with an initial distribution entail a PCSL state formula². We show that the model checking problem for the fragment of PCSL consisting of formulae with only one existential quantifier, such as 2 above, is decidable assuming Schanuel’s Conjecture, a conjecture which generalizes important results in transcendental number theory, including the Lindemann-Weierstrass Theorem. The latter was used in [2] to prove CSL model checking decidability. We also discuss why PCSL model checking decidability is non-trivial and employs a strong number theoretical result.

² Note that this simply allows for checking entailment in a certain state by setting its initial probability to 1.

2 PCSL Syntax and Semantics

We extend the original CSL formulation of Aziz et al by allowing existentially quantified parameters. Sticking to the terminology in [2], we call a “path” through a CTMC M a function on domain $[0, \infty)$ with values in the state set S , which associates to each time step a state and follows the transitions in M . For any state s we denote by U^s the set of paths starting at s . Similarly, we denote by U^M the set of paths starting at any state in M . For any set of paths Γ starting at the same state s we denote its probability by $\mu^s(\Gamma)$.

For any initial distribution π and for any set of paths Φ , not necessarily starting from the same state, we denote its probability by $\mu^\pi(\Phi)$, where the probability of the initial vertex is determined according to the initial distribution π of M :

$$\mu^\pi(\Phi) := \sum_{s \in S} \pi(s) \mu^s(U^s \cap \Phi). \quad (3)$$

We now give the syntax and semantics of our extension, which we name “Parametric Continuous Stochastic Logic” (PCSL). For clarity, we use in PCSL state names instead of state labels in the formulas, while the authors of the original CSL papers use state labels. We also define the satisfaction relation for state formulas over initial distributions instead of states, to allow a wider class of verifiable models. Apart from this, CSL can be seen as the PCSL restriction when using no quantified parameters in the definitions below. We also define PCSL_n , for $n \in \mathbb{N}$, to be the restriction of PCSL with at most n nested existential quantifiers.

2.1 PCSL Syntax

First, let $T = \{t_1, t_2, \dots\}$ be a countably infinite set of free variables to which we have access. These variables will represent existentially quantified real numbers. We define a *parametric term* over a finite set of free variables $T' \subset T$ as a linear combination of free variables in T' with rational coefficients:

1. c is a parametric term, for any $c \in \mathbb{Q}$,
2. $\tau + qt$ is a parametric term, for any parametric term τ , $q \in \mathbb{Q}$, and $t \in T'$.

Let M be a CTMC with state set S . As in CSL, there are two types of PCSL formulas: *state formulas* and *path formulas*.

State formulas are evaluated in states, or over initial distributions³, and their syntax is given by:

1. s , for $s \in S$ (the atomic state formula),
2. If f_1 and f_2 are state formulas, then so are $\neg f_1$ and $f_1 \vee f_2$,⁴
3. If g is a path formula using parametric terms over free variables $\{t_1, \dots, t_r\}$, then $\exists t_1 \in [x_1, y_1] \dots \exists t_r \in [x_r, y_r] \cdot \mathbb{P}_{>c}(g)$ is a state formula, where $c \in \mathbb{Q}$, and for $i = 1, \dots, r$: $x_i, y_i \in \mathbb{Q}$, and $0 \leq x_i \leq y_i$.⁵

Path formulas are evaluated along paths, and their syntax is :

³ In CSL, the state formulas are only evaluated in states. Our extension allows evaluation over initial distributions as well.

⁴ $f_1 \wedge f_2$ and $f_1 \rightarrow f_2$ can be written using only these definitions

⁵ Note that we allow as well no free variables, so no quantifiers at all in such a formula. PCSL differs from CSL specifically by allowing these \exists operators.

7:4 Parametric Continuous Stochastic Logic

1. $f_1 \mathbf{U}_{[a_1, b_1]} f_2 \mathbf{U}_{[a_2, b_2]} \dots f_n$, where f_1, f_2, \dots, f_n are state formulas, and all a_1, \dots, a_{n-1} and b_1, \dots, b_{n-1} are parametric terms over a finite set of free variables.

For any $k \in \mathbb{N}$, the syntax of PCSL_k is the same as of PCSL , with the exception of the path formula rule 1, which for PCSL_k is:

1. $f_1 \mathbf{U}_{[a_1, b_1]} f_2 \mathbf{U}_{[a_2, b_2]} \dots f_n$, where f_1, f_2, \dots, f_n are state formulas, and all a_1, \dots, a_{n-1} and b_1, \dots, b_{n-1} are parametric terms over a set of free variables of size at most k .

2.2 PCSL Semantics

Let M be a CTMC, with state set S , and initial distribution π_0 . Let f and g be PCSL state, respectively path formulas.

We say that a state s satisfies a state formula f if, for a distribution π' such that $\pi'(s) = 1$, we have $M, \pi' \models f$ according to the definitions below. Let us denote by $\llbracket f \rrbracket_M$ the set of states satisfying f . We also denote by $g[t \leftarrow d]$ the path formula obtained by substituting the occurrences of the free variable t in parametric terms of g by the non-negative real d . We define the satisfaction relation $M, \pi \models f$ for a general rational distribution π , using structural induction over the state formula f :

1. f is of the form s ($s \in S$): $M, \pi \models f$ iff $\pi(s) = 1$,
2. f is of the form $\neg f_1$: $M, \pi \models f$ iff $M, \pi \not\models f_1$,
3. f is of the form $f_1 \vee f_2$: $M, \pi \models f$ iff $M, \pi \models f_1$ or $M, \pi \models f_2$,
4. f is of the form $\exists t_1 \in [x_1, y_1] \dots \exists t_r \in [x_r, y_r] \cdot \mathbb{P}_{>c}(g)$: $M, \pi \models f$ iff there exist non-negative reals $c_1 \in [x_1, y_1], \dots, c_r \in [x_r, y_r]$ such that

$$\mu^\pi(\{\rho \in U^M \mid M, \rho \models g[t_1 \leftarrow c_1][t_2 \leftarrow c_2] \dots [t_r \leftarrow c_r]\}) > c,$$

By notation abuse, we define the satisfaction relation $M, \rho \models g$ for path formulas g and for any path ρ :

1. g is a path formula with no free variables (i.e. all parametric terms are numbers) of the form $f_1 \mathbf{U}_{[a_1, b_1]} f_2 \mathbf{U}_{[a_2, b_2]} \dots f_n$ and ρ is a path through M : $M, \rho \models g$ iff there exist positive reals $\alpha_1, \dots, \alpha_{n-1}$ such that for each integer in $[1, n-1]$ we have $a_i \leq \alpha_i \leq b_i$ and for any $\beta \in [\alpha_{i-1}, \alpha_i]$ we have $\pi(\beta) \in \llbracket f_i \rrbracket_M$, and $\pi(\alpha_{n-1}) \in \llbracket f_n \rrbracket_M$.⁶

We further overwrite the satisfaction relation as follows:

- we define $M \models f$ iff $M, \pi_0 \models f$, where π_0 is the initial distribution of M ,
- for any $s \in S$, we define $M, s \models f$ iff $M, \pi' \models f$, where distribution π' is chosen over S such that $\pi'(s) = 1$ and $\pi'(s') = 0$, for any $s' \neq s$.

2.3 PCSL Formulas Examples

The PCSL formula

$$\phi_3 := s_1 \wedge \exists t \in [0, 5] \cdot \mathbb{P}_{>0.5}(\text{true} \mathbf{U}_{[t, t]} s_2)$$

expresses the property that the system is initially in state s_1 and there exists an instantaneous moment $t \leq 5$ during which the probability of being in state s_2 is greater than 0.5. The formula ϕ_3 is in PCSL_1 , but not in CSL . Note that it is different from the CSL formula $\phi_4 := s_1 \wedge \mathbb{P}_{>0.5}(\text{true} \mathbf{U}_{[0, 5]} s_2)$, which expresses the property of being in state s_1 initially and transitioning to s_2 at any moment before 5.0 with probability greater than 0.5.

⁶ The real number α_0 is defined to be 0 for convenience. There are other ways to define the semantics for the path formula, but we want to be consistent with [2].

