

Algebraic Model Checking for Discrete Linear Dynamical Systems

Florian Luca^{1*}, Joël Ouaknine^{2**}, and James Worrell³

¹ School of Mathematics, University of the Witwatersrand, Johannesburg, South Africa

² Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany

`joel@mpi-sws.org`

³ Department of Computer Science, Oxford University, UK

Abstract. Model checking infinite-state systems is one of the central challenges in automated verification. In this survey we focus on an important and fundamental subclass of infinite-state systems, namely discrete linear dynamical systems. While such systems are ubiquitous in mathematics, physics, engineering, etc., in the present context our motivation stems from their relevance to the formal analysis and verification of program loops, weighted automata, hybrid systems, and control systems, amongst many others. Our main object of study is the problem of model checking temporal properties on the infinite orbit of a linear dynamical system, and our principal contribution is to show that for a rich class of properties this problem can be reduced to certain classical decision problems on linear recurrence sequences, notably the Skolem Problem. This leads us to discuss recent advances on the latter and to highlight the prospects for further progress on charting the algorithmic landscape of linear recurrence sequences and linear dynamical systems.

Keywords: Discrete Linear Dynamical Systems · Linear Recurrence Sequences · Model Checking · Orbit Problem · Skolem Problem

1 Introduction

Dynamical systems are a fundamental modelling paradigm in many branches of science, and have been the subject of extensive research for many decades. A (*rational*) *discrete linear dynamical system (LDS)* in ambient space \mathbb{R}^d is given by a square $d \times d$ matrix M with rational entries, together with a starting point $x \in$

* Also affiliated with: the Research Group in Algebraic Structures and Applications, King Abdulaziz University, Jeddah, Saudi Arabia; the Centro de Ciencias Matemáticas UNAM, Morelia, Mexico; and the Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany.

** Also affiliated with Keble College, Oxford as `emmy.network` Fellow, and supported by DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>).

\mathbb{Q}^d . The *orbit* of (M, x) is the infinite trajectory $\mathcal{O}(M, x) := \langle x, Mx, M^2x, \dots \rangle$. An example of a two-dimensional LDS is given in Figure 1. A central concern in the computational theory of dynamical systems is the task of devising algorithms enabling one to decide various kinds of assertions on dynamical-system orbits.

$$M \stackrel{\text{def}}{=} \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix} \quad x \stackrel{\text{def}}{=} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$$

Fig. 1. A two-dimensional discrete linear dynamical system.

One of the most natural and fundamental computational questions concerning linear dynamical systems is the *Point-to-Point Reachability Problem*, also known as the *Kannan-Lipton Orbit Problem*: given a d -dimensional LDS (M, x) together with a point target $y \in \mathbb{Q}^d$, does the orbit of the LDS ever hit the target? The decidability of this question was settled affirmatively in the 1980s in the seminal work of Kannan and Lipton [28,29]. In fact, Kannan and Lipton showed that this problem is solvable in polynomial time, answering an earlier open problem of Harrison from the 1960s on reachability for linear sequential machines [26].

Interestingly, one of Kannan and Lipton’s motivations was to propose a line of attack to the well-known *Skolem Problem*, which had itself been famously open since the 1930s. The Skolem Problem remains unsolved to this day, although substantial advances have recently been made—more on this shortly. Phrased in the language of linear dynamical systems, the Skolem Problem asks whether it is decidable, given (M, x) as above, together with a $(d - 1)$ -dimensional subspace H of \mathbb{R}^d , to determine if the orbit of (M, x) ever hits H . Kannan and Lipton suggested that, in ambient space \mathbb{R}^d of arbitrary dimension, the problem of hitting a low-dimensional subspace might be decidable. Indeed, this was eventually substantiated by Chonev *et al.* for linear subspaces of dimension at most 3 [17,19].

Subsequent research focussed on the decidability of hitting targets of increasing complexity, such as half-spaces [25,33,37,36,38], polytopes [42,18,3], and semialgebraic sets [4,5]. It is also worth noting that discrete linear dynamical systems can equivalently be viewed as linear (or affine) simple, branching-free while loops, where reachability corresponds to loop termination. There is a voluminous literature on the topic, albeit largely focussing on heuristics and semi-algorithms (via spectral methods or the synthesis of ranking functions), rather than exact decidability results. Relevant papers include [44,14,7,8,13,16,21,40,41,27,9,6]. Several of these approaches have moreover been implemented in software verification tools, such as Microsoft’s Terminator [22,23].

In recent years, motivated in part by verification problems for stochastic systems and linear loops, researchers have begun investigating more sophisticated specification formalisms than mere reachability: for example, the paper [1] studies approximate LTL model checking of Markov chains (which themselves

can be viewed as particular kinds of linear dynamical systems), whereas [32] focuses on LTL model checking of low-dimensional linear dynamical systems with semialgebraic predicates.⁴ In [2], the authors solve the semialgebraic model-checking problem for diagonalisable linear dynamical systems in arbitrary dimension against prefix-independent MSO⁵ properties, whereas [31] investigates semialgebraic MSO model checking of linear dynamical systems in which the dimensions of predicates are constrained. For a comprehensive survey of the state of the art on model checking for linear dynamical systems, we refer the reader to [30].

