

On LTL Model Checking for Low-Dimensional Discrete Linear Dynamical Systems

Toghrul Karimov

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany
toghs@mpi-sws.org

Joël Ouaknine

Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany
Department of Computer Science, University of Oxford, UK
joel@mpi-sws.org

James Worrell

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

Abstract

Consider a discrete dynamical system given by a square matrix $M \in \mathbb{Q}^{d \times d}$ and a starting point $s \in \mathbb{Q}^d$. The *orbit* of such a system is the infinite trajectory $\langle s, Ms, M^2s, \dots \rangle$. Given a collection $T_1, T_2, \dots, T_m \subseteq \mathbb{R}^d$ of semialgebraic sets, we can associate with each T_i an atomic proposition P_i which evaluates to *true* at time n if, and only if, $M^n s \in T_i$. This gives rise to the *LTL Model-Checking Problem* for discrete linear dynamical systems: given such a system (M, s) and an LTL formula over such atomic propositions, determine whether the orbit satisfies the formula. The main contribution of the present paper is to show that the LTL Model-Checking Problem for discrete linear dynamical systems is decidable in dimension 3 or less.

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

Keywords and phrases Linear dynamical systems, Orbit Problem, LTL model checking

Digital Object Identifier 10.4230/LIPIcs.MFCS.2020.54

Related Version <http://arxiv.org/abs/2007.02911>

Funding *Joël Ouaknine*: Supported by ERC grant AVS-ISS (648701) and DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>).

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

1 Introduction

A *discrete-time linear dynamical system* consists of a square matrix $M \in \mathbb{Q}^{d \times d}$ and a starting point $s \in \mathbb{Q}^d$. Its trajectory, the *orbit* of s under M , is the infinite sequence $\langle s, Ms, M^2s, \dots \rangle$. Such systems constitute a family of fundamental models, and decision problems associated with their trajectories arise frequently in the analysis of automata, Markov chains, linear recurrence sequences, and linear while loops (see, e.g., [8, 10, 13] and references therein).

One of the earliest decision problems for linear dynamical systems was formulated by Harrison in 1969 [11], and subsequently baptised the “*Orbit Problem*” by Kannan and Lipton, who famously solved it a decade later [12]. The Orbit Problem asks, given a linear dynamical system (M, s) in ambient space \mathbb{R}^d together with a point target $t \in \mathbb{Q}^d$, whether the orbit of s under M reaches t . Kannan and Lipton established polynomial-time decidability of the Orbit Problem in all dimensions. In subsequent work [13], Kannan and Lipton speculated that more complex decision problems might also be decidable; specifically, they considered variants of the Orbit Problem in which the target t is replaced by a linear subspace $T \subseteq \mathbb{R}^d$. They conjectured that for one-dimensional subspaces, reachability might remain decidable,



© Toghrul Karimov, Joël Ouaknine, and James Worrell;
licensed under Creative Commons License CC-BY

45th International Symposium on Mathematical Foundations of Computer Science (MFCS 2020).

Editors: Javier Esparza and Daniel Král'; Article No. 54; pp. 54:1–54:14



Leibniz International Proceedings in Informatics

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

but in the same breath they noted that when T is a $(d - 1)$ -dimensional subspace of \mathbb{R}^d , the corresponding reachability problem is precisely equivalent to the well-known *Skolem Problem* asking whether a linear recurrence sequence has a zero, which itself has been open for many decades [10, 23] (although decidability is known in dimension 4 or less [18, 24]).

The problem of reaching a linear subspace was studied by Chonev *et al.* in [6, 8], in which the authors established decidability for subspaces of dimension up to three (regardless of the dimension of the ambient space). Chonev *et al.* then turned their attention to the *Polyhedron-Hitting Problem* [7], in which the target is an arbitrary polyhedron. Decidability in dimension 3 was established, but the authors showed that in dimensions 4 or higher, solving the Polyhedron-Hitting Problem would necessarily entail major breakthroughs in Diophantine approximation, considered out of reach at the present time. More recently, Almagor *et al.* studied the *Semialgebraic Orbit Problem*, in which the target is an arbitrary semialgebraic set [3]. Once again, decidability in dimension 3 was shown to hold. Finally, in very recent work, and building on earlier results [2], Almagor *et al.* introduced a unifying framework for formulating *reachability* queries for discrete linear dynamical systems [4], subsuming all of the above problems. Roughly speaking, the authors considered instances in which both the source and target are semialgebraic sets, and a specification formalism in which one may quantify over members of these sets. Crucially, however, their *First-Order Orbit Queries* framework only allows *reachability* queries (“is there a positive integer n such that, after n steps, such and such properties hold?”). Almagor *et al.* established decidability in dimension 3.

Main contributions. In this paper, going considerably beyond reachability, we tackle the problem of full *LTL model checking* for orbits of discrete-time linear dynamical systems in dimension 3 (or less). More precisely, we are given a linear dynamical system (M, s) (with a singleton starting point $s \in \mathbb{Q}^3$), together with a collection $T_1, \dots, T_m \subseteq \mathbb{R}^3$ of semialgebraic sets, and an LTL formula φ over atomic propositions P_1, \dots, P_m . The atomic proposition P_i evaluates to *true* at time n if, and only if, the n -th component of the orbit lies in T_i , i.e., $M^n s \in T_i$. Such a framework enables one to formulate vastly more sophisticated and complex properties of orbits than mere reachability. **Our main result is that the LTL Model-Checking Problem for discrete-time linear dynamical systems in three dimensions is decidable, with complexity in exponential space.**

Some remarks are in order:

1. Since we have a single starting point, the orbit consists of a *single* trajectory. The problem we are solving is sometimes referred to in the literature as “*path checking*”, although typical applications in runtime verification and online monitoring involve finite traces, e.g., [15]. Path checking ultimately periodic infinite traces is considered in [16], but the traces arising from linear dynamical systems need not be ultimately periodic (see [1]).
2. Our framework is limited to dynamical systems in three (or fewer) dimensions. As mentioned earlier, it is known that mere polyhedral reachability is “hard” (in a Diophantine-approximation sense) in dimensions 4 and above, and as LTL model checking with semialgebraic targets vastly generalises polyhedral reachability, one cannot reasonably expect to prove decidability in dimensions higher than 3.
3. Beyond the search for maximal versatility and generality, our use of semialgebraic sets – rather than, say, products of intervals or polyhedra – in our specification framework has a practical motivation: in application areas such as program analysis, semialgebraic sets are indispensable to formulate sufficiently expressive properties, whether one seeks to overapproximate a set of reachable states, or to synthesise invariants or barrier certificates; see, e.g., [14].

On a technical level, our approach makes extensive use of spectral techniques and relies on various tools from algebraic and transcendental number theory, notably Baker’s theorem on linear forms in logarithms of algebraic numbers, as well as Kronecker’s theorem in Diophantine approximation.

