O-Minimal Invariants for Discrete-Time Dynamical Systems

SHAULL ALMAGOR, Computer Science Department, Technion

DMITRY CHISTIKOV, Centre for Discrete Mathematics and its Applications (DIMAP) and Department of Computer Science, University of Warwick

JOËL OUAKNINE, Max Planck Institute for Software Systems

JAMES WORRELL, Department of Computer Science, Oxford University

9

Termination analysis of linear loops plays a key rôle in several areas of computer science, including program verification and abstract interpretation. Already for the simplest variants of linear loops the question of termination relates to deep open problems in number theory, such as the decidability of the Skolem and Positivity Problems for linear recurrence sequences, or equivalently reachability questions for discrete-time linear dynamical systems. In this article, we introduce the class of *o-minimal invariants*, which is broader than any previously considered, and study the decidability of the existence and algorithmic synthesis of such invariants as certificates of non-termination for linear loops equipped with a large class of halting conditions. We establish two main decidability results, one of them conditional on Schanuel's conjecture is transcendental number theory.

CCS Concepts: • **Theory of computation** → **Logic and verification**; *Logic and verification*; Finite Model Theory; • **Computing methodologies** → **Algebraic algorithms**; • **Mathematics of computing** → *Continuous mathematics*; Continuous functions; • **Software and its engineering** → *Formal software verification*;

Additional Key Words and Phrases: Invariants, linear loops, linear dynamical systems, non-termination, o-minimality

ACM Reference format:

Shaull Almagor, Dmitry Chistikov, Joël Ouaknine, and James Worrell. 2022. O-Minimal Invariants for Discrete-Time Dynamical Systems. *ACM Trans. Comput. Logic* 23, 2, Article 9 (January 2022), 20 pages.

<https://doi.org/10.1145/3501299>

Joël Ouaknine was supported by ERC grant AVS-ISS (648701) and by DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>). Also affiliated with Keble College, Oxford as emmy.network Fellow. James Worrell was supported by EPSRC Fellowship EP/N008197/1. Shaull Almagor has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 837327. Authors' addresses: S. Almagor, Computer Science Department, Technion, Taub Building, Haifa 3200003, Israel; email: shaull@cs.technion.ac.il; D. Chistikov, Centre for Discrete Mathematics and its Applications (DIMAP) and Department of Computer Science, University of Warwick, Coventry CV4 7AL, United Kingdom; email: d.chistikov@warwick.ac.uk; J. Ouaknine, Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany; email: joel@mpi-sws.org; J. Worrell, Department of Computer Science, Oxford University, Parks Road, Oxford OX1 3QD, United Kingdom; email: jbw@cs.ox.ac.uk.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

1529-3785/2022/01-ART9 \$15.00

<https://doi.org/10.1145/3501299>

1 INTRODUCTION

This article is concerned with the existence and algorithmic synthesis of suitable *invariants* for discrete-time **linear dynamical systems (LDS)**. Invariants are one of the most fundamental and useful notions in the quantitative sciences, and within computer science play a central rôle in areas such as program analysis and verification, abstract interpretation, static analysis, and theorem proving. To this day, automated invariant synthesis remains a topic of active research; see, e.g., [22], and particularly Section 8 therein.

In program analysis, invariants are often invaluable tools enabling one to establish various properties of interest. Our focus here is on simple linear loops, of the following form:

$$P : x \leftarrow s; \text{ while } x \notin F \text{ do } x \leftarrow Ax, \quad (1)$$

where x is a d -dimensional column vector of variables, s is a d -dimensional vector of integer, rational, or real numbers, $A \in \mathbb{Q}^{d \times d}$ is a square rational matrix of dimension d , and $F \subseteq \mathbb{R}^d$ represents the halting condition.

Much research has been devoted to the termination analysis of such loops (and variants thereof); see, e.g., [2, 3, 30]. For $S \subseteq \mathbb{R}^d$, we say that P *terminates* on S if it terminates for all initial vectors $s \in S$. One of the earliest and most famous results in this line of work is due to Kannan and Lipton, who showed polynomial-time decidability of termination in the case where S and F are both singleton vectors with rational entries [20, 21]. This work was subsequently extended to instances in which F is a low-dimensional vector space [7, 9] or a low-dimensional polyhedron [8]. Still starting from a fixed initial vector, the case in which the halting set F is a hyperplane is equivalent to the famous Skolem Problem for linear recurrence sequences, whose decidability has been open for many decades [36, Section 3.9], although once again positive results are known in low dimensions [25, 39]. The case in which F is a half-space corresponds to the Positivity Problem for linear recurrence sequences, likewise famously open in general but for which some partial results also exist [27, 28].

Cases in which the starting set S is infinite have also been extensively studied, usually in conjunction with a halting set F consisting of a half-space. For example, decidability of termination for $S = \mathbb{R}^d$, $S = \mathbb{Q}^d$, and $S = \mathbb{Z}^d$ are known [4, 19, 26, 38]. In the vast majority of cases, however, termination is a hard problem (and often undecidable [41]), which has led researchers to turn to semi-algorithms and heuristics. One of the most popular and successful approaches to establishing termination is the use of ranking functions, on which there is a substantial body of work; see, e.g., [2], which includes a broad survey on the subject.

Observe, for a loop P such as that given in Equation (1), that failure to terminate on a set S corresponds to the existence of some vector $s \in S$ from which P loops forever. It is important to note, however, that the absence of a suitable ranking function does not necessarily entail non-termination, owing to the non-completeness of the method. Yet surprisingly, as pointed out in [18], there has been significantly less research in methods seeking to establish *non-termination* than in methods aimed at proving termination. Most existing efforts for the former have focused on the synthesis of appropriate invariants; see, e.g., [10–12, 14–16, 32–34].

In order to make this notion more precise, let us associate with our loop P a *discrete-time LDS* (A, s) . The *orbit* of this dynamical system is the set $\mathcal{O} = \{A^n s \mid n \geq 0\}$. It is clear that P fails to terminate from s iff $\mathcal{O}$ is disjoint from F . A possible method to establish the latter is therefore to exhibit a set $\mathcal{I} \subseteq \mathbb{R}^d$ such that:

- (1) $\mathcal{I}$ contains the initial vector s , i.e., $s \in \mathcal{I}$;
- (2) $\mathcal{I}$ is invariant under A , i.e., $A\mathcal{I} \subseteq \mathcal{I}$; and
- (3) $\mathcal{I}$ is disjoint from F , i.e., $\mathcal{I} \cap F = \emptyset$.

Fig. 1. The orbit O of (A, s) .

Fig. 2. Trajectory rays of O .

Fig. 3. Trajectory cone for O .

Fig. 4. Invariant set for O .

Indeed, the first two conditions ensure that I contains the entire orbit O , from which the desired claim follows thanks to the third condition.

In instances of non-termination, one notes that the orbit O itself is always an invariant meeting the above conditions. However, since in general one does not know how to algorithmically check Condition (3), such an invariant is of little use. One therefore usually first fixes a suitable class of candidate sets for which the above conditions can be mechanically verified, and within that class, one seeks to determine if an invariant can be found. Examples of such classes include polyhedra [12], algebraic sets [33], and semi-algebraic sets [15].

Main contributions. We focus on loops of the form given in Equation (1) above. We introduce the class of *o-minimal invariants*, which to the best of our knowledge is significantly broader than any of the classes previously considered in the context of linear loops. An o-minimal invariant is one that is definable in some o-minimal expansion of the ordered field $\mathbb{R}_{\text{exp}}$ of real numbers with real exponentiation. We also consider two large classes of halting sets, namely, those definable over the ordered field $\mathbb{R}_0$ of real numbers (i.e., *semi-algebraic sets*) and those definable in $\mathbb{R}_{\text{exp}}$.

Given $s \in \mathbb{Q}^d$, $A \in \mathbb{Q}^{d \times d}$, and $F \subseteq \mathbb{R}^d$, our main results are the following: if F is a semi-algebraic set, it is decidable whether there exists an o-minimal invariant I containing s and disjoint from F , and moreover in positive instances such an invariant can be defined explicitly in $\mathbb{R}_0$. For the more general case in which F is $\mathbb{R}_{\text{exp}}$ -definable, assuming Schanuel's conjecture it is decidable whether there exists an o-minimal invariant I containing s and disjoint from F , and moreover in positive instances such an invariant can be defined explicitly in $\mathbb{R}_{\text{exp}}$.

We illustrate below some of the key ideas from our approach. Consider an LDS (A, s) with $A \in \mathbb{Q}^{3 \times 3}$ whose orbit O is depicted in Figure 1. In our example, O spirals outward at some rate ρ_1 in the x, y -plane, and increases along the z -axis at some rate ρ_2 . Intuitively, ρ_1 and ρ_2 are the moduli of the eigenvalues of A .