There is an intimate connection between linear dynamical systems and linear recurrence sequences. A (*rational*) *linear recurrence sequence (LRS)* $\mathbf{u} = \langle u_n \rangle_{n=0}^{\infty}$ is an infinite sequence of rational numbers satisfying

$$u_{n+d} = c_1 u_{n+d-1} + \cdots + c_{d-1} u_{n+1} + c_d u_n \quad (1)$$

for all $n \in \mathbb{N}$, where the coefficients c_1, \dots, c_d are rational numbers and $c_d \neq 0$. We say that the above recurrence has *order* d . We moreover say that an LRS is *simple* if the characteristic polynomial⁶ of its minimal-order recurrence has no repeated roots. The sequence of Fibonacci numbers $\langle f_n \rangle_{n=0}^{\infty} = \langle 0, 1, 1, 2, 3, 5, \dots \rangle$, which obeys the recurrence $f_{n+2} = f_{n+1} + f_n$, is perhaps the most emblematic LRS, and also happens to be simple. It is a straightforward exercise to show that the orbit $\langle x, Mx, M^2x, \dots \rangle$ of the LDS from Figure 1 consists precisely of successive pairs of consecutive Fibonacci numbers:

$$\langle x, Mx, M^2x, \dots \rangle = \left\langle \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 2 \\ 1 \end{pmatrix}, \dots \right\rangle = \left\langle \begin{pmatrix} f_1 \\ f_0 \end{pmatrix}, \begin{pmatrix} f_2 \\ f_1 \end{pmatrix}, \begin{pmatrix} f_3 \\ f_2 \end{pmatrix}, \dots \right\rangle. \quad (2)$$

Let us now define the following two bivariate predicates:

$$P(y, z) \stackrel{\text{def}}{=} (y^2 - yz - z^2 - 1 = 0) \quad (3)$$

$$Q(y, z) \stackrel{\text{def}}{=} (y^2 - yz - z^2 + 1 = 0). \quad (4)$$

Identifying P and Q with the respective subsets of \mathbb{R}^2 that they represent, one can straightforwardly show that the orbit of (M, x) visits P precisely at even-valued indices, and Q at odd-valued indices (where the first element of the orbit is understood to have index 0). In other words, the LDS (M, x) satisfies the following LTL specification:

$$P \wedge \neg Q \wedge \mathbf{G}(P \Rightarrow \mathbf{X}Q) \wedge \mathbf{G}(Q \Rightarrow \mathbf{X}P). \quad (5)$$

Of course, the general task of determining algorithmically whether a given LDS (in arbitrary dimension) meets a given specification would appear to be

⁴ Semialgebraic predicates are Boolean combinations of polynomial equalities and inequalities.

⁵ Monadic Second-Order Logic (MSO) is a highly expressive specification formalism that subsumes the vast majority of temporal logics employed in the field of automated verification, such as Linear Temporal Logic (LTL).

⁶ The characteristic polynomial associated with recurrence (1) is $X^d - c_1 X^{d-1} - \dots - c_d$.

highly challenging. The principal goal of this paper is to delineate the extent to which this can be achieved automatically when the predicates are built from *algebraic* sets⁷ and the specification formalism is either MSO, or its prefix-independent fragment. Before stating our key results, we need to take a brief detour through the Skolem landscape.

1.1 Skolem Oracles

The celebrated theorem of Skolem, Mahler, and Lech (see [24]) describes the structure of the set $\{n \in \mathbb{N} : u_n = 0\}$ of zero terms of an LRS as follows:

Theorem 1. *Given a linear recurrence sequence $\mathbf{u} = \langle u_n \rangle_{n=0}^\infty$, its set of zero terms is a semilinear set, i.e., it consists of a union of finitely many full arithmetic progressions,⁸ together with a finite set.*

As shown by Berstel and Mignotte [10], in the above one can effectively extract all of the arithmetic progressions; we refer herein to the corresponding procedure as the ‘Berstel-Mignotte algorithm’. Nevertheless, how to compute the leftover finite set of zeros remains open, and is easily seen to be equivalent to the *Skolem Problem*: given an LRS \mathbf{u} , does \mathbf{u} contain a zero term?

Let us therefore introduce the notion of a *Skolem oracle*: given an LRS $\mathbf{u} = \langle u_n \rangle_{n=0}^\infty$, such an oracle returns the finite set of indices of zeros of \mathbf{u} that do not already belong to some infinite arithmetic progression of zeros. Likewise, a *Simple-Skolem oracle* is a Skolem oracle restricted to simple LRS.