In [1], Agrawal *et al.* consider a problem that is closely related to ours, namely the approximate verification of the symbolic dynamics of Markov chains. More specifically, they view a Markov chain as a distribution transformer: a stochastic matrix M and an initial probability distribution s give rise to an orbit $\langle s, Ms, M^2s, \dots \rangle$. They further discretise the probability space into finitely many boxes (products of intervals), which give rise to atomic propositions in exactly the same manner as in our setting. They then consider LTL model checking over the resulting formalism, but observe that the set of infinite words arising as symbolic trajectories of a given Markov chain can fail to be ω -regular; consequently, they switch their attention to “ ϵ -approximations” of the model-checking problem (the precise definitions are technical) and are able to establish decidability in *all* dimensions. This variant of the model-checking problem does not allow to check a specific path and thereby circumvents many of the difficulties arising in the present paper. The two pieces of work are therefore fairly distinct, both in terms of their respective scope and in the mathematical approach taken, despite sharing similar motivations.

2 Mathematical background

A semialgebraic set $T \subseteq \mathbb{R}^n$ is defined by a Boolean combination of polynomial inequalities of the form $p(x_1, \dots, x_n) \geq 0$ and $q(x_1, \dots, x_n) > 0$ for polynomials $p, q \in \mathbb{Z}[x_1, \dots, x_n]$.

2.1 Algebraic numbers

A complex number α is algebraic if it is a root of a polynomial p with integer coefficients. We denote the set of algebraic numbers by \mathbb{A} . For an algebraic number α , its defining polynomial p_α is the unique polynomial of the least degree that has α as a root and coefficients that do not share common factors. Given a polynomial $p \in \mathbb{Z}[x]$, let $\|p\|$ denote the bit length of its representation as a list of coefficients encoded in binary, $\deg(p)$ denote its degree and $H(p)$ denote its height (i.e. the maximum of the absolute value of coefficients of p). Throughout this work we make an extensive use of the facts that for each pair α, β of algebraic numbers, $\deg(\alpha\beta) \leq \deg(\alpha) + \deg(\beta)$ and $H(\alpha\beta) \leq H(\alpha)H(\beta)$.

An algebraic number α can be represented using its defining polynomial p_α together with rational approximations of its real and imaginary parts to sufficient precision. More precisely, α can be represented by $(p_\alpha, a, b, r) \in \mathbb{Z}[x] \times \mathbb{Q}^3$ provided that α is the unique root of p_α in the circle of radius r around $a + bi$. A separation bound due to Mignotte [17] asserts that for roots $\alpha \neq \beta$ of a polynomial $p \in \mathbb{Z}[x]$,

$$|\alpha - \beta| > \frac{\sqrt{6}}{d^{(d+1)/2} H^{d-1}}$$

where d, H are the degree and height of p_α , respectively. Thus if r is less than a quarter of the root separation bound, then the representation is well-defined. Given a polynomial $p \in \mathbb{Z}[x]$, we can compute a standard representation of each of its roots in time polynomial in $\|p\|$ [5]. Thus for an algebraic number α , we denote by $\|\alpha\|$ the bit length of its standard representation.

Given representations of algebraic numbers α, β we can effectively compute representations for the algebraic numbers $\alpha + \beta, \alpha\beta, \frac{1}{\alpha}, |\alpha|, \operatorname{Re}(\alpha), \operatorname{Im}(\alpha)$ in time polynomial in $\|\alpha\| + \|\beta\|$. Efficient algorithms for these tasks can be found in [5, 9].

2.2 Number-theoretic bounds

Throughout the paper, we make an extensive use of the following lemma, which itself is a consequence of the celebrated Baker-Wüstholz theorem.

► **Lemma 1** ([19]). *There exists a constant C such that for algebraic numbers α, β , for every $n \geq 2$ if $\alpha^n \neq \beta$, then $|\alpha^n - \beta| \geq \frac{1}{n^{(\|\alpha\| + \|\beta\|)^C}}$.*

Lemma 2 below states the following: if we start at an arbitrary point $\gamma \in \mathbb{T}$ on the unit circle, and repeatedly apply rotation through $\arg(\lambda)$ radians for $\lambda \in \mathbb{T} \cap \mathbb{A}$ that is not a root of unity, we will enter any open interval $J \subseteq \mathbb{T}$ in at most a certain number of steps that does not depend on the starting point γ but depends on the size of J . Henceforth we denote by $|J|$ the arc length of the interval $J \subseteq \mathbb{T}$ in radians.

► **Lemma 2.** *There exists a constant D such that for every $\lambda \in \mathbb{T} \cap \mathbb{A}$ that is not a root of unity and open subinterval J of \mathbb{T} , for each $\gamma \in \mathbb{T}$, $\gamma\lambda^n \in J$ for some $n < 2\pi \left(\frac{2\pi}{|J|}\right)^{\|\lambda\|^D}$.*

Proof. By the Pigeonhole principle, if there are $N_1 > \frac{2\pi}{|J|}$ points on the unit circle, at least two of them will have arc distance smaller than $|J|$. We select $N_1 = \left\lfloor \frac{2\pi}{|J|} \right\rfloor + 1 > \frac{2\pi}{|J|}$ and consider the sequence $\langle \lambda, \lambda^2, \dots \rangle$. By the preceding argument, there exist $1 \leq k < m \leq N_1$ such that λ^k and λ^m have arc distance smaller than $|J|$. That is, $\arg(\lambda^{m-k}) < |J|$.

Bounding arc length from below with Euclidean distance and using Lemma 1,

$$\arg(\lambda^{m-k}) \geq |\lambda^m - \lambda^k| = |\lambda^{m-k} - 1| \geq \frac{1}{(m-k)(\|\lambda\| + \|1\|)^C} \geq \frac{1}{\left\lfloor \frac{2\pi}{|J|} \right\rfloor^{(\|\lambda\| + \|1\|)^C}}.$$

Now consider the sequence $z_i = \gamma\lambda^{i(m-k)}$ for $i \geq 0$. As the arc distance between any consecutive terms z_i, z_{i+1} is less than $|J|$, this sequence must enter J before winding around the unit circle once. We therefore obtain that for some

$$n_1 \leq \left\lfloor \frac{2\pi - |J|}{\arg(\lambda^{m-k})} \right\rfloor + 1 \leq \left\lfloor (2\pi - |J|) \left\lfloor \frac{2\pi}{|J|} \right\rfloor^{(\|\lambda\| + \|1\|)^C} \right\rfloor + 1 \leq 2\pi \left(\frac{2\pi}{|J|}\right)^{(\|\lambda\| + \|1\|)^C}$$

$z_{n_1} \in J$. Translating this back to the sequence $\langle \lambda, \lambda^2, \dots \rangle$ we have that $\gamma\lambda^n \in J$ for some

$$n \leq \frac{2\pi}{|J|} \cdot 2\pi \left(\frac{2\pi}{|J|}\right)^{(\|\lambda\| + \|1\|)^C} = 2\pi \left(\frac{2\pi}{|J|}\right)^{(\|\lambda\| + \|1\|)^C + 1}.$$

Finally, choosing D such that $(\|\lambda\| + \|1\|)^C + 1 \leq \|\lambda\|^D$ yields the desired result. ◀

3 The LTL Model-Checking Problem

Suppose we are given a matrix $M \in \mathbb{Q}^{3 \times 3}$, a point $s \in \mathbb{Q}^3$ and an LTL formula φ over semialgebraic predicates $T_1, \dots, T_m \subseteq \mathbb{R}^3$ as an input. We associate with each T_i an atomic proposition P_i which evaluates to *true* at time n in case $M^n s \in T_i$. Hence we associate an ω -word w over $2^{\{P_1, \dots, P_m\}}$ with the orbit $\langle s, Ms, M^2s, \dots \rangle$ in the standard manner. The LTL Model-Checking Problem is then to decide whether $w \models \varphi$.