We now consider a “normalised” version of A , with both moduli set to 1. We then connect every point on the normalised orbit with a *trajectory ray* to its corresponding point on O , while respecting the rates ρ_1 and ρ_2 (see Figure 2). One can observe that the normalised orbit is dense in the unit circle. We prove that *any* o-minimal invariant for (A, s) must in fact eventually contain every trajectory ray for every point on the unit circle; we depict the union of these rays, referred to as the *trajectory cone*, in Figure 3. Finally, we show that any o-minimal invariant must in fact contain some truncation of the trajectory cone from below, starting from some height. That is, there is a uniform bound from which all the rays must belong to the invariant. Moreover, we can now synthesise an $\mathbb{R}_{\text{exp}}$ -definable o-minimal invariant by simply adjoining a finite number of orbit points to the truncated trajectory cone, as depicted in Figure 4.

It is worth emphasising that, whilst in general there cannot exist a smallest o-minimal invariant, the family of truncated cones that we define plays the rôle of a “minimal class”, in the sense that *any* o-minimal invariant must necessarily contain some truncated cone. We make all of these notions precise in the main body of the article.

The works that are closest to ours in the literature are [14–16], which consider the same kind of loops as we do here, but restricted to the case in which the halting set F is always a rational singleton. The authors then exhibit procedures for deciding the existence of semi-algebraic invariants ([15, 16]) and semi-linear invariants [14]. The present article has a considerably broader scope, in that we deal with much wider classes both of invariants and halting sets. From a technical standpoint, the present article correspondingly makes heavy use of model-theoretic and number-theoretic tools that are entirely absent from the above articles.

Specifically, we make use of o-minimality in order to reason about the structure of the proposed invariants, as well as quantifier elimination in the case of semi-algebraic targets. On the number-theoretic front, we heavily rely on Baker’s Theorem in order to obtain decidability results. These tools are needed here, as opposed to [15, 16], since, intuitively, when F is not a singleton the general termination problem is not known to be decidable, and in particular, the “interaction” of the orbit with F is not limited to a single orbit point, but may require reasoning about the asymptotics of the orbit.

2 PRELIMINARIES AND MAIN DEFINITIONS

We write $\mathfrak{R}_0$ for the structure $\langle \mathbb{R}, 0, 1, +, \cdot, < \rangle$, i.e., the ordered field of real numbers with constants 0 and 1. A sentence in the corresponding first-order language can be considered as a quantified Boolean combination of atomic propositions of the form $P(x_1, \dots, x_n) > 0$, where P is a polynomial with integer coefficients and $x_1, \dots, x_n$ are variables. Tarski famously showed that the first-order theory of $\mathfrak{R}_0$ admits quantifier elimination [37] and is therefore decidable. In addition to $\mathfrak{R}_0$, we also consider the structure $\mathfrak{R}_{\text{exp}}$, obtained by expanding $\mathfrak{R}_0$ with the real exponentiation function $x \mapsto e^x$. It is an open question whether the theory of the reals with exponentiation is decidable; however, decidability was established subject to Schanuel’s conjecture by MacIntyre and Wilkie [23]. (Schanuel’s conjecture is a unifying conjecture in transcendental number theory that generalises many of the classical results of that subject.) MacIntyre and Wilkie further showed in [23] that decidability of the theory of $\mathfrak{R}_{\text{exp}}$ implies a weak form of Schanuel’s conjecture.

Let $\mathfrak{R}$ be an expansion of the structure $\mathfrak{R}_0$. A set $S \subseteq \mathbb{R}^d$ is *definable* in $\mathfrak{R}$ if there exists a formula $\varphi(x_1, \dots, x_d)$ in $\mathfrak{R}$ with free variables $x_1, \dots, x_d$ such that $S = \{(c_1, \dots, c_d) \in \mathbb{R}^d \mid \mathfrak{R} \models \varphi(c_1, \dots, c_d)\}$. A function $f: B \rightarrow \mathbb{R}^m$ with $B \subseteq \mathbb{R}^n$ is *definable* in $\mathfrak{R}$ if its graph $\Gamma(f) = \{(x, f(x)) \mid x \in B\} \subseteq \mathbb{R}^{n+m}$ is an $\mathfrak{R}$ -definable set. For $\mathfrak{R} = \mathfrak{R}_0$, the ordered field of real numbers, $\mathfrak{R}_0$ -definable sets (respective functions) are known as *semi-algebraic* sets (respective functions).

Remark 1. Our usage of the terms “definable” and “semi-algebraic” corresponds to “definable without parameters” and “semi-algebraic without parameters” in model theory.

Remark 2. Recall that there is a natural first-order interpretation of the field of complex numbers $\mathbb{C}$ in the field of real numbers $\mathbb{R}$. We shall say that a set $S \subseteq \mathbb{C}^d$ is $\mathfrak{R}$ -*definable* if the image $\{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid x + iy \in S\}$ of S under this interpretation is $\mathfrak{R}$ -definable.

A totally ordered structure $\langle M, <, \dots \rangle$ is said to be *o-minimal* if every definable subset of M is a finite union of intervals. Tarski’s result on quantifier elimination [37] implies that $\mathfrak{R}_0$ is o-minimal. The o-minimality of $\mathfrak{R}_{\text{exp}}$ is due to Wilkie [40] and holds unconditionally. An o-minimal expansion $\mathfrak{R}$ of $\mathfrak{R}_0$ satisfies the following useful properties (see [13] for precise definitions and proofs).

- (1) For an $\mathfrak{R}$ -definable set $S \subseteq \mathbb{R}^d$, its topological closure $\bar{S}$ is also $\mathfrak{R}$ -definable.
- (2) For an $\mathfrak{R}$ -definable function $f: S \rightarrow \mathbb{R}$, the number $\inf \{f(x) \mid x \in S\}$ is $\mathfrak{R}$ -definable (as a singleton set).

- (3) O-minimal structures admit *cell decomposition*: Every $\mathfrak{R}$ -definable set $S \subseteq \mathbb{R}^d$ can be written as a finite union of connected components called *cells*. Moreover, each cell is $\mathfrak{R}$ -definable and homeomorphic to $(0, 1)^m$ for some $m \in \{0, 1, \dots, d\}$ (where for $m = 0$ we have that $(0, 1)^0$ is a single point, namely, $\{\vec{0}\} \subseteq \mathbb{R}^d$). The *dimension* of S is defined as the maximal such m occurring in the cell decomposition of S .
- (4) For an $\mathfrak{R}$ -definable function $f: S \rightarrow \mathbb{R}^m$, the dimension of its graph $\Gamma(f)$ is the same as the dimension of S .

As mentioned above, $\mathfrak{R}_0$ is decidable thanks to its effective quantifier elimination procedure. Equivalently, given a semi-algebraic set, we can effectively compute its cell decomposition. Unfortunately, few more expressive theories are known to be unconditionally decidable. Our decidability result in Theorem 6.2 on invariants definable in $\mathfrak{R}_{\text{exp}}$ is subject to Schanuel's conjecture; somewhat surprisingly, however, we exhibit in Theorem 6.4 an unconditional decidability result.

A *discrete-time* LDS consists of a pair (A, s) , where $A \in \mathbb{Q}^{d \times d}$ and $s \in \mathbb{Q}^d$. Its *orbit* $\mathcal{O}$ is the set $\{A^n s \mid n \in \mathbb{N}\}$. An *invariant* for (A, s) is a set $I \subseteq \mathbb{R}^d$ that contains s and is stable under applications of A , i.e., $AI \subseteq I$. Given a set $F \subseteq \mathbb{R}^d$, we say that the invariant I *avoids* F if the two sets are disjoint. An *o-minimal invariant* is one that is definable in an o-minimal expansion of $\mathfrak{R}_{\text{exp}}$.

3 FROM THE ORBIT TO TRAJECTORY CONES AND RAYS

Let (A, s) be an LDS with $A \in \mathbb{Q}^{d \times d}$ and $s \in \mathbb{Q}^d$. We consider the orbit $\mathcal{O} = \{A^n s \mid n \in \mathbb{N}\}$. Write A in Jordan form as $A = PJP^{-1}$ where P is an invertible matrix, and J is a block diagonal matrix of the form $J = \text{diag}(B_1, \dots, B_k)$, where for every $1 \leq i \leq k$, $B_i \in \mathbb{C}^{d_i \times d_i}$ is a Jordan block corresponding to an eigenvalue λ_i :

$$B_i = \begin{pmatrix} \lambda_i & 1 & & \\ & \ddots & \ddots & \\ & & \ddots & 1 \\ & & & \lambda_i \end{pmatrix}.$$

To reflect the block structure of J , we often range over $\{1, \dots, d\}$ via a pair (i, j) , with $1 \leq i \leq k$ and $1 \leq j \leq d_i$, which denotes the index corresponding to row j in block i ; we refer to this notation as *block-row indexing*.

Henceforth, we assume that for all $1 \leq i \leq k$ we have that $\lambda_i \neq 0$ (i.e., that the matrices A and J are invertible). Indeed, if $\lambda_i = 0$, then B_i is a nilpotent block and therefore, for the purpose of invariant synthesis, we can ignore finitely many points of the orbit under A until B_i^n is the 0 block. We can then restrict our attention to the image of A^n , by identifying it with $\mathbb{R}^{d-d_i}$.