As mentioned earlier, the decidability of the Skolem Problem is a longstanding open question [24,39], with a positive answer for LRS of order at most 4 known since the mid-1980s [43,45]. Very recently, two major conditional advances on the Skolem Problem have been made, achieving decidability subject to certain classical number-theoretic conjectures: in [34], Lipton *et al.* established decidability for LRS of order 5 assuming the *Skolem Conjecture* (also known as the *Exponential Local-Global Principle*); and in [11], Bilu *et al.* showed decidability for simple LRS of arbitrary order, subject to both the Skolem Conjecture and the *p-adic Schanuel Conjecture*. It is interesting to note that in both cases, the procedures in question rely on the conjectures *only* for termination; correctness is unconditional. In fact, these procedures are *certifying algorithms* (in the sense of [35]) in that, upon termination, they produce an independent certificate (or witness) that their output is correct. Such a certificate can be checked algorithmically by a third party with no reliance on any unproven conjectures. The authors of [11] have implemented their algorithm within the SKOLEM tool, available online.⁹

In view of the above, Simple-Skolem oracles *can* be implemented with unconditional correctness, and guaranteed termination subject to the Skolem and

⁷ Algebraic sets correspond to positive Boolean combinations of polynomial equalities.

⁸ A full arithmetic progression is a set of non-negative integers of the form $\{a + bm : m \in \mathbb{N}\}$, with $a, b \in \mathbb{N}$.

⁹ <https://skolem.mpi-sws.org/> .

p -adic Schanuel conjectures. Whether full Skolem oracles can be devised is the subject of active research; at the time of writing, to the best of our knowledge, no putative procedure is even conjectured in the general (non-simple) case.

To illustrate the applicability of Skolem oracles to model checking linear dynamical systems, let us return to our running example involving the LDS (M, x) from Figure 1. Recall predicate $P(y, z)$ from Eqn. (3), and identify it with the polynomial it implicitly represents, namely $P(y, z) = y^2 - yz - z^2 - 1$. In view of Eqn. (2), we can write the orbit of (M, x) as follows:

$$\langle M^n x \rangle_{n=0}^\infty = \left\langle \begin{pmatrix} y_n \\ z_n \end{pmatrix} \right\rangle_{n=0}^\infty = \left\langle \begin{pmatrix} f_{n+1} \\ f_n \end{pmatrix} \right\rangle_{n=0}^\infty,$$

where the reader will recall that $\langle f_n \rangle_{n=0}^\infty$ is the LRS of Fibonacci numbers. Evaluating the polynomial $P(y, z)$ at each point of the orbit therefore yields the sequence $\langle f_{n+1}^2 - f_{n+1}f_n - f_n^2 - 1 \rangle_{n=0}^\infty$. Given that LRS (resp. simple LRS) are closed under addition and multiplication, the resulting sequence is immediately seen to be a (simple) LRS. Therefore the Berstel-Mignotte algorithm, together with a (Simple-)Skolem oracle, enable us to compute the set of zeros of this LRS as a semilinear set. In turn, this set is precisely the sequence of indices at which the predicate P holds, i.e., at which the orbit of the LDS (M, x) visits the set represented by P . Since one-dimensional semilinear sets are ultimately periodic, and since every step along the way was effective (assuming the existence of (Simple-)Skolem oracles), evaluating the predicate P on the orbit of (M, x) gives rise to an *effectively ultimately periodic word*. As already noted, this word is indeed in fact $\langle \text{TRUE}, \text{FALSE}, \text{TRUE}, \text{FALSE}, \text{TRUE}, \text{FALSE}, \text{TRUE}, \dots \rangle$.

One can of course repeat the procedure with the predicate Q , so that both P and Q are effectively ultimately periodic. Since MSO over effectively ultimately periodic words is decidable, we have just outlined a general algorithmic process by which one can decide algebraic MSO specifications (such as (5)) on orbits of linear dynamical systems, assuming the existence of Skolem or Simple-Skolem oracles.

Remark 2. As hinted above, it is a general fact that the sequence of values obtained by evaluating a multivariate polynomial on the successive points of the orbit of an LDS is always an LRS; moreover, whenever the LDS is diagonalisable¹⁰, the corresponding LRS is always simple. We provide sketch justifications of these facts in Sec. 2.

1.2 Main Results

We require one final ingredient in order to state the main contributions of this paper. Fix the ambient space to be \mathbb{R}^d , and define the collection \mathcal{C} of subsets of \mathbb{R}^d to be the smallest set containing all algebraic subsets of \mathbb{R}^d , and which is closed under finite union, finite intersection, and complement. In algebraic

¹⁰ An LDS (M, x) is *diagonalisable* provided the matrix M is diagonalisable over the complex numbers.

geometry, \mathcal{C} is usually referred to as the collection of *constructible* subsets of \mathbb{R}^d . We refer to MSO formulas over predicates from \mathcal{C} as *algebraic MSO*, and the corresponding model-checking problem as *algebraic model checking*.