We assume that φ is given in the form

$$\varphi := T \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \mathbf{U} \varphi \mid \varphi \mathbf{R} \varphi \mid \mathbf{X} \varphi$$

where T is an atomic semialgebraic set described via a single polynomial inequality of the form $p(x) > 0$ or $p(x) \geq 0$. Observe that φ does not contain the \neg operator: an arbitrary formula ψ over semialgebraic sets (defined by a Boolean combination of inequalities of the type $p(x) > 0$ or $p(x) \geq 0$) can be translated into an equivalent formula in this form by first translating ψ into negation-normal form and then replacing $\neg p(x) \geq 0$ with $-p(x) > 0$ and $\neg p(x) > 0$ with $-p(x) \geq 0$. This translation incurs at most a linear blowup in size.

Throughout the paper we assume that the polynomials p defining the atomic predicates are given as a list of coefficients including (possibly many) zeros. Hence it is always the case that $\|p\| \geq \deg(p)$.

We analyse the problem based on the eigenvalues of M . Our main result is the following.

► **Theorem 3.** *Given $M \in \mathbb{Q}^{3 \times 3}$, $s \in \mathbb{Q}^3$ and φ , the LTL Model-Checking Problem is decidable in **EXPSpace** in $\|M\| + \|s\| + \|\varphi\|$.*

4 When not all three eigenvalues are real

We first consider the case in which M has complex eigenvalues $\lambda, \bar{\lambda}$ and a real eigenvalue ρ . Moreover, we assume that $\lambda^k \notin \mathbb{R}$ for all $k \in \mathbb{N}$, i.e. $\gamma = \frac{\lambda}{|\lambda|}$ is not a root of unity. This case requires by far the most detailed analysis. However, the final model-checking algorithm is quite simple in that it does not involve any non-trivial manipulations of algebraic numbers or semialgebraic sets.

4.1 Preliminary analysis

In this section we introduce normalised expressions in order to study the set of all values of n for which the term $M^n s$ of the orbit is in an atomic semialgebraic set T . The treatment here mostly mirrors that in [4].

Since M is assumed to have three distinct eigenvalues, we can diagonalise $M = PDP^{-1}$ where $D = \begin{bmatrix} \lambda & 0 & 0 \\ 0 & \bar{\lambda} & 0 \\ 0 & 0 & \rho \end{bmatrix}$. Observe that $M^n s = PD^n P^{-1} s$ and P, P^{-1} contain algebraic entries of constant degree and height polynomial in $\|M\|$. Hence we can write

$$M^n s = \begin{bmatrix} a_1 \lambda^n + \bar{a}_1 \bar{\lambda}^n + c_1 \rho^n \\ a_2 \lambda^n + \bar{a}_2 \bar{\lambda}^n + c_2 \rho^n \\ a_3 \lambda^n + \bar{a}_3 \bar{\lambda}^n + c_3 \rho^n \end{bmatrix}$$

where a_1, a_2, a_3 and c_1, c_2, c_3 are all algebraic numbers with fixed degree and description length polynomial in $\|M\| + \|s\|$.

Let $T = \{x \in \mathbb{R}^3 : p(x) \sim 0\}$, $\sim \in \{>, \geq\}$ be an atomic semialgebraic set. We study the set $\mathcal{Z}(T) = \{n \geq 0 : M^n s \in T\} = \{n \geq 0 : p(M^n s) \sim 0\}$. In the full version of the paper we show that $p(M^n s)$ can be written as

$$\sum_{0 \leq p_1, p_2, p_3 \leq \deg(p)} \alpha_{p_1, p_2, p_3} \lambda^{np_1} \bar{\lambda}^{np_2} \rho^{np_3} + \overline{\alpha_{p_1, p_2, p_3}} \lambda^{np_1} \bar{\lambda}^{np_2} \rho^{np_3} \quad (1)$$

where each α_{p_1, p_2, p_3} is an algebraic number of degree polynomial and height exponential in $\|M\| + \|s\| + \|p\|$. If all the coefficients α_{p_1, p_2, p_3} above are 0 then $p(M^n s) = 0$ for all n , and hence the orbit is either always or never in the semialgebraic set T . In this case, we

can replace T in any LTL formula with **true** or **false**. Otherwise, let $\Lambda = \max\{|\lambda^{p_1} \bar{\lambda}^{p_2} \rho^{p_3}| : \alpha_{p_1, p_2, p_3} \neq 0\}$. Dividing the expression in (1) by Λ we obtain that $p(M^n s) \sim 0$ if and only if

$$\sum_{m=0}^k \beta_m \gamma^{nm} + \overline{\beta_m} \bar{\gamma}^{nm} + r(n) \sim 0$$

where

- $\gamma = \frac{\lambda}{|\lambda|}$ with degree at most 12 and height polynomial in $\|M\|$,
- $k \leq \deg(p)$,
- $r(n) = \sum_{l=1}^{k'} \chi_l \mu_l^n + \overline{\chi_l \mu_l^n}$ with $|\mu_l| < 1$ for every $1 \leq l \leq k'$,
- and all the coefficients and exponents β_m , $0 \leq m \leq k$ and χ_l, μ_l , $1 \leq l \leq k'$ are algebraic.

We refer to $e(n) = \sum_{m=0}^k \beta_m \gamma^{nm} + \overline{\beta_m} \bar{\gamma}^{nm} + r(n)$ as the *normalised expression corresponding to T* , and denote the bit-length of its syntactic representation by $\|e\|$, which can again be shown to be of size polynomial in $\|M\| + \|s\| + \|p\|$.

4.2 On visiting atomic semialgebraic sets

Recall that we defined, for an atomic semialgebraic set T , a matrix M and a starting point s , $\mathcal{Z}(T) = \{n \geq 0 : M^n s \in T\}$. We now study the structure of $\mathcal{Z}(T)$ and show how to compute a useful finite representation for it.

In this section, let $\|\mathcal{I}_T\| = \|M\| + \|s\| + \|p\|$, where p is the polynomial defining T . Let $e(n) = \sum_{m=0}^k \beta_m \gamma^{nm} + \overline{\beta_m} \bar{\gamma}^{nm} + r(n)$ be the normalised expression corresponding to T . We call the function $f(z) = \sum_{m=0}^k \beta_m z^m + \overline{\beta_m} \bar{z}^m$, $f : \mathbb{C} \rightarrow \mathbb{R}$ the *dominant function corresponding to T* . Observe that $e(n) = f(\gamma^n) + r(n)$.