For all $i \in \{1, \dots, k\}$, we can write $\lambda_i = \rho_i \xi_i$ where $\rho_i > 0$ is positive real and ξ_i is a complex number of absolute value 1, with both ρ_i and ξ_i being algebraic.

Observe that now, for every set $F \subseteq \mathbb{R}^d$, we have that $A^n s \in F$ iff $J^n s' \in P^{-1}F$ where $s' = P^{-1}s$. For every $n > d$, $J^n = \text{diag}(B_1^n, \dots, B_k^n)$ with

$$B_i^n = \begin{pmatrix} \lambda_i^n & \frac{n}{\lambda_i} \lambda_i^n & \dots & \frac{\binom{n}{d_i-1}}{\lambda_i^{d_i-1}} \lambda_i^n \\ & \ddots & & \vdots \\ & & & \lambda_i^n \end{pmatrix}.$$

Every coordinate of $J^n s'$ is of the form $\lambda_i^n Q_{i,j}(n) = (\rho_i \xi_i)^n Q_{i,j}(n) = \rho_i^n \xi_i^n Q_{i,j}(n)$ for some $1 \leq i \leq k$ and $1 \leq j \leq d_i$, where $Q_{i,j}$ is a polynomial (possibly with complex coefficients) that depends on J and s' .

Let $R = \text{diag}(\rho_1, \dots, \rho_k)$ and $L = \text{diag}(\xi_1, \dots, \xi_k)$. We define $\mathbb{T}$ to be the subgroup of the torus in $\mathbb{C}^k$ generated by the multiplicative relations of the normalised eigenvalues $\xi_1, \dots, \xi_k$. That is, consider the subgroup $G = \{v = (v_1, \dots, v_k) \in \mathbb{Z}^k \mid \xi_1^{v_1} \cdots \xi_k^{v_k} = 1\}$ of $\mathbb{Z}^k$, and let

$$\mathbb{T} = \{(\alpha_1, \dots, \alpha_k) \in \mathbb{C}^k \mid |\alpha_i| = 1 \text{ for all } i, \text{ and for every } v \in G, \alpha_1^{v_1} \cdots \alpha_k^{v_k} = 1\}.$$

A result by Masser [24] allows to compute a basis for G , and hence a representation of $\mathbb{T}$. Specifically, we use the following form of the result, adapted to our setting from [5, Theorem 3.1]:

THEOREM 3.1 (MASSER). *Denoted by $\|A\|$ the description length of $A \in \mathbb{Q}^{d \times d}$ (i.e., the sum of the lengths of the binary encodings of its entries), then there exists a basis of G , as defined above, such that the absolute value of each entry of the basis vectors is at most $(cd \cdot \text{poly}(\|A\|))^{d-1} d!^{6d}$, where c is some absolute constant.*

Theorem 3.1 allows us to compute a basis by exhaustive search, going over all sets of vectors whose entries are below the given bound. Since we use binary encoding, and the bound in Theorem 3.1 is single-exponential in $\|A\|$, then each basis can be represented in polynomial space (hence, finding a basis can be done in polynomial space).

Using Kronecker's theorem on inhomogeneous simultaneous Diophantine approximation [6] it is shown in [29] that $\{L^n \mid n \in \mathbb{N}\}$ is a dense subset of $\{\text{diag}(\alpha_1, \dots, \alpha_k) \mid (\alpha_1, \dots, \alpha_k) \in \mathbb{T}\}$.

Thus, for every $n \in \mathbb{N}$, we have

$$J^n s' \in \left\{ \left(\rho_1^n p_1 Q_{1,1}(n), \dots, \rho_k^n p_k Q_{k,d_k}(n) \right)^\top \mid (p_1, \dots, p_k) \in \mathbb{T} \right\}.$$

We now define a continuous over-approximation of the expressions ρ_i^n by replacing $n \in \mathbb{N}$ with $\log t$, where $t \geq 1$ is a real variable, so that, writing $b_i := \log \rho_i$, ρ_i^n becomes t^{b_i} . This over-approximation leads to the following definition, which is central to our approach.

Definition 3.2. For every $t_0 \geq 1$, we define the *trajectory cone*¹ for the orbit O as

$$C_{t_0} = \left\{ \left(t^{b_1} p_1 Q_{1,1}(\log t), \dots, t^{b_k} p_k Q_{k,d_k}(\log t) \right)^\top \mid (p_1, \dots, p_k) \in \mathbb{T}, t \geq t_0 \right\}.$$

In particular, we have that $J^n s' \in C_1$.

In order to analyse invariants, we require a finer-grained notion than the entire trajectory cone. To this end, we introduce the following.

Definition 3.3. For every $p = (p_1, \dots, p_k) \in \mathbb{T}$ and every $t_0 \geq 1$, we define the *(trajectory) ray*² $r(p, t_0) = \{(t^{b_1} p_1 Q_{1,1}(\log t), \dots, t^{b_k} p_k Q_{k,d_k}(\log t))^\top \mid t \geq t_0\}$.

Observe that we have $C_{t_0} = \bigcup_{p \in \mathbb{T}} r(p, t_0)$.

Example 3.4. Consider the matrix $A = \text{diag}(5, 2)$ and the initial point $s = (1, 1)^\top$. We then have $\mathbb{T} = \{(1, 1)\}$ and $C_{t_0} = \{(t^{\log 5}, t^{\log 2})^\top \mid t \geq t_0\}$. Observe that this is not an $\mathcal{R}_0$ -definable set, as the quotient $\frac{\log 5}{\log 2}$ is not rational. This shows that even for diagonalizable matrices (where C_{t_0} has a simple form, devoid of the polynomials $Q_{i,j}$), $\mathcal{R}_0$ might not be enough to recover definability of the orbit (in the sense of Theorem 4.1 below).

¹These sets are, of course, not really cones. Nevertheless, if for all i we have $b_i = 1$ and the polynomials $Q_{i,j}$ are constant, then the set C_{t_0} is a conical surface formed by the union of rays going from the origin through all points of $\mathbb{T}$. The initial segments of the rays, of length determined by the parameter t_0 , are removed.

²Likewise, this set is not, strictly speaking, a straight half-line.

4 CONSTRUCTING INVARIANTS FROM TRAJECTORY CONES

We now proceed to show that the trajectory cones defined in Section 3 can be used to characterise o-minimal invariants. More precisely, we show that for an LDS (A, s) with $A = PJP^{-1}$, the image under P of every trajectory cone C_{t_0} , augmented with finitely many points from $\mathcal{O}$, is an invariant. Moreover, we show that such invariants are $\mathfrak{R}_{\text{exp}}$ -definable, and hence o-minimal. Complementing this, we show in Section 5 that every o-minimal invariant must contain some trajectory cone.

In what follows, let $A = PJP^{-1}$, s , as well as the real numbers $b_1, \dots, b_d$ be defined as in Section 3.

THEOREM 4.1. *For every $t_0 \geq 1$, the set $P \cdot C_{t_0} \cup \{A^n s \mid n < \log t_0\}$ is an $\mathfrak{R}_{\text{exp}}$ -definable invariant for the LDS (A, s) .*

The intuition behind Theorem 4.1 is as follows. Clearly, the orbit $\mathcal{O}$ itself is always an invariant for (A, s) . However, it is generally not definable in any o-minimal structure (in particular, since it has infinitely many connected components). In order to recover definability in $\mathfrak{R}_{\text{exp}}$ while maintaining stability under A , the invariants constructed in Theorem 4.1 over-approximate the orbit by the image of the trajectory cone C_{t_0} under the linear transformation P . Finally, a finite set of points from $\mathcal{O}$ is added to this image of the trajectory cone, to fill in the missing points in case t_0 is too large.

The proof of Theorem 4.1 has several parts. First, recall that the trajectory cone itself, C_{t_0} , is an over-approximation of the set $\{J^n P^{-1} s \mid n \in \mathbb{N}\}$. As such, clearly $C_{t_0} \subseteq \mathbb{C}^d$. In comparison, the orbit can be written as $\mathcal{O} = \{PJ^n P^{-1} s \mid n \in \mathbb{N}\} \subseteq \mathbb{R}^d$. We prove in Section 4.1 the following simple lemma, from which it follows that the entire set $P \cdot C_{t_0}$ is also a subset of $\mathbb{R}^d$.

LEMMA 4.2. *For every $p \in \mathbb{T}$ and $t_0 \geq 1$, we have $P \cdot r(p, t_0) \subseteq \mathbb{R}^d$.*

In the second part of the proof of Theorem 4.1, we show that $P \cdot C_{t_0}$ is stable under A . The key ingredient is the following lemma, which characterises the action of J on rays and is proved in Section 4.2.

LEMMA 4.3. *For every $p = (p_1, \dots, p_k) \in \mathbb{T}$ and $t_0 \geq 1$, we have $J \cdot r(p, t_0) = r(Lp, et_0)$.*

The next lemma then lifts Lemma 4.3 to the entire trajectory cone.