Our main results are as follows (precise definitions and statements can be found in Sec. 2):

1. The algebraic model-checking problem for LDS is decidable in arbitrary dimension, subject to the existence of a Skolem oracle.
2. The algebraic model-checking problem for diagonalisable LDS is decidable in arbitrary dimension, subject to the existence of a Simple-Skolem oracle.
 - As an immediate corollary, decidability holds subject to the Skolem and p -adic Schanuel conjectures; moreover, correctness of the model-checking procedure is unconditional, and independent correctness certificates can always be produced upon termination.
3. The algebraic model-checking problem for LDS against prefix-independent specifications is (unconditionally) decidable.

Item 3 above follows from the fact that prefix-independent assertions depend only upon the ultimately periodic components of predicates (see, e.g., [2]), and the latter can be effectively extracted via the Berstel-Mignotte algorithm.

Three further remarks are in order: (i) in ambient space \mathbb{R}^3 , algebraic and even semialgebraic model checking for LDS become unconditionally decidable; this follows immediately from the results of [31], since every predicate in \mathbb{R}^3 belongs to a 3-dimensional subspace (namely \mathbb{R}^3). However: (ii) in ambient space \mathbb{R}^4 , unconditional decidability of algebraic model checking is not known to hold even for diagonalisable LDS, as one can establish hardness for the Skolem Problem at order 5; see [18] for details. (iii) For simplicity, all our results are stated in terms of *rational* linear dynamical systems, living in ambient space \mathbb{R}^d . Nevertheless, it is a straightforward corollary that we can extend our entire framework to *complex-algebraic*¹¹ linear dynamical systems, replacing the ambient space \mathbb{R}^d by \mathbb{C}^d , \mathbb{Q}^d by $\overline{\mathbb{Q}}^d$, and real constructible sets by complex constructible sets.¹² As we sketch in Sec. 2, our main results (as listed above) carry over easily to this more general complex setting. In this extension, it is noteworthy that our Skolem oracles however remain unchanged, i.e., are maintained to apply only to *rational* (rather than complex-algebraic) linear recurrence sequences.

Lastly, it is interesting to note that the algebraic model-checking problem for LDS subsumes not only the original Point-to-Point Reachability Problem, but also the Subspace Reachability Problem suggested by Kannan and Lipton [29] (along with its affine variants), as well as reachability for the *glued vector spaces* of [20].

¹¹ We are referring here to the field of complex algebraic numbers, denoted $\overline{\mathbb{Q}}$.

¹² Complex constructible sets play a central rôle in algebraic geometry; moreover, since the first-order theory of algebraically closed fields admits quantifier elimination, the constructible subsets of \mathbb{C}^d are exactly the subsets of \mathbb{C}^d that are first-order definable over \mathbb{C} .

In the next section, we present a slightly more formal treatment of our framework and results, along with justifications for some of our unsupported assertions. Section 3 concludes with a brief summary and directions for further research.

2 Algebraic Model Checking

Throughout this section, we assume familiarity with the elementary theory of linear recurrence sequences as well as the rudiments of Monadic Second-Order Logic (MSO); there are many excellent references for both topics, such as [24] and [12].

Let us work in fixed ambient space \mathbb{R}^d , and consider a d -dimensional LDS (M, x) (i.e., $M \in \mathbb{Q}^{d \times d}$ and $x \in \mathbb{Q}^d$). Recall that the orbit $\mathcal{O} = \mathcal{O}(M, x)$ of our LDS is the infinite sequence $\langle x, Mx, M^2x, \dots \rangle$ in \mathbb{Q}^d . Let us write $\mathcal{O}[n]$ for the n th term of the orbit, and, for $1 \leq i \leq d$, $\mathcal{O}[n]_i$ for the i th entry of the point $\mathcal{O}[n] \in \mathbb{Q}^d$.

Lemma 3. *Let (M, x) be as above. For any fixed $i \in \{1, \dots, d\}$, the sequence $\langle \mathcal{O}[n]_i \rangle_{n=0}^\infty$ is an LRS whose characteristic polynomial divides the minimal polynomial of M .*

Proof. The fact that the minimal polynomial of M is associated with a recurrence satisfied by the sequence $\langle \mathcal{O}[n]_i \rangle_{n=0}^\infty$ is a straightforward linear-algebraic calculation. It follows that sequence $\langle \mathcal{O}[n]_i \rangle_{n=0}^\infty$ is indeed an LRS having characteristic roots among the eigenvalues of M , with the multiplicity of each characteristic root at most the algebraic multiplicity of the corresponding eigenvalue. The result then immediately follows. \square