A more interesting example is motivated by the following situation. Suppose we have a system and we want a state formula φ to hold with high probability (> 0.8) before time $t = 5$, but we do not want the state formula to be too biased towards any short period, i.e., we do not want there to be any continuous time period of length 0.1 such that φ holds with probability greater than 0.2. This can be modelled in PCSL (in PCSL₁) as:

$$\phi_5 := \mathbb{P}_{>0.8}(\text{true}\mathbf{U}_{[0,5]}\varphi) \wedge \neg\exists t \in [0, 4.9] \cdot \mathbb{P}_{>0.2}(\text{true}\mathbf{U}_{[t,t+0.1]}\varphi).$$

The following formula is in PCSL₂, but (syntactically) not in PCSL₁, as it contains a path formula with two free variables:

$$\phi_6 := \exists t_1 \in [0, 1] \exists t_2 \in [3, 4] \cdot \mathbb{P}_{>0.5}(\text{true}\mathbf{U}_{[t_1,t_1]}s_1 \mathbf{U}_{[t_1,t_2]}\text{true}\mathbf{U}_{[t_2,t_2]}s_2).$$

Formula ϕ_6 expresses the property that there are some moments $t_1 \in [0, 1]$ and $t_2 \in [3, 4]$ such that the probability of being in state s_1 at time t_1 and in state s_2 at time t_2 is greater than 0.5.

3 Mathematical Background

3.1 Exponential Polynomials

► **Definition 1.** *An exponential polynomial is a function $f(t) = \sum_{i=1}^m P_i(t)e^{\alpha_i t}$, where $P_i \in \mathbb{C}[t]$ are polynomials with complex coefficients and $\alpha_1, \dots, \alpha_m$ are complex numbers. We call the coefficients of polynomials P_1, \dots, P_m and the numbers $\alpha_1, \dots, \alpha_m$ the coefficients of the exponential polynomial f .*

Exponential polynomials often arise when writing the explicit solutions of ordinary linear differential equations and when modelling probability distributions in dynamic systems, such as continuous-time Markov chains [4, 2, 3]. We will mainly be concerned with exponential polynomials with algebraic coefficients that are real-valued over reals, i.e., if t is real, then $f(t)$ is real. We simply refer to such functions as real-valued.

The following result, a standard linear algebra result (see [3, 2] for a detailed proof), will later on give the relation between transition probabilities of continuous-time Markov chains and exponential polynomials:

► **Lemma 2.** *Let A be an $n \times n$ matrix with rational (algebraic) entries. Then, for any $\alpha, \beta \in \mathbb{Q}$, the entries of the exponential matrix $f(t) := \exp(A(\alpha t + \beta))$ are real-valued exponential polynomials with algebraic coefficients.*

3.2 Schanuel's Conjecture

Schanuel's Conjecture is a unifying conjecture in the field of transcendental number theory, having as consequences important results about exponential functions over both real and complex numbers, such as in the work of Zilber [12], and in model theory, such as decidability of the first-order theory $Th_{\text{exp}}(\mathbb{R})$ of the field of real numbers with exponentials $\langle \mathbb{R}, +, \times, \exp, 0, 1 \rangle$ [9], and decidability of the Continuous Skolem Problem [4].

Schanuel's Conjecture has the following form:

► **Conjecture 1.** *(Schanuel's Conjecture) Given any n complex numbers z_1, \dots, z_n that are linearly independent over \mathbb{Q} , the extension field $\mathbb{Q}(z_1, \dots, z_n, e^{z_1}, \dots, e^{z_n})$ has transcendence degree at least n over \mathbb{Q} .*

Schanuel's Conjecture states that, for z_i 's as above, among $z_1, \dots, z_n, e^{z_1}, \dots, e^{z_n}$ there are at least n numbers which are not related by any non-trivial polynomial with rational coefficients.

Schanuel's Conjecture is a generalisation of Lindemann-Weierstrass theorem, which lies at the heart of the decidability proof of CSL in [2], the logic that we extend into PCSL.

3.3 The Positivity Problem

► **Definition 3.** *An instance of the Positivity Problem for exponential polynomials is a real-valued exponential polynomial $f(t) = \sum_{j=1}^m P_j(t)e^{\lambda_j t}$ with algebraic coefficients, together with an interval $[c, d]$, where $c, d \in \mathbb{Q}_+$ ($d \geq c \geq 0$). We want to answer the question: does there exist $t \in [c, d]$ such that $f(t) > 0$?*

We show in the appendix that deciding whether a real-valued exponential polynomial with algebraic coefficients is strictly positive at some point in a given bounded interval with rational endpoints is decidable under Schanuel's Conjecture. This decision problem is of particular interest because it arises naturally in continuous linear dynamical systems, as we will see in our CSL extension. We mainly build our proof on top of the one in [4], which shows that we can decide, subject to Schanuel's Conjecture, whether a real-valued exponential polynomial with algebraic coefficients has a zero in a given bounded interval. The additional complexity in the current problem comes from the fact that detecting sign changes for a real-valued exponential polynomial is based on the behaviour of all its factors together, unlike detecting roots, where only the behaviour of one of its factors matters.

► **Theorem 4.** *The Positivity Problem for exponential polynomials is decidable assuming Schanuel's Conjecture.*

The following is a proof outline of Theorem 4; full details can be found in the Appendix. Suppose that we want to decide whether a given real-valued exponential polynomial f is positive throughout an interval $[c, d]$. We reduce this problem to deciding the existence of zeros of exponential polynomials on bounded intervals, which is known to be decidable conditional on Schanuel's Conjecture [4]. In fact the reduction itself uses several of the ideas developed in [4]. To carry out the reduction we first compute the sign of f at the endpoints c and d . Suppose that f is negative at both endpoints. We then compute a factorisation of f in the form $f = f_1^{\alpha_1} \cdots f_k^{\alpha_k}$, where the factors f_i are real-valued exponential polynomials that do not share any common zeros. Then determining whether f changes sign from negative to positive on $[c, d]$ reduces to determining whether one of its factors f_i with odd exponent α_i changes sign. To solve this last problem we give an effectively decidable categorisation of the factors into two types.

We show that factors of the first type are always nonnegative and factors g of second type are such that g and g' have no common zeros, i.e., they always sign at every zero. Thus f becomes positive on $[c, d]$ iff it has a factor of the second type that has a zero in $[c, d]$. The role of Schanuel's conjecture in the above argument is to rule out the existence of common zeros of the different factors of f and common zeros of certain factors and their derivatives.

► **Remark 5.** Given a function f and an interval $[c, d]$ that are an instance of the Positivity Problem for exponential polynomials, a decision procedure for this problem trivially implies a decision procedure for checking in a similar setup if there exists $t \in [c, d]$ such that $f(t) > q$, for any given rational q . This follows as we can let $g(t) := f(t) - q$, so g is an exponential polynomial with algebraic coefficients as well, therefore we can use a decision procedure for the Positivity Problem for exponential polynomials on input function g and interval $[c, d]$ and decide if there exists $t \in [c, d]$ such that $f(t) > q$.

The Positivity Problem for exponential polynomials is a hard problem, as it is trivially inter-reducible with the Non-negativity Problem for exponential polynomials [3], which has the same setup as the Positivity Problem for exponential polynomials, but asks whether for all $t \in [c, d]$ it is true that $f(t) \geq 0$. Concretely, decidability of any of the two problems

implies decidability of the other as well:

$$\forall t \in [c, d]. f(t) \geq 0 \Leftrightarrow \neg \exists t \in [c, d]. (-f(t) > 0), \quad (4)$$

$$\exists t \in [c, d]. f(t) > 0 \Leftrightarrow \neg \forall t \in [c, d]. (-f(t) \geq 0). \quad (5)$$

However, decidability of the Non-negativity Problem for exponential polynomials is open [3], so decidability of the Positivity Problem for exponential polynomials is a hard mathematical task. In fact, comparing exponential polynomials with 0 is a hard task [4, 9, 3], and decision problems related to this would have considerable new implications in both model theory and number theory [4]. This motivates us to work under the assumption of Schanuel's Conjecture, which is often assumed in model theory [4, 12], as the unconditional decidability currently seems out of reach.

4 Model Checking Decidability of PCSL

A central problem in formal verification for any logic describing dynamic systems is the *model checking problem*. Intuitively, it asks whether a model of a certain system satisfies a specification, usually expressed withing a logical formalism.

We introduce the PCSL model checking problem below.

► **Definition 6** (Model checking problem for PCSL). *An instance of the model checking problem for PCSL is given by a continuous-time Markov chain M , a distribution π with rational entries over the states of M , and a PCSL formula φ . We want to answer the question: is it the case that $M, \pi \models \varphi$?*

The model checking problem for PCSL_n , for any $n \in \mathbb{N}$, is defined similarly, with the exception that φ is a PCSL_n formula in Definition 6. We prove in subsection 4.1 that PCSL_1 model checking is decidable assuming Schanuel's Conjecture. We also show in subsection 4.2 that unconditional PCSL_1 model checking is hard from a mathematical point of view, by reducing a well-known hard problem to it.