From [4] we know that f has at most $4k$ zeros in the unit circle \mathbb{T} , which are algebraic numbers with description length polynomial in $\|f\|$, and that

- **Lemma 4.** *There exists $N = 2^{\|e\|^{O(1)}}$ such that for all $n > N$,*
- $f(\gamma^n) \neq 0$, and
- $f(\gamma^n) > 0$ iff $f(\gamma^n) + r(n) > 0$ iff $M^n s \in T$.

Since f is a continuous real-valued function, it maintains its sign between its (at most $4k$) roots on the unit circle. Recalling that $\|e\| = \|\mathcal{I}_T\|^{O(1)}$, we rephrase Lemma 4 as follows:

- **Theorem 5.** *Let T be an atomic semialgebraic set. There exist $N = 2^{\|\mathcal{I}_T\|^{O(1)}}$ and $J \subseteq \mathbb{T}$ that is a union of finitely many open intervals such that for $n > N$, $M^n s \in T$ if and only if $\gamma^n \in J$.*

Such J can be written uniquely as a disjoint union of open arcs in \mathbb{T} . We refer to such intervals as the *component intervals* or *components* of J . Observe that the endpoints of components of J are roots of f . Recall that γ is not a root of unity, and hence by Kronecker's theorem, the sequence $(\gamma^i)_{i \in \mathbb{N}}$ is dense in \mathbb{T} . It follows that the sequence $(\gamma^{N+i})_{i \in \mathbb{N}}$ is likewise dense in \mathbb{T} , and we obtain that J is unique, in the sense of being independent from any N that satisfies the conclusion of Theorem 5. Hence we refer to J as the *finite union of (open) intervals corresponding to T* .

4.3 $\mathcal{Z}(\varphi)$ for general φ

We now study the set $\mathcal{Z}(\varphi) = \{n \geq 0 : \langle M^n s, M^{n+1} s, \dots \rangle \models \varphi\}$, i.e. the set of all suffixes of the original orbit $\langle s, Ms, M^2 s, \dots \rangle$ that satisfy φ , for an arbitrary LTL formula φ . We extend Theorem 5 by showing that $\mathcal{Z}(\varphi)$ also has a corresponding union of finitely many open intervals that can be effectively computed from the finite unions of open intervals corresponding to its subformulas.

In this section, let $\|\mathcal{I}_b\| = \|M\| + \|s\| + \sum_{i=1}^m \|T_i\|$ where T_1, T_2, \dots, T_m are atomic predicates appearing in some formula φ . Intuitively, $\|\mathcal{I}_b\|$ is the “basic length” of the input that doesn’t account for the structure of φ . Our main result is the following:

► **Theorem 6.** *Let φ be an LTL formula. There exists $N = 2^{\|\mathcal{I}_b\|^{O(1)}}$ and a finite union of open intervals $J_\varphi \subseteq \mathbb{T}$ such that for all $n > N$, $n \in \mathcal{Z}(\varphi)$ if and only if $\gamma^n \in J_\varphi$.*

To prove this, we will combine Theorem 5 with the following result:

► **Theorem 7.** *Let semialgebraic sets T_1, T_2, \dots, T_m , time step N and finite unions of open intervals $J_1, J_2, \dots, J_m \subseteq \mathbb{T}$ be such that for all $1 \leq i \leq m$ and time steps $n > N$, $n \in \mathcal{Z}(T_i)$ if and only if $\gamma^n \in J_i$. Then for every LTL formula φ over T_1, T_2, \dots, T_m there exists a finite union of open intervals $J_\varphi \subseteq \mathbb{T}$ such that for all $n > N$, $n \in \mathcal{Z}(\varphi)$ if and only if $\gamma^n \in J_\varphi$. Moreover, such J_φ is unique.*

Proof. The uniqueness of J_φ , if such a finite union of open sets exists, can be established using the same topological argument as the one used in the uniqueness result accompanying Theorem 5. In order to prove existence of J_φ with the desirable properties we proceed by induction on the structure of φ . If $\varphi = T_i$, then $J_\varphi = J_i$ by assumption.

Next, let J_{φ_1} and J_{φ_2} be the finite unions of open intervals corresponding to φ_1 and φ_2 , respectively. Recall from Section 3 that we can assume φ does not contain the \neg operator.

1. Suppose $\varphi = \varphi_1 \vee \varphi_2$. Then $J_\varphi = J_{\varphi_1} \cup J_{\varphi_2}$.
2. Similarly, if $\varphi = \varphi_1 \wedge \varphi_2$ then $J_\varphi = J_{\varphi_1} \cap J_{\varphi_2}$.
3. Consider $\varphi = \mathbf{X}\varphi_1$. Suppose $n > N$. Then

$$\begin{aligned} n \in \mathcal{Z}(\mathbf{X}\varphi_1) &\iff n+1 \in \mathcal{Z}(\varphi_1) \\ &\iff \gamma^{n+1} \in J_{\varphi_1} \\ &\iff \gamma^n \in \gamma^{-1}J_{\varphi_1}. \end{aligned}$$

The first equivalence follows from the semantics of the \mathbf{X} operator, and the second from the fact that $n+1 > N$. Hence $J_\varphi = \gamma^{-1}J_{\varphi_1}$.

4. The main difficulty lies in analysing the case $\varphi = \varphi_1 \mathbf{U} \varphi_2$. If J_{φ_2} is empty, then J_φ will be empty too. Now suppose J_{φ_2} is not empty, and let l be length of a maximal interval in J_{φ_2} . Using Lemma 2 we can effectively compute a bound

$$b = b(\varphi_2) = 2\pi \left(\frac{2\pi}{l} \right)^{\|\gamma\|^D}$$

such that γ^n returns to J_{φ_2} after at most b time steps – that is, for every $n > N$, there exists $0 \leq \Delta \leq b$ such that $\gamma^{n+\Delta} \in J_{\varphi_2}$. Thus we have that for $n > N$,

$$\begin{aligned} n \in \mathcal{Z}(\varphi_1 \mathbf{U} \varphi_2) &\iff \exists \Delta \geq 0. n + \Delta \in \mathcal{Z}(\varphi_2) \wedge \forall m \in [n, n + \Delta). m \in \mathcal{Z}(\varphi_1) \\ &\iff \exists \Delta \in [0, b]. n + \Delta \in \mathcal{Z}(\varphi_2) \wedge \forall m \in [n, n + \Delta). m \in \mathcal{Z}(\varphi_1) \\ &\iff \bigvee_{\Delta=0}^b \left(n + \Delta \in \mathcal{Z}(\varphi_2) \wedge \bigwedge_{m=0}^{\Delta-1} n + m \in \mathcal{Z}(\varphi_1) \right). \end{aligned}$$