LEMMA 4.4. *For every $t_0 \geq 1$, we have $J \cdot C_{t_0} \subseteq C_{t_0}$.*

PROOF. Recall that $C_{t_0} = \bigcup_{p \in \mathbb{T}} r(p, t_0)$. By Lemma 4.3 we have that $J \cdot C_{t_0} = \bigcup_{p \in \mathbb{T}} r(Lp, et_0)$. But $et_0 > t_0$ and $Lp \in \mathbb{T}$ iff $p \in \mathbb{T}$. Hence, we have that $r(Lp, et_0) \subseteq r(p, t_0)$, from which we conclude that $J \cdot C_{t_0} \subseteq \bigcup_{p \in \mathbb{T}} r(p, t_0) = C_{t_0}$. $\square$

The proof of Theorem 4.1 combines all these ingredients together and is given in Section 4.3.

4.1 Proof of Lemma 4.2

Recall that we have $A = PJP^{-1}$, where $J = \text{diag}(B_1, \dots, B_k)$ is a block diagonal matrix with $B_i \in \mathbb{C}^{d_i \times d_i}$ the Jordan block corresponding to eigenvalue λ_i . Write $P = \begin{pmatrix} P_1 & \dots & P_k \end{pmatrix}$, where $P_i \in \mathbb{C}^{d \times d_i}$ for all $i \in \{1, \dots, k\}$. Since the generalized eigenspaces of A respectively corresponding to pairs of complex-conjugate eigenvalues are themselves element-wise conjugate, we can partition the set $\{1, \dots, k\}$ into singletons and pairs of the form $\{i_1, i_2\}$ such that $d_{i_1} = d_{i_2}$, $P_{i_1} = \overline{P_{i_2}}$, and $B_{i_1} = \overline{B_{i_2}}$. In this case, we say that i_1 and i_2 are *conjugate block indices*.

By definition, for conjugate block indices i_1, i_2 we have that for all $j \in \{1, \dots, d_{i_1}\}$ the column of P with block-column index (i_1, j) is conjugate to the column of P with block-column index (i_2, j) .

Likewise, the row of P^{-1} with block-row index (i_1, j) is conjugate to that with index (i_2, j) .³ In particular, for the vector $s' = P^{-1}s$, we have that the entries $s'_{i_1, j}$ and $s'_{i_2, j}$ are complex conjugates.

Let $p \in \mathbb{T}$ and $t_0 \geq 1$. Consider the vector $v := \begin{pmatrix} t^{b_1 p_1 Q_{i_1, 1}(\log t)} \\ \vdots \\ t^{b_k p_k Q_{i_k, d_k}(\log t)} \end{pmatrix} \in r(p, t_0)$. Fix two conjugate block indices $i_1, i_2 \in \{1, \dots, k\}$. We claim that for all $j \in \{1, \dots, d_{i_1}\}$ the entries of v with respective block-row indices (i_1, j) and (i_2, j) , namely $t^{b_1 p_{i_1} Q_{i_1, j}(\log t)}$ and $t^{b_2 p_{i_2} Q_{i_2, j}(\log t)}$, are mutually conjugate.

Towards proving the claim, observe that for every $1 \leq i \leq k$ and $1 \leq j \leq d_i$ we have

$$Q_{i, j}(\log t) = \sum_{m=0}^{d_i-j} \frac{\binom{\log t}{m}}{\lambda_i^m} \cdot s'_{i, j+m},$$

with $(i, j+m)$ being a block-row index.⁴ It follows that the values $Q_{i_1, j}(\log t)$ and $Q_{i_2, j}(\log t)$ are complex conjugates. Note moreover that we have $p_{i_1} p_{i_2} = 1$ since $p \in \mathbb{T}$ and $\xi_{i_1} \xi_{i_2} = 1$. Thus, p_{i_1} and p_{i_2} are also complex conjugates. Finally, we note that $b_{i_1} = \log |\lambda_{i_1}| = \log |\lambda_{i_2}| = b_{i_2}$ and hence, $t^{b_{i_1}} = t^{b_{i_2}}$. The claim follows.

Given the above claim and the fact that for conjugate block indices i_1 and i_2 , for all j the respective rows of P with indices (i_1, j) and (i_2, j) are element-wise conjugate, we conclude that $Pv \in \mathbb{R}$. This concludes the proof. $\square$

4.2 Proof of Lemma 4.3

Let $y = \begin{pmatrix} t^{b_1 p_1 Q_{i_1, 1}(\log t)} \\ \vdots \\ t^{b_k p_k Q_{i_k, d_k}(\log t)} \end{pmatrix} \in r(p, t_0)$. We claim that $Jy = \begin{pmatrix} (et)^{b_1 \xi_1 p_1 Q_{i_1, 1}(\log(et))} \\ \vdots \\ (et)^{b_k \xi_k p_k Q_{i_k, d_k}(\log(et))} \end{pmatrix}$. Note that since $Lp = (\xi_1 p_1, \dots, \xi_k p_k)$, the above suffices to conclude the proof.

Consider a coordinate $m = (i, j)$ of Jy in block-row index, with $j < d_i$. The case of $j = d_i$ is similar and simpler. To simplify notation, we write ξ, ρ , and d instead of ξ_i, ρ_i , and d_i , respectively. Then we have

$$(Jy)_m = \xi \rho t^{b_i} p_i Q_{i, j}(\log t) + t^{b_i} p_i Q_{i, j+1}(\log t).$$

Recall that⁵

$$Q_{i, j}(\log t) = \sum_{c=0}^{d-j} \frac{\binom{\log t}{c}}{(\rho \xi)^c} s'_{i, j+c},$$

with $(i, j+c)$ in block-row index. We can then write

$$(Jy)_m = \xi \rho t^{b_i} p_i \sum_{c=0}^{d-j} \frac{\binom{\log t}{c}}{(\rho \xi)^c} s'_{i, j+c} + t^{b_i} p_i \sum_{c=0}^{d-j-1} \frac{\binom{\log t}{c}}{(\rho \xi)^c} s'_{i, j+c+1}. \quad (2)$$

We now compare this to coordinate m of our claim, namely

$$(et)^{b_i \xi} p_i Q_{i, j}(\log(et)) = (et)^{b_i \xi} p_i \sum_{c=0}^{d-j} \frac{\binom{\log(et)}{c}}{(\rho \xi)^c} s'_{i, j+c}. \quad (3)$$

³Since $\overline{P} = PS$ for S the permutation matrix that interchanges conjugate blocks, we have $\overline{P^{-1}} = S^{-1}P^{-1}$.

⁴Here, for $s \in \mathbb{R}$ and $m \in \mathbb{N}$, one defines $\binom{s}{m} = \frac{1}{m!} \prod_{i=0}^{m-1} (s-i)$, which maintains consistency with the original definition of $Q_{i, j}$ in Section 3.

⁵Here, for $w \in \mathbb{R}$ and $m \in \mathbb{N}$, one defines $\binom{w}{m} = \frac{1}{m!} \prod_{i=0}^{m-1} (w-i)$, which maintains consistency with the original definition of $Q_{i, j}$ in Section 3.

We compare the right-hand sides of Equations (2) and (3) by comparing the coefficients of $s'_{i,q}$ for $q \in \{j, \dots, d\}$ (these being the only ones that appear in the expressions). For $q = j$ we see that in Equation (2) the number $s'_{i,j}$ occurs in the first summand only, and its coefficient is thus $\xi \rho t^{b_i} p_i$, while in Equation (3) it is $(et)^{b_i} \xi p_i = e^{b_i} t^{b_i} \xi p_i = \rho t^{b_i} \xi p_i$, since $b_i = \log \rho$. Thus, the coefficients are equal.

For $q > j$, write $q = j + c$ with $c \geq 1$; the coefficient at $s'_{i,j+c}$ in Equation (2) is then

$$\xi \rho t^{b_i} p_i \frac{\binom{\log t}{c}}{(\rho \xi)^c} + t^{b_i} p_i \frac{\binom{\log t}{c-1}}{(\rho \xi)^{c-1}} = \frac{t^{b_i} \rho \xi p_i}{(\rho \xi)^c} \left(\binom{\log t}{c} + \binom{\log t}{c-1} \right) = \frac{t^{b_i} \xi \rho p_i}{\xi^c} \binom{\log t + 1}{c},$$

where the last equality follows from a continuous version of Pascal's identity. Finally, by noticing that $\log t + 1 = \log(et)$, it is easy to see that this is the same coefficient as in Equation (3).

4.3 Proof of Theorem 4.1

Let $t_0 \geq 1$. By applying Lemma 4.2 to every $p \in \mathbb{T}$, we conclude that $P \cdot C_{t_0} \subseteq \mathbb{R}^d$. It is then easy to see that $P \cdot C_{t_0}$ is definable in $\mathfrak{R}_{\text{exp}}$ (note that the only reason the set C_{t_0} might fail to be $\mathfrak{R}_{\text{exp}}$ -definable is that the underlying domain should be $\mathbb{R}$ and not $\mathbb{C}$).