4.1 Decidability of the model checking problem for PCSL_1 assuming Schanuel's Conjecture

We show that PCSL_1 model checking is decidable assuming Schanuel's Conjecture. For this, we prove that the decidability of the Positivity Problem for exponential polynomials implies PCSL_1 model checking decidability. As discussed in Section 3.3, Schanuel's Conjecture implies decidability of the Positivity Problem for exponential polynomials, therefore we get our result.

Given any Markov chain M with state set $S = \{s_1, \dots, s_k\}$ and rational transition rate matrix Q , and an initial rational distribution π , we proceed by structural induction over PCSL_1 formula φ to show that there exists a model checking procedure to determine if $M, \pi \models \varphi$.

Let us first deal with the trivial cases.

If φ is an atom (state) s , then $M, \pi \models \varphi$ iff $\pi(s) = 1$.

If $\varphi = \varphi_1 \vee \varphi_2$, we have $M, \pi \models \varphi$ iff $M, \pi \models \varphi_1$ or $M, \pi \models \varphi_2$.

If $\varphi = \neg \varphi_1$, we have $M, \pi \models \varphi$ iff $M, \pi \not\models \varphi_1$.

If $\varphi = \mathbb{P}_{>c}(\varphi_1 \mathbf{U}_{[a_1, b_1]} \varphi_2 \mathbf{U}_{[a_2, b_2]} \dots \mathbf{U}_{[a_{n-1}, b_{n-1}]} \varphi_n)$, as we can model check formulas $\varphi_1, \dots, \varphi_n$ by induction, the decidability follows from the same proof used by Aziz et al in [2] to show that classic CSL is decidable, by only using the Lindemann-Weierstrass theorem.

7:8 Parametric Continuous Stochastic Logic

Now, we have to deal with the case

$$\varphi = \exists t \in [a, b] \cdot \mathbb{P}_{>c}(\varphi_1 \mathbf{U}_{[a_1, b_1]} \varphi_2 \mathbf{U}_{[a_2, b_2]} \cdots \mathbf{U}_{[a_{n-1}, b_{n-1}]} \varphi_n),$$

where $a, b, c \in \mathbb{Q}$, and $a_1, b_1, \dots, a_{n-1}, b_{n-1}$ are parametric terms over $\{t\}$ (functions in t of the form $\alpha t + \beta$, with $\alpha, \beta \in \mathbb{Q}$). We therefore need to reason about the quantity

$$f(t) := \mu^\pi(\{\text{paths } \rho \in U^M \mid M, \rho \models \varphi_1 \mathbf{U}_{[a_1, b_1]} \varphi_2 \mathbf{U}_{[a_2, b_2]} \cdots \mathbf{U}_{[a_{n-1}, b_{n-1}]} \varphi_n\}). \quad (6)$$

In fact, we are interested if there exists some $t \in [a, b]$ such that $f(t) > c$. We further show that $f(t)$ is an exponential polynomial that we can algorithmically compute, which gives us the sought conditional decidability result.

Assume for the moment that for any $t \in [a, b]$ we have

$$0 \leq a_1 \leq b_1 \leq \cdots \leq a_{n-1} \leq b_{n-1}. \quad (7)$$

By the structural induction hypothesis, we can compute the sets of states $\llbracket \varphi_1 \rrbracket_M, \dots, \llbracket \varphi_n \rrbracket_M$ satisfying subformulas $\varphi_1, \dots, \varphi_n$. For any subset of states $H \subseteq S$, let its complement be $H^c := S \setminus H$.

We show how to compute the probability function $f(t)$, by similar constructions to the ones in [2]. Let us construct the following matrices, where for any matrix A we refer to its entry on row i and column j as $A(i, j)$.

- Let $Q_{i,i}$ be a transition rate matrix that models states in $\llbracket \varphi_i \rrbracket_M^c$ as absorbing states, and is everywhere else identical to Q :

$$Q_{i,i}(j, k) := \begin{cases} Q(j, k), & \text{if } s_j \in \llbracket \varphi_i \rrbracket_M, \\ 0, & \text{if } s_j \in \llbracket \varphi_i \rrbracket_M^c. \end{cases}$$

This matrix is used to model a run of M which remains in states satisfying φ_i . Also, let $P_{i,i}(t) := \exp(Q_{i,i}t)$ be the transition matrix for time t corresponding to the Markov chain described by $Q_{i,i}$.

- Let $Q_{i,i+1}$ be a transition rate matrix obtained from Q that only models transitions from $\llbracket \varphi_i \rrbracket_M$ to $\llbracket \varphi_i \rrbracket_M \cup \llbracket \varphi_{i+1} \rrbracket_M$, and from $\llbracket \varphi_{i+1} \rrbracket_M$ to $\llbracket \varphi_{i+1} \rrbracket_M$:

$$Q_{i,i+1}(j, k) := \begin{cases} Q(j, k), & \text{if } s_j \in \llbracket \varphi_i \rrbracket_M \text{ and } s_k \in \llbracket \varphi_i \rrbracket_M \cup \llbracket \varphi_{i+1} \rrbracket_M, \\ Q(j, k), & \text{if } s_j \in \llbracket \varphi_{i+1} \rrbracket_M \text{ and } s_k \in \llbracket \varphi_{i+1} \rrbracket_M, \\ 0, & \text{otherwise.} \end{cases}$$

This matrix is used to model transitions from states satisfying φ_i to states satisfying φ_{i+1} . Also, let $P_{i,i+1}(t) := \exp(Q_{i,i+1}t)$ be the transition matrix for time t corresponding to the Markov chain described by $Q_{i,i+1}$.

- Let I_i be an indicator matrix of states in $\llbracket \varphi_i \rrbracket_M$:

$$I_i(j, k) := \begin{cases} 1, & \text{if } s_j = s_k \in \llbracket \varphi_i \rrbracket_M, \\ 0, & \text{otherwise.} \end{cases}$$

This matrix is used to filter out states not satisfying φ_i at certain times.

- Finally, let E_n be a matrix obtained from Q which treats states in $\llbracket \varphi_n \rrbracket_M$ as absorbing states:

$$E_n(j, k) := \begin{cases} 0, & \text{if } s_j \in \llbracket \varphi_n \rrbracket_M, \\ Q(j, k), & \text{otherwise.} \end{cases}$$

This matrix is used at the end of the formula to “collect” all the probability mass of paths which have satisfied the path formula $\varphi_1 \mathbf{U}_{[a_1, b_1]} \varphi_2 \mathbf{U}_{[a_2, b_2]} \cdots \mathbf{U}_{[a_{n-1}, b_{n-1}]} \varphi_n$. Also, let $F_n(t) := \exp(E_n t)$ be the transition matrix for time t corresponding to the Markov chain described by E_n .

It is not hard to see that the probability of paths starting in M according to the initial probability π which satisfy the path formula

$$\varphi_1 \mathbf{U}_{[a_1, b_1]} \varphi_2 \mathbf{U}_{[a_2, b_2]} \cdots \mathbf{U}_{[a_{n-1}, b_{n-1}]} \varphi_n,$$

as defined in (6) above, has the expression:

$$f(t) = \pi^\top \cdot P_{0,0}(a_1) \cdot I_0 \cdot P_{0,1}(b_1 - a_1) \cdot I_1 \cdot P_{1,1}(a_2 - b_1) \cdot I_1 \cdot P_{1,2}(b_2 - a_2) \cdot I_2 \cdots I_{n-1} \cdot F_n(b_n - a_n) \cdot I_n \cdot \mathbf{1}. \quad (8)$$

All matrix functions parameters $(a_1, b_1 - a_1, a_2 - b_1, \dots)$ are of the form $\alpha t + \beta$, for $\alpha, \beta \in \mathbb{Q}$. By Lemma 2 we get that all entries implied in the product at (8) are exponential polynomials with algebraic coefficients. As exponential polynomials with algebraic coefficients are closed under product and sum, we get that $f(t)$ is an exponential polynomial with algebraic coefficients, which we can compute algorithmically, by using classic representation methods of algebraic numbers (see [5, Section 4.2]).

Therefore, by our result - Theorem 4, we get that assuming Schanuel’s Conjecture we can decide if there exists $t \in [a, b]$ such that $f(t) - c > 0$, as $f(t) - c$ is a real-valued exponential polynomial with algebraic coefficients. Therefore, under Schanuel’s Conjecture, we can also model check φ in the case when

$$\varphi = \exists t \in [a, b] \cdot \mathbb{P}_{>c}(\varphi_1 \mathbf{U}_{[a_1, b_1]} \varphi_2 \mathbf{U}_{[a_2, b_2]} \cdots \mathbf{U}_{[a_{n-1}, b_{n-1}]} \varphi_n).$$

In conclusion, we have covered all forming rules of state formulas in PCSL_1 , and proved by structural induction that Schanuel’s Conjecture implies the existence of a model checking procedure for PCSL_1 .