Corollary 4. *Let (M, x) be as above. If M is diagonalisable, then for each fixed $i \in \{1, \dots, d\}$, the LRS $\langle \mathcal{O}[n]_i \rangle_{n=0}^\infty$ is simple.*

Proof. This follows immediately from the well-known fact that a square matrix M is diagonalisable if and only if its minimal polynomial is a product of distinct linear factors (over \mathbb{C}). \square

A set $A \subseteq \mathbb{R}^d$ is *algebraic* if A can be written as a finite positive Boolean combination of polynomial equalities, where all polynomials involved have integer coefficients. The collection $\mathcal{C} \subseteq 2^{\mathbb{R}^d}$ of *constructible* subsets of \mathbb{R}^d is the smallest set that includes all algebraic sets and is closed under finite Boolean combinations (including complementation). Any constructible set $C \in \mathcal{C}$ can therefore be represented in conjunctive normal form, i.e., as an expression of the form $C = \bigcap_{i=1}^a \bigcup_{j=1}^b B_{i,j}$, where each $B_{i,j}$ is either an algebraic set or the complement of one.

Let $\mathcal{C} = \{C_1, \dots, C_\ell\}$ be a finite list of constructible sets (not necessarily disjoint), giving rise to an alphabet $\Sigma \stackrel{\text{def}}{=} 2^{\mathcal{C}}$. The orbit \mathcal{O} of our LDS (M, x) then naturally gives rise to an infinite *characteristic word* $w \in \Sigma^\omega$, as follows: writing $w[n]$ for the n th letter of w , we have $C_i \in w[n]$ iff $M^n x \in C_i$.

Proposition 5. *Characteristic words over constructible predicates are ultimately periodic.*

Proof. Let us write $\mathcal{O} = \langle x_n \rangle_{n=0}^\infty$ to denote the orbit of our LDS (M, x) , and let us further write each x_n as $(x_n^{(1)}, \dots, x_n^{(d)})^T$. Fix a polynomial $f \in \mathbb{Z}[X_1, \dots, X_d]$, and consider the sequence $\langle f(x_n) \rangle_{n=0}^\infty$. Since sums and products of LRS are LRS, by Lemma 3 the sequence $\langle f(x_n) \rangle_{n=0}^\infty$ is an LRS, and by the Skolem-Mahler-Lech theorem its set of zeros is therefore semilinear, hence ultimately periodic. Since pointwise Boolean combinations (including complementation) of ultimately periodic words are ultimately periodic, and taking account of the fact that any constructible set is a finite Boolean combination of algebraic sets, the result immediately follows. \square

We say that a word is *effectively* ultimately periodic if one can compute an integer threshold beyond which the word in question becomes periodic. Stringing everything together:

Corollary 6.

1. *Assume the existence of a Skolem oracle. Then characteristic words over constructible predicates are effectively ultimately periodic.*
2. *Suppose (M, x) is a diagonalisable LDS, and assume the existence of a Simple-Skolem oracle. Then any characteristic word associated with the orbit of (M, x) over constructible predicates is effectively ultimately periodic.*

Proof. The first item follows directly from Prop. 5, together with the Berstel-Mignotte algorithm. So does the second item, invoking in addition Cor. 4, and taking account of the fact that simple LRS are closed under sums and products, along with the fact that constant sequences are themselves simple LRS. \square

Let (M, x) and \mathcal{C} be as above, and let φ be an MSO formula with atomic predicates drawn from \mathcal{C} —we refer to such formulas as *algebraic MSO specifications*. The question of whether the characteristic word w associated with the orbit of (M, x) satisfies specification φ —which is usually written as $(M, x) \models \varphi$ —is the *algebraic model-checking problem* for discrete linear dynamical systems.

In addition, we say that φ is *prefix-independent* if the infinite words that satisfy it are closed under the operations of insertion and deletion of finitely many letters. Prefix-independent properties can be used to describe asymptotic behaviours (e.g., “does the orbit enter C_1 infinitely often?”), but not reachability.

We are now ready to formally state our main results:

Theorem 7.

1. *The algebraic model-checking problem for LDS is decidable, subject to the existence of a Skolem oracle.*
2. *The algebraic model-checking problem for diagonalisable LDS is decidable, subject to the existence of a Simple-Skolem oracle.*
3. *The algebraic model-checking problem for LDS against prefix-independent specifications is (unconditionally) decidable.*

Proof. Item 1 is an immediate consequence of Büchi’s seminal work [15] establishing the decidability of MSO, together with Cor. 6(1) and the observation that effectively ultimately periodic predicates can be algorithmically translated to ordinary MSO by encoding the predicates as formulas.

The same holds for Item 2, invoking Cor. 6(2) in lieu of Cor. 6(1).