Next, by Lemma 4.4 we have that $J \cdot C_{t_0} \subseteq C_{t_0}$. Applying P from the left, we get $PJ \cdot C_{t_0} \subseteq P \cdot C_{t_0}$. Thus, we have $AP \cdot C_{t_0} = PJP^{-1}P \cdot C_{t_0} = PJ \cdot C_{t_0} \subseteq P \cdot C_{t_0}$.

Finally, observe that $\{A^n s \mid n \geq \log t_0\} \subseteq P \cdot C_{t_0}$. Since any finite subset of $\mathcal{O}$ can be described in $\mathfrak{R}_0$, we conclude that the set $\{A^n s \mid n < \log t_0\} \cup P \cdot C_{t_0}$ is an $\mathfrak{R}_{\text{exp}}$ -definable invariant for (A, s) .

5 O-MINIMAL INVARIANTS MUST CONTAIN TRAJECTORY CONES

In this section, we consider invariants definable in o-minimal extensions of $\mathfrak{R}_{\text{exp}}$. Fix such an extension $\mathfrak{R}$ for the remainder of this section.

THEOREM 5.1. *Consider an $\mathfrak{R}$ -definable invariant $\mathcal{I}$ for the LDS (A, s) . Then there exists $t_0 \geq 1$ such that $P \cdot C_{t_0} \subseteq \mathcal{I}$.*

To prove Theorem 5.1, we begin by making the following claims of increasing strength:

CLAIM 1. *For every $p \in \mathbb{T}$ there exists $t_0 \geq 1$ such that $P \cdot r(p, t_0) \subseteq \mathcal{I}$ or $P \cdot r(p, t_0) \cap \mathcal{I} = \emptyset$.*

CLAIM 2. *For every $p \in \mathbb{T}$ there exists $t_0 \geq 1$ such that $P \cdot r(p, t_0) \subseteq \mathcal{I}$.*

CLAIM 3. *There exists $t_0 \geq 1$ such that for every $p \in \mathbb{T}$, we have $P \cdot r(p, t_0) \subseteq \mathcal{I}$.*

PROOF OF CLAIM 1. Fix $p \in \mathbb{T}$. Then the set

$$\{t \geq 0 : P(t^{b_1} p_1 Q_{1,1}(\log t), \dots, t^{b_k} p_k Q_{k,d_k}(\log t))^T \in \mathcal{I}\},$$

is $\mathfrak{R}$ -definable and hence comprises a finite union of intervals. If this set contains an unbounded interval then there exists t_0 such that $P \cdot r(p, t_0) \subseteq \mathcal{I}$; otherwise there exists t_0 such that $P \cdot r(p, t_0) \cap \mathcal{I} = \emptyset$. $\square$

Before proceeding to Claim 2, we prove an auxiliary lemma, which is an adaptation of a similar result from [15]. For a set X , we write $\overline{X}$ to refer to the topological closure of X . We use the usual topology on $\mathbb{R}^n$, $\mathbb{C}^n$, and the (usual) subspace topology on their subsets.

LEMMA 5.2. *Let $S, F \subseteq \mathbb{T}$ be $\mathfrak{R}$ -definable⁶ sets such that $\overline{S} = \overline{F} = \mathbb{T}$. Then $F \cap S \neq \emptyset$.*

⁶Recall that, in order to reason about $\mathbb{T} \subseteq \mathbb{C}^k$ in $\mathfrak{R}$, we identify $\mathbb{C}$ with $\mathbb{R}^2$.

PROOF. We start by stating two properties of the dimension of a definable set in an o-minimal theory $\mathfrak{R}$. First, for any $\mathfrak{R}$ -definable set $X \subseteq \mathbb{R}^n$, we have $\dim(X) = \dim(\overline{X})$ [13, Chapter 4, Theorem 1.8]. Secondly, if $X \subseteq Y$ are $\mathfrak{R}$ -definable subsets of $\mathbb{R}^n$ that have the same dimension, then X has a non-empty interior in Y [13, Chapter 4, Corollary 1.9]. In the situation at hand, since $\dim(F) = \dim(\overline{F})$, it follows that F has a non-empty interior with respect to the subspace topology on $\overline{F} = \overline{S}$. But then S is dense in $\overline{S}$, while F has a non-empty interior in $\overline{S}$, and thus $S \cap F \neq \emptyset$. $\square$

PROOF OF CLAIM 2. We strengthen Claim 1. Assume by way of contradiction that there exist $p \in \mathbb{T}$ and $t_0 \in \mathbb{R}$ such that $P \cdot r(p, t_0) \cap \mathcal{I} = \emptyset$, and consider $J^{-1} \cdot r(p, t_0)$. Let $q \in \mathbb{T}$ be $L^{-1}p = (\frac{p_1}{\xi_1}, \dots, \frac{p_k}{\xi_k})$ and let $t_1 = \frac{t_0}{e}$. Then, $p = Lq$ and $t_0 = et_1$ and, by Lemma 4.3, $Jr(q, t_1) = r(Lq, et_1) = r(p, t_0)$. Since J is invertible, we conclude that $J^{-1}r(p, t_0) = r(q, t_1)$.

We now claim that $P \cdot r(q, t_1) \cap \mathcal{I} = \emptyset$. Recall that $P \cdot r(p, t_0) \cap \mathcal{I} = \emptyset$. Applying $A^{-1} = PJ^{-1}P^{-1}$, we have by the above that $P \cdot r(q, t_1) \cap A^{-1}\mathcal{I} = \emptyset$. Since $A\mathcal{I} \subseteq \mathcal{I}$, then $\mathcal{I} \subseteq A^{-1}\mathcal{I}$, so we have $P \cdot r(q, t_1) \cap \mathcal{I} \subseteq P \cdot r(q, t_1) \cap A^{-1}\mathcal{I} = \emptyset$.

Clearly $t_1 \leq t_0$ and $r(q, t_0) \subseteq r(q, t_1)$, so, in particular, $P \cdot r(q, t_0) \cap \mathcal{I} = \emptyset$. Thus, assuming $P \cdot r(p, t_0) \cap \mathcal{I} = \emptyset$, we have just proved that $P \cdot r(L^{-1}p, t_0) \cap \mathcal{I} = \emptyset$; repeating this argument, we get that for every $n \in \mathbb{N}$, the point $s = L^{-n}p$ satisfies $P \cdot r(s, t_0) \cap \mathcal{I} = \emptyset$.

Let $S = \{L^{-n}p \mid n \in \mathbb{N}\}$. Then S is dense in $\mathbb{T}$, since the group of multiplicative relations defined by the eigenvalues of L^{-1} is the same as the one defined by those of L . Define $S' = \{s \in \mathbb{T} \mid P \cdot r(s, t_0) \cap \mathcal{I} = \emptyset\}$. Then S' is $\mathfrak{R}$ -definable, and we have $S \subseteq S' \subseteq \mathbb{T}$. Moreover, $\overline{S} = \mathbb{T}$, so $\overline{S'} = \mathbb{T}$.

We now prove that, in fact, $S' = \mathbb{T}$. Assuming (again by way of contradiction) that there exists $q \in \mathbb{T} \setminus S'$, then by the definition of S' we have $P \cdot r(q, t_0) \cap \mathcal{I} \neq \emptyset$. It follows that for every $n \in \mathbb{N}$, the point $q' = L^n q$ also satisfies $P \cdot r(q', t_0) \cap \mathcal{I} \neq \emptyset$. Define $F = \{L^n q \mid n \in \mathbb{N}\}$, then F is dense in $\mathbb{T}$. But then the set $F' = \{q \in \mathbb{T} \mid P \cdot r(q, t_0) \cap \mathcal{I} \neq \emptyset\}$ satisfies $F \subseteq F' \subseteq \mathbb{T}$ and $\overline{F'} = \mathbb{T}$. Now the sets S' and F' are both definable in $\mathfrak{R}$, and the topological closure of each of them is $\mathbb{T}$. It follows from Lemma 5.2 that $F' \cap S' \neq \emptyset$, which is clearly a contradiction. Therefore, there is no $q \in \mathbb{T} \setminus S'$; that is, $S' = \mathbb{T}$.