By the induction hypothesis, $n + \Delta \in \mathcal{Z}(\varphi_2) \iff \gamma^{n+\Delta} \in J_{\varphi_2}$, which is equivalent to $\gamma^n \in \gamma^{-\Delta} J_{\varphi_2}$. Similarly, $n + m \in \mathcal{Z}(\varphi_1) \iff \gamma^n \in \gamma^{-m} J_{\varphi_1}$. Hence we obtain that

$$n \in \mathcal{Z}(\varphi_1 \mathbf{U} \varphi_2) \iff \gamma^n \in \bigvee_{\Delta=0}^b \left(\gamma^{-\Delta} J_{\varphi_2} \wedge \bigwedge_{m=0}^{\Delta-1} \gamma^{-m} J_{\varphi_1} \right)$$

and therefore, the union of open intervals corresponding to φ is

$$J_\varphi = \bigcup_{\Delta=0}^b \left(\gamma^{-\Delta} J_{\varphi_2} \cap \bigcap_{m=0}^{\Delta-1} \gamma^{-m} J_{\varphi_1} \right).$$

5. Finally, suppose $\varphi = \varphi_1 \mathbf{R} \varphi_2$. If $J_{\varphi_1 \wedge \varphi_2} = J_{\varphi_1} \cap J_{\varphi_2}$ is empty, then $J_\varphi = \mathbb{T}$ if $J_{\varphi_2} = \mathbb{T}$ and J_φ is empty otherwise. If, on the other hand, $J_{\varphi_1 \wedge \varphi_2}$ is not empty, then $J_\varphi = J_{\varphi_2} \mathbf{U} (\varphi_1 \wedge \varphi_2)$ which can be computed using the preceding analysis. \blacktriangleleft

From the construction described above we can observe that the endpoints of components of J_φ come from those of its immediate subformulas. For example, the endpoints in $J_{\varphi_1 \mathbf{U} \varphi_2}$ are all endpoints of either φ_1 or φ_1 which have been multiplied by γ^{-1} for at most $b(\varphi_2)$ steps. In general, the endpoints of component intervals in J_φ are of the form $\gamma^{-n}z$ where n is an integer and z is a zero of a dominant function (as defined in Section 4.2) corresponding to some semialgebraic target set T_i appearing in φ .

To prove Theorem 6, recall from Theorem 5 that we already know that for each atomic T_i , there exist $N_i = 2^{(\|M\| + \|s\| + \|T_i\|)^{O(1)}}$ and J_i such that for $n > N$, $n \in \mathcal{Z}(T_i)$ if and only if $\gamma^n \in J_i$. Taking $N = \max_{1 \leq i \leq m} N_i = 2^{\|\mathcal{I}_b\|^{O(1)}}$ we obtain the desired result.

4.4 Analysing the inductive construction of $\mathcal{Z}(T)$ quantitatively

We now study how small the component intervals in the set J_φ corresponding to a formula φ can be. Our aim with this analysis is to be able to bound the *return time* $T(\varphi)$ of φ with respect to the orbit $\pi = \langle s, Ms, M^2s, \dots \rangle$, defined as

$$T(\varphi) = \sup\{t_2 - t_1 \mid t_1, t_2 \in \mathbb{N} \wedge \pi[t_2, \infty) \models \varphi \wedge \pi[t_1, \infty) \not\models \varphi \text{ for every } t_1 \leq t < t_2\}.$$

Informally, $T(\varphi)$ denotes the longest time φ remains false in $\pi = \langle s, Ms, M^2s, \dots \rangle$ before becoming true.

Recall from Section 4.3 that the endpoints of intervals in J_φ are of the form $\gamma^{-n}z$ for some z that is a root of a dominant function corresponding to an atomic predicate T_i appearing in φ . Hence for an endpoint $u \in \mathbb{T}$ we define the *retraction depth* of u to be the smallest integer n such that $u = \gamma^{-n}z$ for such z . Next, for an LTL formula φ over T_1, \dots, T_m , define

- $d(\varphi)$ to be the length of a smallest maximal interval in the finite union J_φ of intervals corresponding to φ ,
- $R(\varphi)$ to be the maximum of retraction depths of endpoints in J_φ , called the *retraction depth* of φ , and
- $D(\varphi)$ denote the maximum nesting depth of temporal operators in φ , with atomic formulas having depth 0.

Further, throughout this section let $\|\mathcal{I}_b\| = \|M\| + \|s\| + \sum_{i=1}^m \|T_i\|$ and $\|\mathcal{I}\| = \|M\| + \|s\| + \|\varphi\|$ where T_1, \dots, T_m are the atomic semialgebraic sets appearing in the input formula φ . We link the quantities defined above by analysing the inductive construction described in Theorem 7.

- **Lemma 8.** *For every φ , $d(\varphi) \geq \frac{1}{(R(\varphi)+2)\|\mathcal{I}_b\|^{O(1)}}$.*

Proof. We proceed by bounding how close two endpoints of an interval in J_φ can be given the retraction depth $R(\varphi)$ of J_φ . Let z_1, z_2 be roots of dominant functions corresponding to T_1, T_2 that appear in φ and $\gamma^{-n_1} z_1, \gamma^{-n_2} z_2$ two endpoints of component intervals in J_φ . We show that $\|\gamma^{-n_1} z_1 - \gamma^{-n_2} z_2\| \geq \frac{1}{(N+2)\|\mathcal{I}_b\|^{O(1)}}$ where $N = |n_1 - n_2|$. The statement of the lemma then follows from the fact that $N \leq R(\varphi)$ by definition.

For simplicity assume $n_1 \geq n_2$. Recall that the roots z_1, z_2 have degree $\|\mathcal{I}_b\|^{O(1)}$ and height $2\|\mathcal{I}_b\|^{O(1)}$ whereas γ has degree at most 12 and height $\|\mathcal{I}_b\|^{O(1)}$. Next, observe that

$$\|\gamma^{-n_1} z_1 - \gamma^{-n_2} z_2\| = \|z_1 - \gamma^{n_1-n_2} z_2\| = \left\| \frac{z_1}{z_2} - \gamma^{n_1-n_2} \right\| = \|z' - \gamma^{n_1-n_2}\|$$

where $z' \in \mathbb{T}$ is also of degree $\|\mathcal{I}_b\|^{O(1)}$ and height $2\|\mathcal{I}_b\|^{O(1)}$.

- If $n_1 - n_2 < 2$, then we use the root separation bound given in Section 2.1 to obtain that $\|z_1 - \gamma^{n_1-n_2} z_2\| \geq \frac{1}{2\|\mathcal{I}_b\|^{O(1)}}$. This can be done by separating the roots of the product polynomial $p_1 p_2$ where p_1, p_2 are polynomials that have z' and $\gamma^{n_1-n_2}$ as roots. Observe that $p_1 p_2$ itself is also of degree $\|\mathcal{I}_b\|^{O(1)}$ and height $2\|\mathcal{I}_b\|^{O(1)}$.
- If $n_1 - n_2 \geq 2$, then by a direct application of Lemma 1 we obtain that

$$\|z_1 - \gamma^{n_1-n_2} z_2\| \geq \frac{1}{(n_1 - n_2)(\|\gamma\| + \|z'\|)^C} = \frac{1}{(n_1 - n_2)\|\mathcal{I}_b\|^{O(1)}}.$$

Combining the two bounds yields the desired result. \blacktriangleleft

We now move onto analysing the retraction depth of φ . If φ does not contain any temporal operators then all the endpoints in J_φ are roots of dominant functions themselves. Hence, by definition, $R(\varphi) = 0$. For general φ , on the other hand, we have the following result.