► **Remark 7.** Let us briefly discuss the assumption (7). We assumed that the parametric terms $a_1, b_1, \dots, a_{n-1}, b_{n-1}$ in $\{t\}$, which are linear functions in t , satisfy for all $t \in [a, b]$: $a_1 \leq b_1 \leq \dots \leq a_{n-1} \leq b_{n-1}$. Let us write them explicitly as $a_i = x_i(t), b_i = y_i(t)$, for $i = 1, \dots, n - 1$. First, it is easy to see that, as all parametric terms are linear functions in t with rational coefficients, there is some maximal interval $[c, d] \subseteq [a, b]$, with $t, c \in \mathbb{Q}$, such that all $0 \leq x_1(t) \leq y_1(t), 0 \leq x_2(t) \leq y_2(t), \dots, 0 \leq x_{n-1}(t) \leq y_{n-1}(t)$ hold for all $t \in [c, d]$, and at least one of them doesn’t hold for any $t \in [a, b] \setminus [c, d]$. Then, we can just seek some value of t in $[c, d]$ that satisfies the formula, as outside this interval the formula is not syntactically valid. Now, in order to be able to also assume the inequalities $y_i(t) \leq x_{i+1}(t)$, we can split the interval $[c, d]$ in intervals in which either $y_i(t) \leq x_{i+1}(t)$, or $y_i(t) \geq x_{i+1}(t)$ holds, and deal with all possible cases separately. The full details for this part are rather technical, and mostly follow the technique in [2].

4.2 Hardness of the model checking problem for PCSL_1

To show hardness of deciding the model checking problem for PCSL_1 , therefore showing hardness of deciding the model checking problem for PCSL implicitly, we proceed in two steps.

First, we introduce a hard decision problem - *the Positivity Problem for matrix exponentials*. This is a hard problem as it is inter-reducible with the Positivity Problem for exponential polynomials (by simple algebraic manipulation, see [3] for details). We have discussed in

subsection 3.3 why the Positivity Problem for exponential polynomials is a hard problem: it would imply decidability of the Non-negativity Problem for exponential polynomials, which is currently open [3]. We show that the Positivity Problem for matrix exponentials is reducible to a decision problem regarding CTMC properties - *the Threshold Problem for continuous-time Markov chains*.

Second, we show that PCSL₁ model checking decidability implies decidability of the Threshold Problem for continuous-time Markov chains. Therefore, decidability of the model checking problem for PCSL₁ implies decidability of the Positivity Problem for matrix exponentials. This stands as hardness evidence for a PCSL₁ model checking decision procedure, and for a PCSL model checking decision procedure as well, because PCSL₁ is a fragment of PCSL.

4.2.1 Reduction of a Hard Problem to the Threshold Problem for Continuous-Time Markov Chains

We introduce below the Threshold Problem for continuous-time Markov chains and the Positivity Problem for matrix exponentials.

► **Definition 8** (Threshold Problem for continuous-time Markov chains). $I = (\langle \mathbf{u}, \mathbf{R}, \mathbf{v} \rangle, \langle a, b \rangle)$ is an instance of the **Threshold Problem for continuous-time Markov chains**, where $\mathbf{u} \in \mathbb{Q}^k$ is a stochastic vector⁷, $\mathbf{v} \in \{0, 1\}^k$, $\mathbf{R} \in \mathbb{Q}^{k \times k}$ is a rate matrix (for some $k \in \mathbb{N}$), and $a, b \in \mathbb{Q}$ such that $0 \leq a \leq b$. We want to answer the question: does there exist some real $t \in [a, b]$ such that $\mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v} > \frac{1}{2}$?

Intuitively, in the Threshold Problem for continuous-time Markov chains, \mathbf{u} represents the initial distribution and \mathbf{R} represents the rate matrix of a CTMC. Then, we ask if at some moment during a given interval $[a, b]$ the probability of being in a state from a given set, that is described by 1-entries of \mathbf{v} , is greater than $\frac{1}{2}$.

► **Definition 9** (Positivity Problem for matrix exponentials). $I = (\langle \mathbf{u}, \mathbf{A}, \mathbf{v} \rangle, \langle a, b \rangle)$ is an instance of the **Positivity Problem for matrix exponentials**, where $\mathbf{u}, \mathbf{v} \in \mathbb{Q}^k$, $\mathbf{A} \in \mathbb{Q}^{k \times k}$ (for some $k \in \mathbb{N}$), and $a, b \in \mathbb{Q}^+$, with $0 \leq a \leq b$. We want to answer the question: does there exist some real $t \in [a, b]$ such that $\mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} > 0$?

Note that the Positivity Problem for matrix exponentials can be seen as a generalization of the Threshold Problem for continuous-time Markov chains, both because the former has much more general instances, but also because its decidability implies decidability of the latter. To see this, let $(\langle \mathbf{u}, \mathbf{R}, \mathbf{v} \rangle, \langle a, b \rangle)$ be an instance of the Threshold Problem for continuous-time Markov chains, then $e^{\mathbf{R}t}$ is a stochastic matrix⁸, so $\mathbf{u}^\top e^{\mathbf{R}t}$ is a stochastic row vector, therefore $\mathbf{u}^\top e^{\mathbf{R}t} \mathbf{1} = 1$. Then, we could obtain a decision procedure for this problem by applying a decision procedure for the Positivity Problem for matrix exponentials on instance $(\langle \mathbf{u}, \mathbf{R}, \mathbf{v} - \frac{1}{2} \mathbf{1} \rangle, \langle a, b \rangle)$:

$$\begin{aligned} \exists t \in [a, b] \text{ such that } \mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v} > \frac{1}{2} &\Leftrightarrow \\ \exists t \in [a, b] \text{ such that } \mathbf{u}^\top e^{\mathbf{R}t} (\mathbf{v} - \frac{1}{2} \mathbf{1}) > 0. \end{aligned}$$

⁷ Has positive entries that sum up to 1.

⁸ Its rows are probability distributions.

Furthermore, as $e^{\mathbf{R}t}$ is a stochastic matrix, we would expect the Threshold Problem for continuous-time Markov chains to be considerably easier, as eigenvalues of stochastic matrices are well-behaved⁹.

Surprisingly, we show that the Positivity Problem for matrix exponentials is reducible to the Threshold Problem for continuous-time Markov chains, thus making the two decision problems equivalently hard.

► **Theorem 10.** *The Positivity Problem for matrix exponentials is reducible to the Threshold Problem for continuous-time Markov chains.*

Proof. The full proof is given in the appendix. Using algebraic manipulations, we construct a rate transition matrix \mathbf{O} and some vectors \mathbf{u}_1 and \mathbf{v}_3 , and then a rate transition matrix \mathbf{R} , a stochastic vector $\tilde{\mathbf{u}}$ and a vector $\tilde{\mathbf{v}}$ with only 0 and 1 entries such that

$$\begin{aligned} \exists t \in [a, b] \text{ such that } \mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} > 0 &\iff \\ \exists t \in [a, b] \text{ such that } \mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_3 > \frac{1}{2} &\iff \\ \exists t \in [a, b] \text{ such that } \tilde{\mathbf{u}}^\top e^{\mathbf{R}t} \tilde{\mathbf{v}} > \frac{1}{2}. & \end{aligned} \quad (9)$$

However, a decision procedure for the Threshold Problem for continuous-time Markov chains would specifically allow us to answer queries such as $\exists t \in [a, b]$ such that $\tilde{\mathbf{u}}^\top e^{\mathbf{R}t} \tilde{\mathbf{v}} > \frac{1}{2}$, therefore it would give a decision procedure for the Positivity Problem for matrix exponentials as well. ◀

4.2.2 Expressing the Threshold Problem for Continuous-Time Markov Chains in PCSL₁

Let $I = (\langle \mathbf{u}, \mathbf{R}, \mathbf{v} \rangle, \langle a, b \rangle)$ be an instance of the Threshold Problem for continuous-time Markov chains. Recall that $\mathbf{u} \in \mathbb{Q}^k$ is a stochastic vector, $\mathbf{R} \in \mathbb{Q}^{k \times k}$ is a rate matrix, $\mathbf{v} \in \{0, 1\}^k$, and $0 \leq a \leq b$ are rationals, and we want to answer whether there exists $t \in [a, b]$ such that $\mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v} > \frac{1}{2}$.