Finally, Item 3 follows from the fact that prefix-independent assertions depend only upon the ultimately periodic components of predicates (see, e.g., [2]), and the latter can be effectively extracted via the Berstel-Mignotte algorithm. \square

As noted earlier, since Simple-Skolem oracles can be implemented into provably correct certifying procedures which terminate subject to classical number-theoretic conjectures [11], we have:

Corollary 8. *The algebraic model-checking problem for diagonalisable LDS is decidable, assuming the Skolem Conjecture and the p -adic Schanuel Conjecture. Moreover, correctness of the attendant procedure is unconditional, and independent correctness certificates can be produced upon termination.*

Corollary 8 is arguably our most consequential and interesting contribution. Given that promising experimental results are reported in [11] regarding the implementation of a Simple-Skolem algorithm, it appears rather plausible that one could likewise build an efficient and practical model checker for diagonalisable LDS against algebraic specifications.

Finally, let us record, as already noted in the Introduction, that both Thm. 7 and Cor. 8 can be extended *mutatis mutandis* to complex-algebraic linear dynamical systems, whilst only invoking Skolem oracles for rational linear recurrence sequences. We sketch a short justification of this claim below.

Let (M, x) be a complex-algebraic LDS in ambient space \mathbb{C}^d , i.e., $M \in \overline{\mathbb{Q}}^{d \times d}$ and $x \in \overline{\mathbb{Q}}^d$, and let $f \in \overline{\mathbb{Q}}[X_1, \dots, X_d]$ be an arbitrary polynomial with complex-algebraic coefficients. Let $Z \stackrel{\text{def}}{=} \{n \in \mathbb{N} : f(M^n x) = 0\}$ be the set of indices at which the orbit of (M, x) lies in the algebraic set $f^{-1}(0)$. The crux of our claim boils down to the following lemma:

Lemma 9. *Let (M, x) , f , and Z be as above. Then one can effectively construct a rational LRS $\mathbf{u} = \langle u_n \rangle_{n=0}^\infty$ such that Z is precisely the zero set of \mathbf{u} . Moreover, if M is diagonalisable, then \mathbf{u} is a simple rational LRS.*

Proof. Let $\mathbb{K} \subset \overline{\mathbb{Q}}$ be the smallest number field containing all the entries of M and x , as well as all the coefficients of f . As in the proof of Prop. 5, one easily shows that the sequence $\mathbf{v} = \langle f(M^n(x)) \rangle_{n=0}^\infty$ is an LRS lying entirely in \mathbb{K} .

Recall the notion of *norm* $\mathcal{N}_{\mathbb{K}} : \mathbb{K} \rightarrow \mathbb{Q}$ from algebraic number theory, defined as $\mathcal{N}_{\mathbb{K}}(x) = \prod_{\sigma : \mathbb{K} \rightarrow \mathbb{C}} \sigma(x)$, where the product is indexed by the collection of field embeddings of \mathbb{K} into \mathbb{C} . Since \mathbf{v} is an LRS, then so is $\sigma(\mathbf{v}) = \langle \sigma(v_n) \rangle_{n=0}^\infty$ for any

embedding $\sigma : \mathbb{K} \rightarrow \mathbb{C}$, and moreover such applications will also preserve simplicity of LRS (this can be seen by inspecting the effect of σ on the exponential-polynomial closed-form representation of \mathbf{v}). Write $\mathbf{u} \stackrel{\text{def}}{=} \mathcal{N}_{\mathbb{K}}(\mathbf{v}) = \langle \mathcal{N}_{\mathbb{K}}(v_n) \rangle_{n=0}^{\infty}$. Since products of (simple) LRS are again (simple) LRS, we have that \mathbf{u} is a (simple) LRS lying entirely in \mathbb{Q} . Moreover, since field embeddings fix 0, \mathbf{u} and \mathbf{v} have precisely the same zero set, as required. \square

3 Conclusion

This paper has demonstrated that solving the Skolem Problem is key to model checking a rich class of algebraic properties on linear dynamical systems. We have formulated our results in terms of the existence of oracles for the Skolem Problem and a special subcase, the Simple-Skolem Problem. Implementing such oracles is the subject of ongoing research. As remarked earlier, we have recently devised an algorithm for the Simple-Skolem Problem whose output comes with an easily checkable correctness certificate (namely a set of zeros of the given sequence and a certificate that the remaining terms of the sequence are non-zero) and which terminates subject to certain classical number-theoretic conjectures. We are currently investigating whether a similar approach can be devised in the general (non-simple) case.

In this note we have concentrated on logical specifications built over constructible predicates, that is, those that are defined by Boolean combinations of polynomial equalities. In many applications, one is also interested in the more general class of *semialgebraic* predicates, that is, those defined by Boolean combinations of polynomial *inequalities*. The task of model checking MSO formulas over such predicates appears vastly more complex. Already the question of whether the orbit of an LDS remains within a prescribed halfspace—or equivalently whether all terms of an LRS are non-negative, known as the *Positivity Problem*—is highly challenging: decidability is known only for sequences of order at most 5, whereas for sequences of order 6 a solution to the Positivity Problem would entail major breakthroughs in number theory [37,39].