► **Lemma 9.** *For every formula φ with temporal operator depth $D(\varphi) > 0$ there exist $k \leq D(\varphi)$ formulas $\varphi_1, \dots, \varphi_k$ over the same atomic predicates as φ such that $D(\varphi_1) < D(\varphi_2) < \dots < D(\varphi_k) < D(\varphi)$ and*

$$R(\varphi) \leq k + 2\pi \sum_{i=1}^k \left(\frac{2\pi}{d(\varphi_i)} \right)^{\|\gamma\|^D}.$$

Proof. Let φ be a formula with $D(\varphi) > 0$. We first show that there exist a subformula φ' of φ with smaller temporal operator depth and a formula φ'' (possibly not a subformula of φ) with $D(\varphi'') < D(\varphi)$ such that

$$R(\varphi) \leq 1 + R(\varphi') + 2\pi \left(\frac{2\pi}{d(\varphi'')} \right)^{\|\gamma\|^D}.$$

The statement of the lemma then follows by repeatedly applying the inequality to $R(\varphi')$ at most $D(\varphi)$ times. We proceed by a case analysis on the structure of φ .

1. If $\varphi = \varphi_1 \wedge \varphi_2$ or $\varphi = \varphi_1 \vee \varphi_2$, then $R(\varphi) \leq \max\{R(\varphi_1), R(\varphi_2)\}$. Hence we can take φ' to be the immediate subformula with the larger retraction depth and φ'' to be any subformula of φ ;
2. If $\varphi = \mathbf{X}\varphi_1$, then $R(\varphi) \leq 1 + R(\varphi')$ for a smaller subformula φ' (namely φ_1);
3. If $\varphi = \varphi_1 \mathbf{U}\varphi_2$, and J_{φ_2} is empty, then so is J_φ and $R(\varphi) = 0$. If J_{φ_2} is not empty, then the inductive construction shows that the endpoints of J_φ are all endpoints of J_{φ_1} or J_{φ_2} multiplied by γ^{-1} for at most $b = b(\varphi_2) = 2\pi \left(\frac{2\pi}{d(\varphi_2)} \right)^{\|\gamma\|^D}$ steps (Theorem 7, case

- 4). Hence $R(\varphi) \leq \max\{R(\varphi_1), R(\varphi_2)\} + b(\varphi_2)$ and we can take φ' to be the immediate subformula of φ with larger retraction depth and φ'' to be φ_2 (which indeed does have a smaller temporal operator depth and also happens to be a subformula of φ).
4. Finally, suppose $\varphi = \varphi_1 \mathbf{R} \varphi_2$. The only non-trivial case arises from non-empty $J_{\varphi_1 \wedge \varphi_2}$, where the endpoints of J_φ are all endpoints of $J_{\varphi_1 \wedge \varphi_2}$ or J_{φ_2} multiplied by γ^{-1} for at most $b(\varphi_1 \wedge \varphi_2)$ steps. In this case, similarly to the above, $R(\varphi) \leq R(\varphi') + 2\pi \left(\frac{2\pi}{d(\varphi')} \right)^{\|\gamma\|^D}$ where φ' is the immediate subformula of φ with smaller temporal operator depth and φ'' is a formula with smaller temporal operator depth (namely, $\varphi_1 \wedge \varphi_2$, which is not a subformula of φ). ◀

We now combine Lemmas 8 and 9 in the following way. Let q be a polynomial such that $\|\gamma\|^D \leq q(\|\mathcal{I}_b\|)$ and $d(\varphi) \geq \frac{1}{(R(\varphi)+2)^{q(\|\mathcal{I}_b\|)}}$. We obtain that for every φ with temporal operator depth $D(\varphi) > 0$ there exist $k \leq D(\varphi)$ formulas $\varphi_1, \dots, \varphi_k$ with $D(\varphi_1) < D(\varphi_2) < \dots < D(\varphi_k) < D(\varphi)$ such that

$$d(\varphi) \geq \frac{1}{(R(\varphi) + 2)^{q(\|\mathcal{I}_b\|)}} \geq \frac{1}{\left(k + 2 + 2\pi \sum_{i=1}^k \left(\frac{2\pi}{d(\varphi_i)} \right)^{q(\|\mathcal{I}_b\|)} \right)^{q(\|\mathcal{I}_b\|)}}.$$

In the full version of this paper we analyse this recursive relation and show the following.

► **Theorem 10.** *For any LTL formula φ , $d(\varphi) = \frac{1}{2^{2(\|\mathcal{I}_b\| + D(\varphi))^{O(1)}}}$. In particular, $d(\varphi) = \frac{1}{2^{2\|\mathcal{I}\|^{O(1)}}}$, where $\|\mathcal{I}\| = \|M\| + \|s\| + \|\varphi\|$.*

The interpretation is that all intervals in J_φ have length bounded below by the reciprocal of quantity whose magnitude is doubly exponential in the length of the input. In particular, we can compute a uniform lower bound that only depends on the encoding length of the atomic predicates and the depth of the temporal operators (in addition to $\|M\|$ and $\|s\|$) and not on the structure of φ .

We are now in a position to apply our quantitative analysis to the model-checking problem via the return time, as discussed at the beginning of this section.

► **Theorem 11.** *The return time $T(\varphi)$ of any LTL formula φ with respect to an orbit $\langle s, Ms, M^2s, \dots \rangle$ is $2^{2(\|\mathcal{I}_b\| + D(\varphi))^{O(1)}}$. In particular, $T(\varphi) = 2^{2\|\mathcal{I}\|^{O(1)}}$.*

Proof. Let $N = 2^{\|\mathcal{I}_b\|^{O(1)}}$ be the time after which whether the suffix $\langle M^n s, M^{n+1} s, \dots \rangle \models \varphi$ depends only on γ^n , as described in Theorem 7. Applying Lemma 2 to Theorem 10 we obtain that the return time $T'(\varphi)$ of φ with respect to $\langle M^{N+1} s, M^{N+2} s, \dots \rangle$ is $2^{2(\|\mathcal{I}_b\| + D(\varphi))^{O(1)}}$. Hence the return time with respect to the original orbit is at most $N + T'(\varphi) = 2^{2(\|\mathcal{I}_b\| + D(\varphi))^{O(1)}}$. ◀

We will use this result in Section 6 to construct, given an input formula, an equivalent formula (with respect to the given orbit) that only has bounded quantifiers and then proceed to solve the resulting finitary model-checking problem.