From this, however, it follows that $P \cdot C_{t_0} \cap \mathcal{I} = \emptyset$, which is again a contradiction, since $P \cdot C_{t_0} \cap \mathcal{O} \neq \emptyset$ and $\mathcal{O} \subseteq \mathcal{I}$, so we are done. $\square$

PROOF OF CLAIM 3. Consider the function $f: \mathbb{T} \rightarrow \mathbb{R}$ defined by $f(p) = \inf\{t \in \mathbb{R} \mid P \cdot r(p, t) \subseteq \mathcal{I}\}$. By Claim 2 this function is well-defined. Since $P \cdot r(p, t)$ is $\mathfrak{R}$ -definable, then so is f . Moreover, its graph $\Gamma(f)$ has finitely many connected components, and the same dimension as $\mathbb{T}$. Thus, there exists an open set $K \subseteq \mathbb{T}$ (in the induced topology on $\mathbb{T}$) such that f is continuous on K . Furthermore, K is homeomorphic to $(0, 1)^m$ for some $0 \leq m \leq k$, and thus we can find sets $K'' \subseteq K' \subseteq K$ such that K'' is open, and K' is closed.⁷ Since f is continuous on K , it attains a maximum on K' . Consider the set $\{L^n \cdot K'' \mid n \in \mathbb{N}\}$. By the density of $\{L^n \mid n \in \mathbb{N}\}$ in $\mathbb{T}$, this is an open cover of $\mathbb{T}$, and hence there is a finite subcover $\{L^{n_1}K'', \dots, L^{n_m}K''\}$. Since $K'' \subseteq K'$, it follows that $\{L^{n_1}K', \dots, L^{n_m}K'\}$ is a finite closed cover of $\mathbb{T}$.

We now show that, for all $p \in \mathbb{T}$, we have $f(Lp) \leq e \cdot f(p)$. Indeed, consider any $p \in \mathbb{T}$ and $t > 0$ such that $P \cdot r(p, t) \subseteq \mathcal{I}$. Applying $A = PJP^{-1}$, we get $PJ \cdot r(p, t) \subseteq A\mathcal{I} \subseteq \mathcal{I}$. By Lemma 4.3, $J \cdot r(p, t) = r(Lp, et)$, so we can conclude that $P \cdot r(Lp, et) \subseteq \mathcal{I}$. This means that $P \cdot r(p, t) \subseteq \mathcal{I}$ implies $P \cdot r(Lp, et) \subseteq \mathcal{I}$; therefore, $f(Lp) \leq e \cdot f(p)$.

Now denote $s_0 = \max_{p \in K'} f(p)$. Then for every $1 \leq i \leq m$ we have $\max_{p \in L^{n_i}K'} f(p) \leq e^{n_i} s_0$; so $f(p)$ is indeed bounded on $\mathbb{T}$. $\square$

⁷In case $m = 0$, the proof actually follows immediately from Claim 2, since $\mathbb{T}$ is finite.

Finally, we conclude from Claim 3 that there exists $t_0 \geq 1$ such that $P \cdot C_{t_0} \subseteq I$. This completes the proof of Theorem 5.1.

6 DECIDING THE EXISTENCE OF O-MINIMAL INVARIANTS

We now turn to the algorithmic aspects of invariants and present our two main results, Theorems 6.2 and 6.3.

Let $\mathfrak{R}$ be either $\mathfrak{R}_0$ or $\mathfrak{R}_{\text{exp}}$. We consider the following problem: given an LDS (A, s) , with $A \in \mathbb{Q}^{d \times d}$ and $s \in \mathbb{Q}^d$, and given an $\mathfrak{R}$ -definable halting set $F \subseteq \mathbb{R}^d$, we wish to decide whether there exists an o-minimal invariant I for (A, s) that avoids F , and to compute such an invariant if it exists. We term this question the *O-Minimal Invariant Synthesis Problem for $\mathfrak{R}$ -Definable Halting Sets*.

The following is a consequence of Theorems 4.1 and 5.1.

LEMMA 6.1. *Let (A, s) and $\mathfrak{R}$ be as above, and let F be $\mathfrak{R}$ -definable. Then there exists an o-minimal invariant I for (A, s) that avoids F iff there is some $t_0 \geq 1$ such that $P \cdot C_{t_0} \cap F = \emptyset$ and such that $A^n s \notin F$ for every $0 \leq n \leq \log t_0$.*

PROOF. By Theorem 5.1, if an o-minimal invariant I for (A, s) exists, then there exists $t_0 \geq 1$ such that $P \cdot C_{t_0} \subseteq I$. Moreover, $I \cap F = \emptyset$ implies $O \cap F = \emptyset$, so that $A^n s \notin F$ for every $n \in \mathbb{N}$, and in particular for $0 \leq n \leq \log t_0$.

Conversely, let there be $t_0 \geq 1$ such that $P \cdot C_{t_0} \cap F = \emptyset$, and for every $0 \leq n \leq \log t_0$, it holds that $A^n s \notin F$. Let $t'_0 \in \mathbb{Q}$ be such that $t'_0 \geq t_0$. By Theorem 4.1, the set $I = P \cdot C_{t'_0} \cup \{A^n s \mid 0 \leq n \leq \log t'_0\}$ is an $\mathfrak{R}_{\text{exp}}$ -definable invariant that avoids F . $\square$

Observe that the formula $\exists t_0 \geq 1 : P \cdot C_{t_0} \cap F = \emptyset$ is a sentence in $\mathfrak{R}_{\text{exp}}$, and by Lemma 6.1, deciding the existence of an invariant amounts to determining the truth value of this sentence.

6.1 Decidability for $\mathfrak{R}_{\text{exp}}$ -definable Halting Sets Assuming Schanuel's Conjecture

Applying Theorem 4.1, we note that an invariant for (A, s) that avoids F —if one exists—can always be defined in $\mathfrak{R}_{\text{exp}}$.

THEOREM 6.2. *The O-Minimal Invariant Synthesis Problem for $\mathfrak{R}_{\text{exp}}$ -Definable Halting Sets is decidable, assuming Schanuel's conjecture. Moreover, in positive instances, we can explicitly define such an invariant in $\mathfrak{R}_{\text{exp}}$.*

PROOF. Assume Schanuel's conjecture. Then by [23], the first-order theory of the structure $\mathfrak{R}_{\text{exp}}$ is decidable. Thus we can decide whether there exists $t_0 \geq 1$ such that $P \cdot C_{t_0} \cap F = \emptyset$. If the sentence is false, then by Lemma 6.1 there is no invariant, and we are done. If the sentence is true, however, it still remains to check whether $A^n s \notin F$ for every $0 \leq n \leq \log t_0$. While we can decide whether $A^n s \in F$ for a fixed n , observe that we do not have an *a priori* bound on t_0 . Hence we proceed as follows: For every $n \in 1, 2, \dots$, check both whether $A^n s \in F$ and, for $t_0 = e^n$, whether $P \cdot C_{t_0} \cap F = \emptyset$. In case $A^n s \in F$, then clearly there is no invariant, since $O \cap F \neq \emptyset$, and we are done. On the other hand, if $P \cdot C_{t_0} \cap F = \emptyset$, then return the $\mathfrak{R}_{\text{exp}}$ -definable invariant as per Lemma 6.1.

We claim that the above procedure always halts. Indeed, we know that there exists t_0 for which $P \cdot C_{t_0} \cap F = \emptyset$. Thus, either for some $n < \log t_0$, it holds that $A^n s \in F$, in which case there is no invariant and we halt when we reach n , or we proceed until we reach $n \geq \log t_0$, in which case we halt and return the invariant. $\square$

Remark 3. It is interesting to note that, should Schanuel's conjecture turn out to be false, the above procedure could still never return a “wrong” invariant. The worst that could happen is that

decidability of $\mathfrak{R}_{\text{exp}}$ fails in that the putative algorithm of [23] simply never halts, so no verdict is ever returned.

6.2 Unconditional Decidability for Semi-algebraic Halting Sets

In this section, we restrict attention to semi-algebraic halting sets. Our main result is as follows.

THEOREM 6.3. *The O-Minimal Invariant Synthesis Problem for Semi-Algebraic Halting Sets is decidable. Moreover, in positive instances, we can explicitly define such an invariant in $\mathfrak{R}_0$.*

Theorem 6.3 may come as a double surprise: First, as shown by Lemma 6.1, deciding the existence of an o-minimal invariant amounts to determining the truth value of the sentence $\exists t_0 \geq 1 : P \cdot C_{t_0} \cap F = \emptyset$. Since $P \cdot C_{t_0}$ might not be definable in $\mathfrak{R}_0$, then this sentence is only $\mathfrak{R}_{\text{exp}}$ -definable, even when F is $\mathfrak{R}_0$ -definable. Therefore, determining the truth value of this sentence is not immediate. Second, even if we do manage to determine the truth value of this sentence, the synthesized invariant as per Theorem 4.1 (namely $P \cdot C_{t_0}$ along with a finite “tail”) is $\mathfrak{R}_{\text{exp}}$ definable, but might not be $\mathfrak{R}_0$ definable. Thus, the synthesized invariant in Theorem 6.3 must be more elaborate than $P \cdot C_{t_0}$.

We prove Theorem 6.3 in parts, first addressing the decidability of the existence of an o-minimal invariant (Theorem 6.4 below), and then showing how to synthesize, in positive instances, an $\mathfrak{R}_0$ -definable invariant (Section 6.2.1).