Let M be the continuous-time Markov chain corresponding to rate matrix \mathbf{R} , with initial probability distribution $\pi_0 := \mathbf{u}$, and with state set S such that $|S| = k$. The probability distribution over states at time t is given by $\mathbf{u}^\top e^{\mathbf{R}t}$. Therefore, as $\mathbf{v} \in \{0, 1\}^k$, we can see the expression $\mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v}$ as summing up the probability distribution at time t of states corresponding to 1-entries in \mathbf{v} .

Let the states from S that correspond to 1-entries of \mathbf{v} be $S' = \{s_1, \dots, s_i\}$. If S' is empty, then $\mathbf{v} = \mathbf{0}$, and we have $\mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v} = 0$, so the Threshold Problem for continuous-time Markov chains instance is a negative instance, and we trivially have:

$$\exists t \in [a, b] \text{ such that } \mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v} > \frac{1}{2} \iff M \models s \wedge \neg s, \text{ for some } s \in S. \quad (10)$$

Otherwise, if S' is not empty, we get:

$$\begin{aligned} \exists t \in [a, b] \text{ such that } \mathbf{u}^\top e^{\mathbf{R}t} \mathbf{v} > \frac{1}{2} &\iff \\ M, \pi_0 \models \exists t \in [a, b] \cdot \mathbb{P}_{>1/2}(\text{true} \mathbf{U}_{[t,t]}(s_1 \vee \dots \vee s_i)). & \end{aligned} \quad (11)$$

⁹ Standard linear algebra results imply that 1 is always an eigenvalue of any stochastic matrix, and all the eigenvalues have absolute value less or equal to 1.

Thus, for any instance I of the Threshold Problem for continuous-time Markov chains, there exists a continuous-time Markov chain M and some PCSL₁ formula φ such that I is a yes-instance of the Threshold Problem for continuous-time Markov chains if and only if the PCSL₁ satisfaction relation $M \models \varphi$ holds.

In conclusion, a PCSL₁ model checking procedure (that decides if PCSL₁ statements of the form $M \models \varphi$ hold) would yield the existence of a decision procedure for the Threshold Problem for continuous-time Markov chains. As we have shown in Section 4.2.1, this would imply the decidability of the Positivity Problem for matrix exponentials, which, as discussed, is a hard problem and is not currently known to be decidable. Therefore, the unconditional decidability of PCSL₁ model checking, and thus of PCSL model checking, seems to be a hard problem.

5 Conclusions

5.1 Overview

We introduced PCSL, a powerful parametric logic for formally expressing temporal properties of continuous-time Markov chains. We investigated the model checking problem of our logic, proving that its unconditional decidability is a hard problem, and showed that Schanuel's Conjecture implies decidability of the model checking problem for an expressive fragment of PCSL. The last result relies on a technical proof that Schanuel's Conjecture implies the decidability of the Positivity Problem for exponential polynomials, which is an important achievement in the field.

The logic could have simply been extended to allow operators of the form $\text{Pr}_{\&c}$, where $\&$ could be any of \leq , \geq , $<$, $>$, $=$, or \neq , instead of only allowing $\text{Pr}_{>c}$. All our results would still hold, as [4] proves the conditional decidability of verifying whether exponential polynomials are equal to a given constant in some given interval, and this together with our proofs would suffice for obtaining the same consequences about model checking PCSL. We restricted our attention to PCSL using only operators of the form $\text{Pr}_{>c}$, which makes our arguments more concise, while presenting all the fundamental mathematical problems we have tackled.

5.2 Future Work

We propose two main directions for future work on our project.

5.2.1 General Conditional Decidability of PCSL

We have shown decidability of PCSL₁ model checking assuming Schanuel's Conjecture, by proving conditional decidability of the Positivity Problem for exponential polynomials. In general, the decidability of model checking PCSL _{n} reduces to the decidability of the Positivity Problem for exponential polynomials in n variables. In fact, we found out that using the polynomial resultant for eliminating variables in the two variable case reduces the decidability of the model checking problem for PCSL₂ to a purely algebraic problem. We believe that decidability of model checking PCSL _{n} also follows assuming Schanuel's Conjecture, and therefore we propose seeking a general proof for conditional decidability of model checking PCSL.

5.2.2 Practical Model Checking of PCSL

We have mainly been concerned with the fundamental problem of PCSL decidability, however in practice we expect that the malicious cases we encountered theoretically should not represent too much of a risk in real-life applications. As we have seen interesting classes of properties that are expressible in PCSL, it is worthy to further investigate the practical aspects of model checking PCSL and possible optimizations for an actual procedure.

References

- 1 S Akshay, Timos Antonopoulos, Joël Ouaknine, and James Worrell. Reachability problems for Markov chains. *Information Processing Letters*, 115(2):155–158, 2015.
- 2 Adnan Aziz, Kumud Sanwal, Vigyan Singhal, and Robert Brayton. Model-checking continuous-time Markov chains. *ACM Transactions on Computational Logic (TOCL)*, 1(1):162–170, 2000.
- 3 Paul Bell, Jean-Charles Delvenne, Raphael Jungers, and Vincent D. Blondel. The continuous Skolem-Pisot problem: On the complexity of reachability for linear ordinary differential equations. *Theoretical Computer Science*, 411(40–42):3625–3634, 2010.
- 4 Ventsislav Chonev, Joel Ouaknine, and James Worrell. On the Skolem problem for continuous linear dynamical systems. *43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 100:1–100:13, 2016.
- 5 Henri Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1993.
- 6 Paul M. Cohn. *Basic Algebra: Groups, Rings and Fields*. Springer, second edition, 2004.
- 7 Bettina Just. Integer relations among algebraic numbers. In *International Symposium on Mathematical Foundations of Computer Science*, pages 314–320. Springer, 1989.
- 8 Marta Kwiatkowska, Gethin Norman, and David Parker. PRISM 4.0: Verification of probabilistic real-time systems. *Proc. 23rd International Conference on Computer Aided Verification (CAV’11), volume 6806 of LNCS, pages 585–591, Springer, 2011.*
- 9 Angus Macintyre and Alex J Wilkie. On the decidability of the real exponential field. In (ed. Piergiorgio Odifreddi) *Kreiseliana: About and Around Georg Kreisel.*, 1996.
- 10 Ivan Niven. *Irrational Numbers*. The Mathematical Association of America, fifth edition, 2005.
- 11 Paul S Wang. Factoring multivariate polynomials over algebraic number fields. *Mathematics of Computation*, 30(134):324–336, 1976.
- 12 Boris Zilber. Exponential sums equations and the Schanuel conjecture. *Journal of the London Mathematical Society*, 65(1):27–44, 2002.

A Proof of Theorem 4

Let $f(t) = \sum_{j=1}^m P_j(t)e^{\lambda_j t}$, together with the interval $[c, d]$ be an instance of the Positivity Problem for exponential polynomials. Let \mathbb{K} be the number field generated by the coefficients of polynomials P_1, \dots, P_m and by $\lambda_1, \dots, \lambda_m$ over \mathbb{Q} . We can algorithmically determine a basis $\{a_1, \dots, a_r\}$ over \mathbb{Q} of the real parts of λ_i 's, and a basis $\{b_1, \dots, b_s\}$ over \mathbb{Q} of the imaginary parts of λ_i 's [7].

Without loss of generality, assume that all real and imaginary parts of $\lambda_1, \dots, \lambda_m$ can be written as linear combinations of $\{a_1, \dots, a_r\}$, respectively of $\{b_1, \dots, b_s\}$ that use integer coefficients instead of rational coefficients (this follows as we can pick a suitable $N \in \mathbb{N}$ and write $f_1(t) := f(Nt) = \sum_{j=1}^m P_j(Nt)e^{(\lambda_j N)t}$).

It follows that we can write $f(t)$ as a polynomial in the field of Laurent polynomials \mathcal{R} , with multiplicative units the non-zero monomials in $y_1, \dots, y_r, z_1, \dots, z_s$:

$$\mathcal{R} := \mathbb{K}[x, y_1, y_1^{-1}, \dots, y_r, y_r^{-1}, z_1, z_1^{-1}, \dots, z_s, z_s^{-1}].$$

We write $f(t) = P(t, e^{a_1 t} \dots, e^{a_r t}, e^{b_1 t} \dots, e^{b_s t})$, where P is a polynomial with non-negative power in its first argument, and with any integer power in the others.