5 The remaining cases

Now suppose M has three real eigenvalues or eigenvalues $\lambda, \bar{\lambda}, \rho$ with $\gamma = \frac{\lambda}{|\lambda|}$ a root of unity. In the full version of this paper we show that in both cases, for an atomic semialgebraic target T , $\mathcal{Z}(T) = \{n \geq 0 : \langle M^n s, M^{n+1} s, \dots \rangle \in T\}$ is a semilinear set for which an explicit representation can be computed. In particular,

■ **Listing 1** Recursive model-checking algorithm for formulas with only bounded temporal operators.

```

ModelCheck(formula F, starting point n)
  case F = Until(F1, F2, upper bound B):
    for i=0 to B do
      if ModelCheck(F2, n+i) return true
      if not ModelCheck(F1, n+i) return false
    return false
  case F = Release(F1, F2, upper bound B):
    for i=0 to B do
      if not ModelCheck(F2, n+i) return false
      if ModelCheck(F1, n+i) return true
    return true
  case F = Next(F1)
    return ModelCheck(F1, n+1)
  case F = And(F1, F2):
    l = ModelCheck(F1, n)
    r = ModelCheck(F2, n)
    return l and r
  case F = Or(F1, F2):
    l = ModelCheck(F1, n)
    r = ModelCheck(F2, n)
    return l or r
  case F = atomic semialgebraic T:
    return Oracle(T, n)

```

► **Theorem 12.** *Given a semialgebraic set T , a square matrix $M \in \mathbb{Q}^{3 \times 3}$ with three real eigenvalues, and a starting point $s \in \mathbb{Q}$, there exists an integer $N = 2^{\|\mathcal{I}_T\|^{O(1)}}$ and a computable $X \subseteq \{0, 1\}$ such that for all $n > N$, $M^n s \in S$ if and only if $n \bmod 2 \in X$.*

► **Theorem 13.** *Given a semialgebraic set T , a square matrix $M \in \mathbb{Q}^{3 \times 3}$ with eigenvalues $\lambda, \bar{\lambda}, \rho$ where $\gamma = \frac{\lambda}{|\lambda|}$ is a root of unity, and a starting point $s \in \mathbb{Q}$, there exists an integer $N = 2^{\|\mathcal{I}_T\|^{O(1)}}$ and a computable $X \subseteq \{0, 1, \dots, 287\}$ such that for all $n > N$, $M^n s \in S$ if and only if $n \bmod 288 \in X$.*

Here once again $\|\mathcal{I}_T\| = \|p\| + \|M\| + \|s\|$, where p is the polynomial defining T . In the next section we discuss how to utilise Theorems 12 and 13 in order to obtain a decision procedure for the relevant cases of the LTL Model-Checking Problem.

6 Model-checking algorithm and its complexity

In this section we summarize our algorithmic contribution. Suppose we are given $M \in \mathbb{Q}^{3 \times 3}$, $s \in \mathbb{Q}^3$ and an LTL formula φ over semialgebraic T_1, \dots, T_m as the input. We describe a decision procedure for determining whether $\langle s, Ms, M^2s, \dots \rangle \models \varphi$.

Let us first consider the complexity of determining, for a given n , $M \in \mathbb{Q}^{3 \times 3}$, $s \in \mathbb{Q}^3$ and a semialgebraic target T defined via $p(x) \sim 0$, whether $M^n s \in T$. Using iterated squaring we can encode the statement $p(M^n s) \sim 0$ into the existential theory of real numbers using a formula of size $O(\|M\| \log n + \|p\| + \|s\|)$. If the input is M , s and an LTL formula φ containing T , this can be written as $O(\|\mathcal{I}\| + \log n)$. Since the existential theory of real numbers can be decided in polynomial space (see, e.g., [21]), an oracle for determining whether $M^n s$ is in a target set T can be implemented using space polynomial in $\|\mathcal{I}\| + \log n$.

We now move onto the main algorithm. As the first step, determine whether M has three real eigenvalues ρ_1, ρ_2, ρ_3 or two complex eigenvalues $\lambda, \bar{\lambda}$ and a real eigenvalue ρ . If the latter is the case, additionally determine whether $\gamma = \frac{\lambda}{|\lambda|}$ is a root of unity or not.

If M only has real eigenvalues or γ is a root of unity, we proceed by computing an explicit representation for the semilinear set $\mathcal{Z}(\varphi) = \{n \geq 0 : \langle M^n s, M^{n+1} s, \dots \rangle \models \varphi\}$. We illustrate how this can be done by using Theorem 13 and repeatedly combining semilinear sets in case where γ is a root of unity. In case M has three real eigenvalues the same procedure can be applied to Theorem 12.

We first compute an explicit representation for $\mathcal{Z}(T_i) = \{n \geq 0 : \langle M^n s, M^{n+1} s, \dots \rangle \in T_i\}$ for each atomic T_i in φ . To this end, we compute the value of $N_i = 2^{\|\mathcal{I}_{T_i}\|^{O(1)}}$ described in Section 5 for each $1 \leq i \leq m$ and then take the maximum $N = \max_{1 \leq i \leq m} N_i$. Next we determine $F_i = \{n \leq N : M^n s \in T_i\}$ and compute $X_i \subseteq \{0, 1, \dots, 287\}$ such that for $n > N$, $M^n s \in T_i$ if and only if $n \bmod 288 \in X_i$. These sets can be determined by making queries to the oracle of the form $M^n s \in T_i$ for $0 \leq n \leq N + 288$, requiring $2^{\|\mathcal{I}\|^{O(1)}}$ space in total. Finally, from sets F_i, X_i for $1 \leq i \leq m$ we can construct, for arbitrary formula φ , sets F and X such that for all $n \leq N$, $\langle M^n s, M^{n+1} s, \dots \rangle \models \varphi$ if and only if $n \in F$ and for all $n > N$, $\langle M^n s, M^{n+1} s, \dots \rangle \models \varphi$ if and only if $n \bmod 288 \in X$. It only remains to check whether $0 \in F$. Hence we have a decision procedure that is in **EXPSpace** in $\|\mathcal{I}\|$.

If, on the other hand, γ is not a root of unity, then we proceed by replacing each **R** and **U** operator in φ with a bounded one. Suppose $\varphi_1 \mathbf{U} \varphi_2$ is a subformula of φ . Using Theorem 11 we can compute an upper bound B on return time $T(\varphi_2)$ of φ_2 with respect to $\langle s, Ms, M^2 s, \dots \rangle$. We then simply replace $\varphi_1 \mathbf{U} \varphi_2$ in φ with $\varphi_1 \mathbf{U}^{\leq B} \varphi_2$ (“ φ_1 remains true until φ_2 is true, and φ_2 becomes true within the first B steps”), with the justification that at any time step n , if the formula φ_2 remains false for all suffixes $\langle M^{n+\delta} s, M^{n+\delta+1} s, \dots \rangle$, $0 \leq \delta \leq B$, then φ_2 will remain false for all $\langle M^{n+\delta} s, M^{n+\delta+1} s, \dots \rangle$, $\delta \geq 0$. Similarly, for a subformula of the form $\varphi_1 \mathbf{R} \varphi_2$ we first compute bounds B_1 and B_2 on the return times $T(\varphi_1 \wedge \varphi_1)$ and $T(\neg \varphi_2)$, respectively, and set $B = \max\{B_1, B_2\}$. Observe that B is at most the bound stipulated in Theorem 11 on the return time of $\varphi_1 \mathbf{R} \varphi_2$ as the latter has higher temporal operator depth. Finally, we replace $\varphi_1 \mathbf{R} \varphi_2$ with the bounded version $\varphi_1 \mathbf{R}^{\leq B} \varphi_2$ with the semantics that either φ_1 successfully releases φ_2 within the first B steps or φ_2 remains true for the first B steps.