THEOREM 6.4. *The O-Minimal Invariant Synthesis Problem for Semi-Algebraic Halting Sets is decidable. Moreover, in positive instances, we can explicitly define such an invariant in $\mathfrak{R}_{\text{exp}}$.*

By Lemma 6.1, in order to prove Theorem 6.4, it is enough to decide the truth value of the $\mathfrak{R}_{\text{exp}}$ -sentence $\exists t_0 \geq 1 : P \cdot C_{t_0} \cap F = \emptyset$. Indeed, as $A^n s \in \mathbb{Q}^d$, one can always check unconditionally whether for a given n the vector $A^n s$ belongs to the semi-algebraic set F . The algorithm is then otherwise the same as that presented in the proof of Theorem 6.2. The proof of Theorem 6.4 therefore boils down to the following lemma.

LEMMA 6.5. *For F a semi-algebraic set, it is decidable whether there exists $t_0 \geq 1$ such that $P \cdot C_{t_0} \cap F = \emptyset$.*

Our key tool is the following celebrated result from transcendental number theory:

THEOREM 6.6 (BAKER’S THEOREM [1]). *Let $\alpha_1, \dots, \alpha_m \in \mathbb{C}$ be algebraic numbers different from 0 and let $b_1, \dots, b_m \in \mathbb{Z}$ be integers. Write $\Lambda = b_1 \log \alpha_1 + \dots + b_m \log \alpha_m$. There exists a number C effectively computable from $b_1, \dots, b_m, \alpha_1, \dots, \alpha_m$ such that if $\Lambda \neq 0$ then $|\Lambda| > H^{-C}$, where H is the maximum height of $\alpha_1, \dots, \alpha_m$.*

Recall that the subgroup $\mathbb{T}$ of the torus defined by the multiplicative relations of the eigenvalues of A is a semi-algebraic set. Write $\tau(t) = (t^{b_1} Q_{1,1}(\log t), \dots, t^{b_k} Q_{k,d_k}(\log t))$, and consider the set

$$U = \{(z_1, \dots, z_d)^\top \in \mathbb{C}^d \mid \forall (p_1, \dots, p_d) \in \mathbb{T}, P(z_1 p_1, \dots, z_d p_d)^\top \in \mathbb{R}^d \setminus F\}.$$

It is enough to decide whether there exists $t_0 \geq 1$ such that for all $t \geq t_0$, $\tau(t) \in U$.

Observe that U is a semi-algebraic set (see Remark 2 on Page 4).

Using cell decomposition, describe U as a finite union of connected components, each of which is given by a conjunction of the form $\bigwedge_{l=1}^m R_l(u_1, \dots, u_d, v_1, \dots, v_d) \sim_l 0$. Here, for every $1 \leq l \leq m$, $\sim_l \in \{>, =\}$ and R_l is a polynomial with integer coefficients in variables $u_1, \dots, u_d, v_1, \dots, v_d$; for each i , the variables u_i and v_i represent $\text{Re } z_i$ and $\text{Im } z_i$, the real and imaginary parts of z_i , respectively.

We claim that we can restrict our attention to a single connected component. Indeed, the first note that by substituting $\tau(t)$ for $(z_1, \dots, z_d)$ in the conjunction $\bigwedge_{l=1}^m R_l \sim_l 0$, we get a constraint

on t expressible in $\mathfrak{R}_{\text{exp}}$. By o-minimality of $\mathfrak{R}_{\text{exp}}$, the set of all $t \in \mathbb{R}$ satisfying this conjunction is a finite union of points and (possibly unbounded) intervals. Therefore, since the number of connected components is finite, the following two statements are equivalent: (i) there exists $t_0 \geq 1$ such that for all $t \geq t_0$ it holds that $\tau(t) \in U$, and (ii) there exists a single connected component of U for which this holds (perhaps with a larger value of t_0).

Thus, we now need to decide whether we can find $t_0 \geq 1$ such that for every $t \geq t_0$ it holds that $R_l(\tau(t)) \sim_l 0$ for every $1 \leq l \leq m$. Fix $1 \leq l \leq m$. Recall that we consider every vector in $\mathbb{C}^d$ as a vector in $\mathbb{R}^{2d}$; thus, the polynomial R_l has the form

$$\sum_i a_i \cdot u_1^{n'_{i,1}} \cdot \dots \cdot u_d^{n'_{i,d}} \cdot v_1^{n''_{i,1}} \cdot \dots \cdot v_d^{n''_{i,d}},$$

with $a_i \in \mathbb{Z}$ and $n'_{i,s}, n''_{i,s} \in \mathbb{Z}_{\geq 0}$. Therefore, $R_l(\tau(t))$ is the sum of terms of the form

$$a_i \cdot t^{(n'_{i,1}+n''_{i,1}) \cdot b_1 + \dots + (n'_{i,d}+n''_{i,d}) \cdot b_d} \cdot (\text{Re } Q_{1,1}(\log t))^{n'_{i,1}} \cdot \dots \cdot (\text{Re } Q_{k,d_k}(\log t))^{n'_{i,k}} \cdot (\text{Im } Q_{1,1}(\log t))^{n''_{i,1}} \cdot \dots \cdot (\text{Im } Q_{k,d_k}(\log t))^{n''_{i,k}},$$

where $Q_{i,j}(\cdot)$, as above, are polynomials from the definition of trajectory cones. Note that all $Q_{i,j}$ are only evaluated at real points, and hence, it is easy for us to refer to $\text{Re } Q_{i,j}$ and $\text{Im } Q_{i,j}$; these are polynomials in one real variable with real algebraic coefficients. We rewrite $R_l(\tau(t))$ in the form

$$\sum_i t^{n_{i,1} \cdot b_1 + \dots + n_{i,k} \cdot b_k} \cdot f_i(\log t),$$

where each $f_i(\cdot)$ is also a polynomial with real algebraic coefficients, and $b_1, \dots, b_k$ are distinct logarithms of the moduli of the eigenvalues of A . We can compute all these polynomials f_i , eliminating from the sum of all terms that have $f_i \equiv 0$.

Observe that $R_l(\tau(t))$ is a function of a single variable $t > 0$. In order to reason about the sign of this expression as $t \rightarrow \infty$, we need to determine its leading term. To that end, we first need to decide for every $i \neq j$ whether the two new exponents $n_{i,1}b_1 + \dots + n_{i,k}b_k$ and $n_{j,1}b_1 + \dots + n_{j,k}b_k$ are equal and, if not, which is larger. (If the exponents are equal, we aggregate the polynomials f_i and f_j accordingly.) By rearranging the terms, it is enough to decide whether $n_1b_1 + \dots + n_kb_k > 0$ for some $n_1, \dots, n_k \in \mathbb{Z}$. Recall that $b_j = \log \rho_j$. By Baker's theorem, there exists an effectively computable $\epsilon > 0$ such that either $n_1b_1 + \dots + n_kb_k = 0$ or $|n_1b_1 + \dots + n_kb_k| > \epsilon$.

We now proceed by computing an approximation Δ of $n_1b_1 + \dots + n_kb_k$ with additive error at most $\frac{\epsilon}{3}$. This is easily done, as we are dealing with computable quantities. We then have that $\Delta \in [-\frac{\epsilon}{3}, \frac{\epsilon}{3}]$ iff $n_1b_1 + \dots + n_kb_k = 0$, and otherwise we have $\text{sign}(\Delta) = \text{sign}(n_1b_1 + \dots + n_kb_k)$. Thus, we can sort the exponents $n_{i,1} \cdot b_1 + \dots + n_{i,k} \cdot b_k$ in descending order and, using the same procedure, compare each of them to 0.

Now consider the term that has the largest exponent, m ; suppose this term is $t^m \cdot f_i(\log t)$. Then the sign of $R_l(\tau(t))$ as $t \rightarrow \infty$ is determined by the sign of the leading term of the polynomial $f_i(\cdot)$; only if the sum is empty can the sign of $R_l(\tau(t))$ be 0 for all sufficiently large t .

The argument above shows that we can compute the leading terms of the expressions $R_l(\tau(t))$ and decide whether the conjunction $\bigwedge_{l=1}^m R_l \sim_l 0$ holds for all $t \geq t_0$ starting from some t_0 . This completes the proof.

6.2.1 Existence of Semi-Algebraic Invariants for Semi-Algebraic Halting Sets. We now proceed to show that in the case of a semi-algebraic halting set, the existence of an o-minimal invariant implies the existence of a semi-algebraic invariant (note that clearly the other implication is trivial).

Intuitively, the trajectory cone C_{t_0} is not already a semi-algebraic set for two "reasons:" the $\log t$ factors, and the possibly-irrational exponents b_i . In the following, we over-approximate these

factors by semi-algebraic components. However, as we will show, the approximation must carefully take into account the relationships between the exponents.

Let $\epsilon \in (0, 1)$, and consider the trajectory cone C_{t_0} . We define an over-approximating set of C_{t_0} by replacing the $\log t$ factors by an interval ranging from some constant lower bound μ to t^ϵ . That is, define for $\mu \in \mathbb{Q}$ and $t_0 \in \mathbb{R}$

$$\tilde{C}_{t_0, \epsilon, \mu} = \left\{ \begin{pmatrix} t^{b_1} p_1 Q_{1,1}(s) \\ \vdots \\ t^{b_k} p_k Q_{k,d_k}(s) \end{pmatrix} : (p_1, \dots, p_d) \in \mathbb{T}, t \geq t_0, \mu \leq s \leq t^\epsilon \right\}.$$

Next, we modify the irrational b_i exponents into rational ones. This is done in two parts. First, we approximate b_i by lower and upper rational bounds, next we enforce additive relationships among the approximations.