Being a localisation of the polynomial ring $\mathbb{A} := \mathbb{K}[x, y_1, \dots, y_r, z_1, \dots, z_s]$, \mathcal{R} is a unique factorisation domain (this is a standard result; see [6, Theorem 10.3.7]) and has an effective procedure for factoring it into irreducible polynomials [11]. We extend the conjugation over \mathcal{R} , by defining a ring automorphism $(\cdot)^*$ which acts on P to yield P^* as below:

$$\begin{aligned} P(x, y_1, \dots, y_r, z_1, \dots, z_s) &= \sum_i \alpha_i x^{\beta_i} y_1^{\gamma_{1,i}} \dots y_r^{\gamma_{r,i}} z_1^{\delta_{1,i}} \dots z_s^{\delta_{s,i}} \\ P^*(x, y_1, \dots, y_r, z_1, \dots, z_s) &:= \sum_i \overline{\alpha_i} x^{\beta_i} y_1^{\gamma_{1,i}} \dots y_r^{\gamma_{r,i}} z_1^{-\delta_{1,i}} \dots z_s^{-\delta_{s,i}}. \end{aligned} \quad (12)$$

The motivation behind this definition is that for $f(t) = P(t, e^{a_1 t} \dots, e^{a_r t}, e^{b_1 t} \dots, e^{b_s t})$ we have $f(\bar{t}) = P^*(t, e^{a_1 t} \dots, e^{a_r t}, e^{b_1 t} \dots, e^{b_s t})$. For such a real-valued f , we have $P = P^*$.

As $(\cdot)^*$ is a ring automorphism over the unique factorization domain \mathcal{R} , we get that if a polynomial Q in \mathcal{R} divides P , then there exists some R in \mathcal{R} such that $P = QR$, so $P^* = Q^*R^*$. Therefore, Q^* also divides P^* , but as $P = P^*$ we have that Q^* divides P . Therefore, factors of P come in $*$ -conjugated pairs.

We will use Schanuel's Conjecture through the following result, which follows from it by using the concept of *resultant* of two polynomials and basic algebraic manipulations [4]:

► **Lemma 11.** *Let r, s be non-negative integers, and let $\{a_1, \dots, a_r\}$ and $\{b_1, \dots, b_s\}$ be \mathbb{Q} -linearly independent sets of algebraic numbers. Let $P, Q \in \mathcal{R}$ be polynomials with algebraic coefficients that are coprime in \mathcal{R} . Then the following equations have no common solution $t \in \mathbb{R} \setminus \{0\}$: $P(t, e^{a_1 t}, \dots, e^{a_r t}, e^{b_1 t}, \dots, e^{b_s t}) = 0, Q(t, e^{a_1 t}, \dots, e^{a_r t}, e^{b_1 t}, \dots, e^{b_s t}) = 0$*

We say that two polynomials $P, Q \in \mathcal{R}$ are *associates* if $Q = \mathbf{z}^{\mathbf{u}} P$, where $\mathbf{z}^{\mathbf{u}}$ is a monomial in z_1, \dots, z_s (note that the associate relation is symmetric by the definition of \mathcal{R}).

We have seen that we can write the exponential polynomial as a Laurent polynomial in \mathcal{R} : $f(t) = P(t, e^{a_1 t} \dots, e^{a_r t}, e^{b_1 t} \dots, e^{b_s t})$. As $f(t)$ is real-valued, it can be factored in irreducible polynomials from \mathbb{K} that are either real valued, or come with their conjugate pair in the factorization of f . Therefore, there exist some irreducible polynomials Q_1, \dots, Q_k in \mathcal{R} that come in pair with their conjugates and some irreducible polynomials R_1, \dots, R_l in \mathcal{R} and positive integers $\alpha_1 \dots \alpha_k, \beta_1, \dots, \beta_l$ such that we can write $P = \prod_{i=1}^k (Q_i Q_i^*)^{\alpha_i} \cdot \prod_{j=1}^l R_j^{\beta_j}$.

Define the functions $u_i(t) := Q_i(t, e^{a_1 t} \dots, e^{a_r t}, e^{b_1 t} \dots, e^{b_s t})$, which are not real-valued and come in pairs with their conjugates; and $v_j(t) := R_j(t, e^{a_1 t} \dots, e^{a_r t}, e^{b_1 t} \dots, e^{b_s t})$, which are real-valued. Let $w_i(t) := u_i(t) \overline{u_i(t)}$, for $t \in \mathbb{R}$. Then $f(t) = \prod_{i=1}^k w_i(t)^{\alpha_i} \cdot \prod_{j=1}^l v_j(t)^{\beta_j}$, where functions $w_1, \dots, w_k, v_1, \dots, v_l$ are real-valued, analytic functions.

Recall that the decision problem asks whether $f(t)$ is strictly positive for some value of t in $[c, d]$, for some given $c, d \in \mathbb{Q}$. If $f(t)$ has a trivial form (i.e., if $f(t)$ is a polynomial in $\mathbb{K}[x]$, with no exponentials) we can easily decide this problem by approximating its roots in $[c, d]$ and classifying the sign on f between them (using, for example, the Sturm sequence of the polynomial). Otherwise, by Lindemann-Weierstrass Theorem (see [2]), we get that $f(t) = \sum_{j=1}^m P_j(t) e^{\lambda_j t}$ can be 0 in an algebraic point t if and only if $P_j(t) = 0$, for all $j \in \{1, \dots, m\}$. We can use standard factorization algorithms for computing common algebraic roots of the P_j polynomials. If there is any common algebraic root t^* , then $f(t^*) = 0$. As all the derivatives of f are also exponential polynomials, we can determine in a similar way the smallest M such that the M^{th} derivative of f is non-zero at t^* . By Taylor's Theorem, for any t , there exists some ϵ between t and t^* such that $f(t) = f(t^*) + \frac{f^{(M)}(\epsilon)}{M!} (t - t^*)^M$.

If M is odd, then f changes sign at t^* . Otherwise, there is no sign change at t^* . We can therefore deal with common roots of all polynomials $P_j(t)$. Then, we can assume without loss of generality that the P_j polynomials have no common root in (c, d) .

We can trivially get rid of the case when all the polynomials $P_j(t)$ have c as a common root by dividing them by the highest power of $(t - c)$ that divides all of them. This can be done safely, as changing signs at c does not make sense within the interval $[c, d]$ and as $(t - c)$ is always positive for $t > c$, so this division does not affect potential sign changes of $f(t)$ anywhere in $[c, d]$. The same holds for d . Therefore, we can assume without loss of generality that the P_j polynomials have no common root in $[c, d]$.

As P_j 's cannot all be 0 at the same (algebraic) point, by Lindemann-Weierstrass Theorem we get that f is non-zero in any rational point. In particular, $f(c)$ and $f(d)$ are non-zero, so we can use any standard approximation procedure until we can compare $f(c), f(d)$ with 0 (for example, see [2, Lemma 2]). If either of $f(c), f(d)$ is strictly greater than 0, then we are done. Therefore, assume from now on that both $f(c)$ and $f(d)$ are strictly negative.

No two different functions from $w_1(t), \dots, w_k(t), v_1(t), \dots, v_l(t)$ can have a common real zero in $[c, d]$, as this would imply that two of the polynomials $Q_1, \dots, Q_k, R_1, \dots, R_l$ have a common solution of the form $(t, e^{a_1 t}, \dots, e^{a_r t}, e^{b_1 i t}, \dots, e^{b_s i t})$, which contradicts Lemma 11, as all the listed polynomials are irreducible (and not associates) and as $t = 0$ cannot be a solution of such functions, because of Lindemann-Weierstrass Theorem (as we have dealt with algebraic roots above, which are common roots of all P_j 's). This means that it is enough to decide whether there exists some function among w_i 's and v_j 's with odd exponent in f that changes its sign in $[c, d]$, as we can just consider the one which changes its sign at the least $\tau \in [c, d]$. If we let this function be g , we have $g(\tau) = 0$ and we then know that no other function among the w_i 's and v_j 's has a solution at τ , so there is some interval $I = (\tau - \epsilon, \tau + \epsilon)$ such that g has exactly opposite signs on $(\tau - \epsilon, \tau)$ and $(\tau, \tau + \epsilon)$, and no other function equals 0 on I . It is easy then to see that deciding whether $f(t) > 0$ for some $t \in [c, d]$ is equivalent to deciding whether any of the real-valued functions with odd exponent in f among w_i 's and v_j 's changes its sign in $[c, d]$. This follows easily as the real-valued functions with even exponents are always non-negative and cannot change sign. Therefore, we can assume without loss of generality that all exponents are 1, so $f(t) = \prod_{i=1}^k w_i(t) \cdot \prod_{j=1}^l v_j(t)$.

In general, classical numerical algorithms should work in most of the cases for our decision problem. However, when an exponential polynomial has a tangential zero, detecting it through such procedures requires infinite precision. The difficulty in solving our problem comes exactly from dealing with cases of such tangential zeros.

Let us now see how to decide if any of the v_j 's or w_i 's changes its sign on $[c, d]$.