References

1. Agrawal, M., Akshay, S., Genest, B., Thiagarajan, P.S.: Approximate verification of the symbolic dynamics of Markov chains. *J. ACM* **62**(1), 2:1–2:34 (2015)
2. Almagor, S., Karimov, T., Kelmendi, E., Ouaknine, J., Worrell, J.: Deciding ω -regular properties on linear recurrence sequences. *Proc. ACM Program. Lang.* **5**(POPL), 1–24 (2021)
3. Almagor, S., Ouaknine, J., Worrell, J.: The polytope-collision problem. In: 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017. *LIPIcs*, vol. 80, pp. 24:1–24:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017)
4. Almagor, S., Ouaknine, J., Worrell, J.: The semialgebraic Orbit Problem. In: 36th International Symposium on Theoretical Aspects of Computer Science, STACS

2019. LIPIcs, vol. 126, pp. 6:1–6:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
5. Almagor, S., Ouaknine, J., Worrell, J.: First-order orbit queries. *Theory Comput. Syst.* **65**(4), 638–661 (2021)
 6. Ben-Amram, A.M., Doménech, J.J., Genaim, S.: Multiphase-linear ranking functions and their relation to recurrent sets. In: *Static Analysis - 26th International Symposium, SAS 2019. Lecture Notes in Computer Science*, vol. 11822, pp. 459–480. Springer (2019)
 7. Ben-Amram, A.M., Genaim, S.: On the linear ranking problem for integer linear-constraint loops. In: *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13*. pp. 51–62. ACM (2013)
 8. Ben-Amram, A.M., Genaim, S.: Ranking functions for linear-constraint loops. *J. ACM* **61**(4), 26:1–26:55 (2014)
 9. Ben-Amram, A.M., Genaim, S.: On multiphase-linear ranking functions. In: *Computer Aided Verification - 29th International Conference, CAV 2017. Lecture Notes in Computer Science*, vol. 10427, pp. 601–620. Springer (2017)
 10. Berstel, J., Mignotte, M.: Deux propriétés décidables des suites récurrentes linéaires. *Bull. Soc. math. France* **104**, 175–184 (1976)
 11. Bilu, Y., Luca, F., Nieuwveld, J., Ouaknine, J., Purser, D., Worrell, J.: Skolem meets Schanuel. In: Szeider, S., Ganian, R., Silva, A. (eds.) *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22–26, 2022, Vienna, Austria. LIPIcs*, vol. 241, pp. 62:1–62:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022)
 12. Börger, E., Grädel, E., Gurevich, Y.: *The Classical Decision Problem. Perspectives in Mathematical Logic*, Springer (1997)
 13. Bradley, A.R., Manna, Z., Sipma, H.B.: Termination analysis of integer linear loops. In: *CONCUR 2005 - Concurrency Theory, 16th International Conference, CONCUR 2005. Lecture Notes in Computer Science*, vol. 3653, pp. 488–502. Springer (2005)
 14. Braverman, M.: Termination of integer linear programs. In: *Computer Aided Verification, 18th International Conference, CAV 2006. Lecture Notes in Computer Science*, vol. 4144, pp. 372–385. Springer (2006)
 15. Büchi, J.R.: Weak second order arithmetic and finite automata. *Zeit. für Math. Logik und Grund. der Math.* **6**(1-6), 66–92 (1960)
 16. Chen, H.Y., Flur, S., Mukhopadhyay, S.: Termination proofs for linear simple loops. *Int. J. Softw. Tools Technol. Transf.* **17**(1), 47–57 (2015)
 17. Chonev, V., Ouaknine, J., Worrell, J.: The Orbit Problem in higher dimensions. In: *Symposium on Theory of Computing Conference, STOC'13*. pp. 941–950. ACM (2013)
 18. Chonev, V., Ouaknine, J., Worrell, J.: The polyhedron-hitting problem. In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*. pp. 940–956. SIAM (2015)
 19. Chonev, V., Ouaknine, J., Worrell, J.: On the complexity of the Orbit Problem. *J. ACM* **63**(3), 23:1–23:18 (2016)
 20. Colcombet, T., Petrişan, D.: Automata in the category of glued vector spaces. In: Larsen, K.G., Bodlaender, H.L., Raskin, J. (eds.) *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21–25, 2017 - Aalborg, Denmark. LIPIcs*, vol. 83, pp. 52:1–52:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2017)