We have now reduced the original problem of checking whether the orbit $\langle s, Ms, M^2 s, \dots \rangle$ satisfies φ to determining whether it satisfies φ' with all operators bounded by at most $2^{2^{\|\mathcal{I}\|^{O(1)}}}$ steps. Moreover, note that our algorithm so far does not involve any manipulation of semialgebraic sets or algebraic numbers. In Listing 1 we describe a simple recursive procedure for determining whether a path satisfies such a formula φ' with only bounded temporal operators starting from a time step n .

To analyse the complexity of our main algorithm, let B be a maximum bound on a temporal operator in φ' (i.e. maximum bound on the return time of a subformula of φ). Observe that during the run of the model-checking algorithm, all calls to the oracle are for time steps $n \leq D(\varphi')B = 2^{2^{\|\mathcal{I}\|^{O(1)}}}$, where $D(\varphi)$ is the temporal operator depth of φ as defined in Section 4.4. Therefore, the total space required by the oracle is $O(\|\mathcal{I}\| + \log(D(\varphi)B)) = 2^{\|\mathcal{I}\|^{O(1)}}$. With respect to the oracle, our algorithm operates in $O(D(\varphi) \cdot \log(D(\varphi)B)) = 2^{\|\mathcal{I}\|^{O(1)}}$ space as it simply maintains at most $D(\varphi)$ -many counters with $D(\varphi)B$ bits. Adding the two space requirements we conclude that our decision procedure lies in **EXPSpace** in $\|\mathcal{I}\|$.

7 Conclusion

We have given an algorithm to model check an LTL formula on the orbit of a linear dynamical system in dimension at most 3. The procedure reduces the LTL Model-Checking Problem to an equivalent bounded model-checking problem, which can be solved directly. The heart of the proof is the effective upper bound, given in Theorem 11, of the so-called return time of an LTL formula on a given orbit. Establishing this bound requires the use of several number-theoretic tools. As we have noted in the introduction, there are formidable obstacles to generalising this result to matrices of higher dimensions, since the LTL Model-Checking Problem generalises numerous longstanding open decision problems on linear dynamical systems. Another direction for further work is to consider the problem of model checking MSO, i.e., to generalise the logic. Here we plan to explore connections with the respective frameworks of Semenov [22] and Rabinovitch [20] on decidable extensions of MSO with almost periodic predicates. Finally, in this work we have considered the unique orbit determined by a fixed starting point. But many situations ask to quantify over different orbits, e.g., one could ask whether there is a neighbourhood of a given point such that all orbits starting in the neighbourhood satisfy a given LTL formula – see [4] and [1] for work in this direction.

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.
- 2 S. Almagor, J. Ouaknine, and J. Worrell. The Polytope-Collision Problem. In *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10–14, 2017, Warsaw, Poland*, pages 24:1–24:14, 2017.
- 3 Shaull Almagor, Joël Ouaknine, and James Worrell. The Semialgebraic Orbit Problem. In *36th International Symposium on Theoretical Aspects of Computer Science, STACS 2019, March 13–16, 2019, Berlin, Germany*, pages 6:1–6:15, 2019.
- 4 Shaull Almagor, Joël Ouaknine, and James Worrell. First-order orbit queries. *Theory Comput Syst*, 2020.
- 5 S. Basu, R. Pollack, and M-F. Roy. *Algorithms in real algebraic geometry*, volume 20033. Springer, 2005.
- 6 V. Chonev, J. Ouaknine, and J. Worrell. The Orbit Problem in higher dimensions. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 941–950. ACM, 2013.
- 7 V. Chonev, J. Ouaknine, and J. Worrell. The Polyhedron-Hitting Problem. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 940–956. SIAM, 2015.
- 8 V. Chonev, J. Ouaknine, and J. Worrell. On the complexity of the Orbit Problem. *J. ACM*, 63(3):23:1–23:18, 2016.
- 9 H. Cohen. *A course in computational algebraic number theory*, volume 138. Springer Science & Business Media, 2013.
- 10 V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. Skolem’s Problem – on the border between decidability and undecidability. Technical Report 683, Turku Centre for Computer Science, 2005.
- 11 M. A Harrison. Lectures on linear sequential machines. Technical report, DTIC Document, 1969.
- 12 R. Kannan and R. J. Lipton. The Orbit Problem is decidable. In *Proceedings of the twelfth annual ACM symposium on Theory of computing*, pages 252–261. ACM, 1980.
- 13 R. Kannan and R. J. Lipton. Polynomial-time algorithm for the Orbit Problem. *Journal of the ACM (JACM)*, 33(4):808–821, 1986.

- 14 Zachary Kincaid, John Cyphert, Jason Breck, and Thomas W. Reps. Non-linear reasoning for invariant synthesis. *Proc. ACM Program. Lang.*, 2(POPL):54:1–54:33, 2018.
- 15 Martin Leucker and Christian Schallhart. A brief account of runtime verification. *J. Log. Algebr. Program.*, 78(5):293–303, 2009.
- 16 Nicolas Markey and Philippe Schnoebelen. Model checking a path. In *CONCUR 2003 - Concurrency Theory, 14th International Conference, Proceedings*, volume 2761 of *Lecture Notes in Computer Science*, pages 248–262. Springer, 2003.
- 17 M. Mignotte. Some useful bounds. In *Computer algebra*, pages 259–263. Springer, 1983.
- 18 M. Mignotte, T. Shorey, and R. Tijdeman. The distance between terms of an algebraic recurrence sequence. *J. für die reine und angewandte Math.*, 349, 1984.
- 19 J. Ouaknine and J. Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In *International Colloquium on Automata, Languages, and Programming*, pages 330–341. Springer, 2014.
- 20 Alexander Rabinovich and Wolfgang Thomas. Decidable theories of the ordering of natural numbers with unary predicates. In Zoltán Ésik, editor, *Computer Science Logic*, pages 562–574, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 21 J. Renegar. A faster PSPACE algorithm for deciding the existential theory of the reals. In *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*, pages 291–295, 1988.
- 22 Alexei Semenov. *Decidability of monadic theories*, volume 176, pages 162–175. Springer Berlin Heidelberg, April 2006.
- 23 T. Tao. *Structure and randomness: pages from year one of a mathematical blog*. American Mathematical Soc., 2008.
- 24 N. K. Vereshchagin. Occurrence of zero in a linear recursive sequence. *Mathematical notes of the Academy of Sciences of the USSR*, 38(2):609–615, 1985.