Case 1: Decide if some v_j changes sign on $[c, d]$.

Recall that $v_j(t) = R_j(t, e^{a_1 t}, \dots, e^{a_r t}, e^{b_1 t}, \dots, e^{b_s t})$ is a real-valued function ($R_j = R_j^*$). Also, recall that we ruled out the case of $f(c) = 0$, so we can approximate arbitrarily close $v_j(c)$ to decide if it is positive or negative. Assume without loss of generality that $v_j(c) < 0$. Then, as $v_j(d) \neq 0$, if we get by approximating it that $v_j(d) > 0$, we are trivially done, so assume that $v_j(d) < 0$. We want to decide if there exists some $t \in [c, d]$ such that $v_j(t) > 0$.

In this case, we claim that deciding if there is some zero of v_j in $[c, d]$ is equivalent to deciding if it changes sign, i.e., the equations $v_j(t) = 0, v_j'(t) = 0$ have no common solution. So, if $v_j(t) = 0$ for some $t \in [c, d]$, there is a least such t (by continuity on a bounded interval), and it is easy to see that, if $v_j'(t) \neq 0$, we get that v_j' changes sign at t , from negative to positive.

7:16 Parametric Continuous Stochastic Logic

To see that $v_j(t) = 0, v'_j(t) = 0$ have no common solution, write $v'_j(t)$ as a polynomial in $t, e^{a_1 t}, \dots, e^{a_r t}, e^{b_1 t}, \dots, e^{b_s t}$ and get by a simple degree chasing argument that it is coprime with the polynomial $R_j(t, e^{a_1 t}, \dots, e^{a_r t}, e^{b_1 t}, \dots, e^{b_s t}) = v_j(t)$, thus getting a contradiction with Lemma 11 (see Type-2 polynomial argument in [4] for details).

In conclusion, we can use the decision procedure described in [4] for zero finding for the purpose of deciding sign changing.

Case 2: Decide if some w_i changes sign on $[c, d]$.

Recall that, for $t \in \mathbb{R}$:

$$w_i(t) = u_i(t) \overline{u_i(t)}.$$

Note that $w_i(t)$ cannot change sign at any real t , therefore this case is trivial, as $w_i(t) \geq 0$.

In conclusion, the Positivity Problem for exponential polynomials is decidable assuming Schanuel's Conjecture.

B Proof of Theorem 10

Proof. Let $(\langle \mathbf{u}, \mathbf{A}, \mathbf{v} \rangle, \langle a, b \rangle)$ be an instance of the Positivity Problem for matrix exponentials. Let $\mathbf{D} \in \mathbb{Q}^{k \times k}$ be a diagonal matrix such that $\mathbf{D} = \text{diag}(d_1, \dots, d_k)$, where $d_i := 1$ if $\mathbf{v}_i = 0$ and $d_i := v_i$, otherwise. Note that $d_i \geq 0$ for any i .

Now, by letting $\bar{\mathbf{v}} \in \mathbb{Q}^k$ be such that if $\mathbf{v}_i = 0$ then $\bar{\mathbf{v}}_i := 0$, and otherwise $\bar{\mathbf{v}}_i := 1$, it is clear that $\mathbf{v} = \mathbf{D}\bar{\mathbf{v}}$. By denoting $\mathbf{B} := \mathbf{D}^{-1}\mathbf{A}\mathbf{D}$ and $\bar{\mathbf{u}} := \mathbf{D}^\top \mathbf{u}$, we get:

$$\mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} = \mathbf{u}^\top \mathbf{D} \mathbf{D}^{-1} e^{\mathbf{A}t} \mathbf{D} \bar{\mathbf{v}} = \mathbf{u}^\top \mathbf{D} e^{\mathbf{D}^{-1} \mathbf{A} \mathbf{D} t} \bar{\mathbf{v}} = \bar{\mathbf{u}}^\top e^{\mathbf{B}t} \bar{\mathbf{v}} \quad (13)$$

We adopt the following construction and map used in [1] for a related reduction in the discrete case: let $\mathbf{P} \in \mathbb{Q}^{2k \times 2k}$ be a matrix obtained by replacing each entry b_{ij} of \mathbf{B} by the symmetric matrix $\begin{bmatrix} p_{ij} & q_{ij} \\ q_{ij} & p_{ij} \end{bmatrix}$, where $p_{ij} = \max\{b_{ij}, 0\}$ and $q_{ij} = \max\{-b_{ij}, 0\}$. Let ρ be a map which sends $\begin{bmatrix} a & b \\ b & a \end{bmatrix}$ to $a - b$ and, applied to a matrix which can be partitioned in blocks of the form before, sends each block to the according difference. It is easy to check that ρ is a (surjective) homomorphism from the ring of matrices in $\mathbb{Q}^{2k \times 2k}$ (which can be partitioned in 2×2 blocks of the form $\begin{bmatrix} a & b \\ b & a \end{bmatrix}$) to the ring of matrices in $\mathbb{Q}^{k \times k}$.

By looking at the power series expansion of the matrix exponential $e^{\mathbf{X}}$, because of its convergence we get $e^{\rho(\mathbf{M})} = \rho(e^{\mathbf{M}})$. Recall (13): $\mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} = \bar{\mathbf{u}}^\top e^{\mathbf{B}t} \bar{\mathbf{v}}$. As $\rho(\mathbf{P}) = \mathbf{B}$, we get:

$$\mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} = \bar{\mathbf{u}}^\top e^{\rho(\mathbf{P})t} \bar{\mathbf{v}} = \bar{\mathbf{u}}^\top \rho(e^{\mathbf{P}t}) \bar{\mathbf{v}}. \quad (14)$$

Write $\bar{\mathbf{u}} := (\alpha_1, \dots, \alpha_k)^\top$ and $\bar{\mathbf{v}} := (\beta_1, \dots, \beta_k)^\top$. Given $w_1, \dots, w_k \in \mathbb{Q}$, define $\mathbf{x} \in \mathbb{Q}^{2k}$ by

$$\mathbf{x} := (\alpha_1 + w_1, w_1, \alpha_2 + w_2, w_2, \dots, \alpha_k + w_k, w_k)^\top.$$

Let us also define $\mathbf{y} \in \mathbb{Q}^{2k}$ by

$$\mathbf{y} := (\beta_1, -\beta_1, \beta_2, -\beta_2, \dots, \beta_k, -\beta_k)^\top.$$

▷ **Claim 12.** For all $w_1, \dots, w_k \in \mathbb{Q}$ it holds that $\mathbf{x}^\top e^{\mathbf{P}t} \mathbf{y} = \mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v}$ for all $t \in \mathbb{R}$.

Proof. Let us fix a positive real t . Denote the elements of $e^{\mathbf{A}t} =: \begin{bmatrix} e_{11} & e_{12} & \dots & e_{1k} \\ \vdots & \vdots & \ddots & \vdots \\ e_{k1} & e_{k2} & \dots & e_{kk} \end{bmatrix}$ and, as $e^{\mathbf{A}t} = \rho(e^{\mathbf{P}t})$, we can write:

$$e^{\mathbf{P}t} = \begin{bmatrix} f_{11} & g_{11} & f_{12} & g_{12} & \dots & f_{1k} & g_{1k} \\ g_{11} & f_{11} & g_{12} & f_{12} & \dots & g_{1k} & f_{1k} \\ \vdots & \vdots & \ddots & \vdots & & & \\ f_{k1} & g_{k1} & f_{k2} & g_{k2} & \dots & f_{kk} & g_{kk} \\ g_{k1} & f_{k1} & g_{k2} & f_{k2} & \dots & g_{kk} & f_{kk} \end{bmatrix}, \quad (15)$$

where $f_{ij} - g_{ij} = e_{ij}$ for all i, j . Then, we get:

$$\begin{aligned} \mathbf{x}^\top e^{\mathbf{P}t} \mathbf{y} &= \sum_{i=1}^k \sum_{j=1}^k ((\alpha_i + w_i) f_{ij} \beta_j - (\alpha_i + w_i) g_{ij} \beta_j + w_i g_{ij} \beta_j - w_i f_{ij} \beta_j) \\ &= \sum_{i=1}^k \sum_{j=1}^k (\alpha_i (f_{ij} - g_{ij}) \beta_j) = \sum_{i=1}^k \sum_{j=1}^k (\alpha_i e_{ij} \beta_j) \\ &= \mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v}. \end{aligned}$$

As $t \in \mathbb{R}$ was arbitrary, we get that the claim holds. \square

Choose w_i 's such that \mathbf{x} has only positive entries: $w_1 = \dots = w_k := \max(|\alpha_1|, \dots, |\alpha_k|) + 1$. Let $S > 0$ be the sum of \mathbf{x} 's entries; and let $\mathbf{z} := \frac{1}{S} \mathbf{x}$. Then, by Claim 12, we have: $\mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} > 0 \iff \mathbf{x}^\top e^{\mathbf{P}t} \mathbf{y} > 0 \iff (\frac{1}{S} \mathbf{x})^\top e^{\mathbf{P}t} \mathbf{y} > 0 \iff \mathbf{z}^\top e^{\mathbf{P}t} \mathbf{y} > 0$.