21. Colón, M., Sipma, H.: Synthesis of linear ranking functions. In: Tools and Algorithms for the Construction and Analysis of Systems, 7th International Conference, TACAS 2001. Lecture Notes in Computer Science, vol. 2031, pp. 67–81. Springer (2001)
22. Cook, B., Podelski, A., Rybalchenko, A.: Termination proofs for systems code. In: Proceedings of the ACM SIGPLAN 2006 Conference on Programming Language Design and Implementation. pp. 415–426. ACM (2006)
23. Cook, B., Podelski, A., Rybalchenko, A.: Terminator: Beyond safety. In: Computer Aided Verification, 18th International Conference, CAV 2006. Lecture Notes in Computer Science, vol. 4144, pp. 415–418. Springer (2006)
24. Everest, G., van der Poorten, A.J., Shparlinski, I.E., Ward, T.: Recurrence Sequences, Mathematical surveys and monographs, vol. 104. American Mathematical Society (2003)
25. Halava, V., Harju, T., Hirvensalo, M.: Positivity of second order linear recurrent sequences. *Discret. Appl. Math.* **154**(3), 447–451 (2006)
26. Harrison, M.A.: Lectures on Linear Sequential Machines. Academic Press, New York (1969)
27. Hosseini, M., Ouaknine, J., Worrell, J.: Termination of linear loops over the integers. In: 46th International Colloquium on Automata, Languages, and Programming, ICALP 2019. LIPIcs, vol. 132, pp. 118:1–118:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2019)
28. Kannan, R., Lipton, R.J.: The Orbit Problem is decidable. In: Proceedings of the 12th Annual ACM Symposium on Theory of Computing 1980. pp. 252–261. ACM (1980)
29. Kannan, R., Lipton, R.J.: Polynomial-time algorithm for the Orbit Problem. *J. ACM* **33**(4), 808–821 (1986)
30. Karimov, T., Kelmendi, E., Ouaknine, J., Worrell, J.: What’s decidable about discrete linear dynamical systems? *CoRR* **abs/2206.11412** (2022)
31. Karimov, T., Lefauchaux, E., Ouaknine, J., Purser, D., Varonka, A., Whiteland, M.A., Worrell, J.: What’s decidable about linear loops? *Proc. ACM Program. Lang.* **6**(POPL), 1–25 (2022)
32. Karimov, T., Ouaknine, J., Worrell, J.: On LTL model checking for low-dimensional discrete linear dynamical systems. In: 45th International Symposium on Mathematical Foundations of Computer Science, MFCS 2020. LIPIcs, vol. 170, pp. 54:1–54:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2020)
33. Laohakosol, V., Tangsupphathawat, P.: Positivity of third order linear recurrence sequences. *Discret. Appl. Math.* **157**(15), 3239–3248 (2009)
34. Lipton, R.J., Luca, F., Nieuwveld, J., Ouaknine, J., Worrell, D.P.J.: On the Skolem Problem and the Skolem Conjecture. In: 37th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2022, Haifa, Israel, August 2 - August 5, 2022. ACM (2022)
35. McConnell, R.M., Mehlhorn, K., Näher, S., Schweitzer, P.: Certifying algorithms. *Comput. Sci. Rev.* **5**(2), 119–161 (2011)
36. Ouaknine, J., Worrell, J.: On the positivity problem for simple linear recurrence sequences. In: Automata, Languages, and Programming. pp. 318–329. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
37. Ouaknine, J., Worrell, J.: Positivity problems for low-order linear recurrence sequences. In: Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014. pp. 366–379. SIAM (2014)

38. Ouaknine, J., Worrell, J.: Ultimate positivity is decidable for simple linear recurrence sequences. In: Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014. Lecture Notes in Computer Science, vol. 8573, pp. 330–341. Springer (2014)
39. Ouaknine, J., Worrell, J.: On linear recurrence sequences and loop termination. ACM SIGLOG News **2**(2), 4–13 (2015)
40. Podelski, A., Rybalchenko, A.: A complete method for the synthesis of linear ranking functions. In: Verification, Model Checking, and Abstract Interpretation, 5th International Conference, VMCAI. Lecture Notes in Computer Science, vol. 2937, pp. 239–251. Springer (2004)
41. Podelski, A., Rybalchenko, A.: Transition invariants. In: 19th IEEE Symposium on Logic in Computer Science (LICS 2004). pp. 32–41. IEEE Computer Society (2004)
42. Tarasov, S.P., Vyalyi, M.N.: Orbits of linear maps and regular languages. In: Computer Science - Theory and Applications - 6th International Computer Science Symposium in Russia, CSR 2011. Lecture Notes in Computer Science, vol. 6651, pp. 305–316. Springer (2011)
43. Tijdeman, R., Mignotte, M., Shorey, T.N.: The distance between terms of an algebraic recurrence sequence. Journal für die reine und angewandte Mathematik **349**, 63–76 (1984)
44. Tiwari, A.: Termination of linear programs. In: Computer Aided Verification, 16th International Conference, CAV 2004. Lecture Notes in Computer Science, vol. 3114, pp. 70–82. Springer (2004)
45. Vereshchagin, N.: The problem of appearance of a zero in a linear recurrence sequence. Mat. Zametki **38**(2), 609–615 (1985)