Note that we reduced the Positivity Problem for matrix exponentials to the one above, where \mathbf{z} is a stochastic vector. Also, \mathbf{y} 's entries are either -1 , 0 , or 1 .

Let the entries of \mathbf{P} be p_{ij} . Let us now pick a number r that is greater than the sum of any row of \mathbf{P} : $r > \sum_{j=1}^{2k} p_{ij}$, for each $1 \leq i \leq 2k$. Let $q_i := r - \sum_{j=1}^{2k} p_{ij}$, for each $1 \leq i \leq 2k$, and let \mathbf{Q} be a diagonal matrix: $\mathbf{Q} := \text{diag}(q_1, \dots, q_{2k})$.

Let us define $\mathbf{O} := \begin{bmatrix} \mathbf{P} - r\mathbf{I} & \mathbf{Q} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$, where each of the four blocks of \mathbf{O} is a $2k \times 2k$ matrix. We note that \mathbf{O} is a rate matrix. By inspection of the block multiplications, it is easy to see that the top left block of \mathbf{O}^n is $(\mathbf{P} - r\mathbf{I})^n$. Hence, by the power series expansion of $e^{\mathbf{X}}$ and by setting $\mathbf{u}_1 := \begin{bmatrix} \mathbf{z} \\ \mathbf{0} \end{bmatrix}$, $\mathbf{v}_1 := \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix}$, we get: $\mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_1 = \mathbf{z}^\top e^{(\mathbf{P}-r\mathbf{I})t} \mathbf{y} = \frac{1}{e^{rt}} \mathbf{z}^\top e^{\mathbf{P}t} \mathbf{y}$, so

$$\mathbf{z}^\top e^{\mathbf{P}t} \mathbf{y} > 0 \iff \mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_1 > 0.$$

Thus, we reduced the initial problem to the existence of a t in $[a, b]$ such that $\mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_1 > 0$, where \mathbf{u}_1 is stochastic, \mathbf{O} is a rate matrix and \mathbf{v}_1 has entries in $\{-1, 0, 1\}$.

By letting $\mathbf{v}_2 := \mathbf{v}_1 + \mathbf{1}$: $\mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_1 > 0 \iff \mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_2 > \mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{1} = 1$. Furthermore, $\mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_2 > 1 \iff \mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_3 > \frac{1}{2}$, where $\mathbf{v}_3 := \frac{1}{2} \mathbf{v}_2$, so \mathbf{v}_3 's entries are in $\{0, \frac{1}{2}, 1\}$.

We have reduced the problem whether there exists some t in $[a, b]$ such that $\mathbf{u}^\top e^{\mathbf{A}t} \mathbf{v} > 0$, where \mathbf{u}, \mathbf{v} are any vectors and \mathbf{A} is any matrix, to the problem whether there exists some t in $[a, b]$ such that $\mathbf{u}_1^\top e^{\mathbf{O}t} \mathbf{v}_3 > \frac{1}{2}$, where \mathbf{u}_1 is stochastic, \mathbf{O} is a rate matrix and \mathbf{v}_3 has entries in $\{0, \frac{1}{2}, 1\}$.

7:18 Parametric Continuous Stochastic Logic

This last problem asks for a given continuous-time Markov chain whether there exists a moment t in $[a, b]$ such that summing the probabilities of being in certain states at time t with fixed weights in $\{0, \frac{1}{2}, 1\}$ yields a result greater than $\frac{1}{2}$. We reduce this problem to a similar one with coefficients in $\{0, 1\}$ by splitting the states in the former problem that have weight $\frac{1}{2}$ in two identical states that, seen together as a black box, act as the original state.

More formally, if a state s_i has associated coefficient $\frac{1}{2}$ in \mathbf{v}_3 , we split it into states $s_{i,1}$ and $s_{i,2}$ and modify the transition rates:

- for any state s_j having a strictly positive transition rate $r_{j,i}$ to s_i : delete this transition and add two new transition rates to $s_{i,1}$ and $s_{i,2}$, both with rate $\frac{r_{j,i}}{2}$,
- for any state s_j such that there is a strictly positive transition rate $r_{i,j}$ from s_i to s_j : delete this transition and add two new transition rates from $s_{i,1}$ and $s_{i,2}$ to s_j , both with rate $r_{j,i}$.

Note that by being in state $s_{i,1}$ or $s_{i,2}$ in the new Markov chain we get the same behaviour as being in s_i in the original Markov chain (all the outgoing outgoing rates from $s_{i,1}$ or $s_{i,2}$ stay the same as the outgoing rates from s_i). We also modify the initial distribution: if the initial probability of s_i was p_i , set the new initial probability in s_i to 0 and both initial probabilities of $s_{i,1}$ and $s_{i,2}$ to $\frac{p_i}{2}$.

By regarding the cluster of states $s_{i,1}$ and $s_{i,2}$ as a “black-box state”, it behaves equivalently to state s_i in the original Markov chain. Because of the symmetry, the probability of being in $s_{i,1}$ at time t equals the probability of being in state $s_{i,2}$ at time t , which is equal to half of the probability of being in state s_i at time t in the original Markov chain.

Starting from the CTMC with rate transition matrix \mathbf{O} , initial distribution \mathbf{u}_1 and coefficient vector \mathbf{v}_3 , we can iteratively apply the described splitting process, by going through all the states having weights in the original formulation equal to $\frac{1}{2}$. We then get a sequence of new continuous-time Markov chains $\mathcal{M}_1, \dots, \mathcal{M}_N$ with rate matrices $\mathbf{R}_1, \dots, \mathbf{R}_N$ and with initial distributions $\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_N$, and new weight vectors $\tilde{\mathbf{v}}_1, \dots, \tilde{\mathbf{v}}_N$ defined as follows:

- 0/1 in the corresponding positions in $\tilde{\mathbf{v}}_{i+1}$ of states having previous coefficients 0/1 in $\tilde{\mathbf{v}}_i$,
- 0 in the corresponding positions in $\tilde{\mathbf{v}}_{i+1}$ of the most recent split state having previous coefficient 1/2 in $\tilde{\mathbf{v}}_i$,
- 1 in the corresponding position in $\tilde{\mathbf{v}}_{i+1}$ of the first newly created state by splitting (of the form $s_{i,1}$) and 0 to the second such state (of the form $s_{i,2}$),
- $\frac{1}{2}$ in the corresponding position in $\tilde{\mathbf{v}}_{i+1}$ of all other states having previous coefficients $\frac{1}{2}$ in $\tilde{\mathbf{v}}_i$.

We have the invariant $\tilde{\mathbf{u}}_i^\top e^{\mathbf{R}_i t} \tilde{\mathbf{v}}_i = \tilde{\mathbf{u}}_{i+1}^\top e^{\mathbf{R}_{i+1} t} \tilde{\mathbf{v}}_{i+1}$. Let $\mathcal{M} := \mathcal{M}_N$ be the CTMC obtained after the iterative splitting process described above, with rate matrix $\mathbf{R} := \mathbf{R}_N$ and initial distribution $\tilde{\mathbf{u}} := \tilde{\mathbf{u}}_N$, and let the final coefficient vector be $\tilde{\mathbf{v}} := \tilde{\mathbf{v}}_N$. It is clear now that $\mathbf{u}_1^\top e^{\mathbf{O} t} \mathbf{v}_3 = \tilde{\mathbf{u}}^\top e^{\mathbf{R} t} \tilde{\mathbf{v}}$.

Consequently, $\exists t \in [a, b]$ s.t. $\mathbf{u}^\top e^{\mathbf{A} t} \mathbf{v} > 0 \iff \exists t \in [a, b]$ s.t. $\tilde{\mathbf{u}}^\top e^{\mathbf{R} t} \tilde{\mathbf{v}} > \frac{1}{2}$, where $\tilde{\mathbf{u}}$ is a stochastic vector, \mathbf{R} is a rate matrix and $\tilde{\mathbf{u}}$ has entries in $\{0, 1\}$. We conclude that the Positivity Problem for matrix exponentials is reducible to the Threshold Problem for continuous-time Markov chains. ◀