Consider vectors $\ell = (\ell_1, \dots, \ell_k), \mathbf{u} = (u_1, \dots, u_k) \in \mathbb{Q}^k$ with $\ell \leq \mathbf{b} \leq \mathbf{u}$ and define

$$\text{Box}(\ell, \mathbf{u}) = \{\mathbf{c} \in \mathbb{R}^k : \ell \leq \mathbf{c} \leq \mathbf{u}\}.$$

Second, let $\mathbf{b} = (b_1, \dots, b_k)$, and define

$$\mathbb{S} = \{\mathbf{c} \in \mathbb{R}^k : \forall \mathbf{z} \in \mathbb{Z}^k, \mathbf{b} \cdot \mathbf{z} = 0 \text{ implies } \mathbf{c} \cdot \mathbf{z} = 0\},$$

to be the set of vectors that maintain the integer additive relationships among the b_i .

We are now ready to define the *fat trajectory cone*, which is a semi-algebraic set that approximates $\tilde{C}_{t_0, \epsilon, \mu}$. Given t_0, ϵ, μ, ℓ , and $\mathbf{u}$, we define

$$\mathcal{F}_{t_0, \epsilon, \mu}^{\ell, \mathbf{u}} = \left\{ \begin{pmatrix} t^{c_1} p_1 Q_{1,1}(s) \\ \vdots \\ t^{c_k} p_k Q_{k,d_k}(s) \end{pmatrix} : (p_1, \dots, p_d) \in \mathbb{T}, t \geq t_0, \mu \leq s \leq t^\epsilon, \mathbf{c} \in \mathbb{S} \cap \text{Box}(\ell, \mathbf{u}) \right\}.$$

That is, we replace the exponents vector $\mathbf{b}$ with a vector $\mathbf{c}$ that is close enough to $\mathbf{b}$, and maintains its additive relations.

It is not immediate from the definition of $\mathcal{F}_{t_0, \epsilon, \mu}^{\ell, \mathbf{u}}$ that it is indeed a semi-algebraic set, nor that we can compute a representation of it. Indeed, we cannot quantify the exponents $\mathbf{c}$ in the first-order theory of the reals, and it is not clear that the set $\mathbb{S}$ can be finitely represented. We start by addressing these issues.

LEMMA 6.7. $\mathcal{F}_{t_0, \epsilon, \mu}^{\ell, \mathbf{u}}$ is definable in $\mathfrak{R}_0$, and we can compute a representation of it.

PROOF. Recall that $\mathbf{b} = (b_1, \dots, b_k)$ where $b_i = \log \rho_i$ for every i . Consider the abelian group $L = \{\mathbf{z} \in \mathbb{Z}^k : \varrho^{\mathbf{z}} = 1\}$ where $\varrho^{\mathbf{z}} = \rho_1^{z_1} \cdots \rho_k^{z_k}$. By [24] we can compute a finite basis $\{\mathbf{z}^1, \dots, \mathbf{z}^m\}$ for L .

Note that for every $\mathbf{z} \in \mathbb{Z}^k$, we have that $\mathbf{b} \cdot \mathbf{z} = 0$ iff $\varrho^{\mathbf{z}} = 1$. Thus, we can write

$$\mathbb{S} = \{\mathbf{c} \in \mathbb{R}^k : \forall \mathbf{z} \in \mathbb{Z}^k, \mathbf{z} \in L \text{ implies } \mathbf{c} \cdot \mathbf{z} = 0\} = \left\{ \mathbf{c} \in \mathbb{R}^k : \bigwedge_{i=1}^m \mathbf{c} \cdot \mathbf{z}^i = 0 \right\}.$$

Let $\mathbb{S}' = \{\mathbf{r} \in \mathbb{R}^k : \bigwedge_{i=1}^m \mathbf{r} \cdot \mathbf{z}^i = 1\}$ where, as before, $\mathbf{r} \cdot \mathbf{z} = r_1^{z_1} \cdots r_k^{z_k}$. Consider some $t > 1$. For every $\mathbf{c} \in \mathbb{R}^k$, denote $t^{\mathbf{c}} = (t^{c_1}, \dots, t^{c_k})$, then clearly $\mathbf{c} \in \mathbb{S}$ iff $t^{\mathbf{c}} \in \mathbb{S}'$. In addition, $\mathbf{c} \in \text{Box}(\ell, \mathbf{u})$ iff $t^{\mathbf{c}} \in \text{Box}(t^\ell, t^{\mathbf{u}})$. We conclude that we can write

$$\mathcal{F}_{t_0, \epsilon, \mu}^{\ell, \mathbf{u}} = \left\{ \begin{pmatrix} r_1 p_1 Q_{1,1}(s) \\ \vdots \\ r_k p_k Q_{k,d_k}(s) \end{pmatrix} : (p_1, \dots, p_d) \in \mathbb{T}, t \geq t_0, \mu \leq s \leq t^\epsilon, \mathbf{r} \in \mathbb{S}' \cap \text{Box}(t^\ell, t^{\mathbf{u}}) \right\}.$$

Since $\mathbb{T}$, $\mathbb{S}'$ and $\text{Box}(t^\ell, t^u)$ are all semi-algebraic sets (the latter due to ℓ and u being rational vectors), then so is $\mathcal{F}_{t_0, \epsilon, \mu}^{\ell, u}$. $\square$

The following lemma is the main result of this section, and states that if $P \cdot C_{t_0}$ is disjoint from a semi-algebraic halting set, then we can approximate it by some appropriate fat cone.

LEMMA 6.8. *Let Y be a semi-algebraic set such that $P \cdot C_{t_0} \cap Y = \emptyset$ for some $t_0 \in \mathbb{R}$, then there exist $\mu, t_1 \in \mathbb{R}$, $\epsilon > 0$, and $\ell, u \in \mathbb{Q}^d$ such that $P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u} \cap Y = \emptyset$ and $A \cdot P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u} \subseteq P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u}$.*

The rest of the section is devoted to the proof of Lemma 6.8. We start by showing that the fat cone is invariant under J .

LEMMA 6.9. *For every ϵ , there exists t_0 such that for every $t_1 \geq t_0$ and for every μ, ℓ , and u we have that $J \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u} \subseteq \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u}$.*

PROOF. Following the proof of Lemma 4.3, we see that for $y = \begin{pmatrix} t^{c_1} p_1 Q_{1,1}(s) \\ \vdots \\ t^{c_k} p_k Q_{k,d_k}(s) \end{pmatrix}$ we have that $Jy = \begin{pmatrix} e^{b_1 t^{c_1}} \xi_1 p_1 Q_{1,1}(s+1) \\ \vdots \\ e^{b_k t^{c_k}} \xi_k p_k Q_{k,d_k}(s+1) \end{pmatrix}$. Thus, it suffices to show that for large enough t_0 , for every $t \geq t_0$ the following hold:

- (1) If $\mu \leq s \leq t^\epsilon$ then $\mu \leq s+1 \leq (et)^\epsilon$.
- (2) $(e^{b_1} t^{c_1}, \dots, e^{b_k} t^{c_k}) = ((et)^{c'_1}, \dots, (et)^{c'_k})$ for some $c' \in \mathbb{S} \cap \text{Box}(\ell, u)$.

For item 1, observe that since $\mu \leq s$, we clearly have $\mu \leq s+1$. For the second inequality, since $\epsilon > 0$, we have that $e^\epsilon - 1 > 0$. Let t_0 be large enough such that $(e^\epsilon - 1)t^\epsilon \geq 1$ for every $t \geq t_0$, then for every $t \geq t_0$ we have that $t^\epsilon + 1 \leq (et)^\epsilon$. It follows that $s+1 \leq t^\epsilon + 1 \leq (et)^\epsilon$, as desired.

For item 2, define $c' = (c'_1, \dots, c'_k)$ by setting $c'_i = \frac{b_i + \log(t)c_i}{1 + \log(t)}$. It is easy to check that $e^{b_i} t^{c_i} = (et)^{c'_i}$ for every $1 \leq i \leq k$. Furthermore, since $\mathbb{S} \cap \text{Box}(\ell, u)$ is a convex set, and c' is a convex combination of b and c , it follows that $c' \in \mathbb{S} \cap \text{Box}(\ell, u)$, and we are done. $\square$

It remains to prove that we can choose a fat cone that is disjoint from the target set.

LEMMA 6.10. *Let Y be a semi-algebraic set, and let $t_0 \in \mathbb{R}$ be such that $PC_{t_0} \cap Y = \emptyset$, then there exist $\epsilon > 0$, $\mu \in \mathbb{R}$, $t_1 \in \mathbb{R}$, and $\ell, u \in \mathbb{Q}^k$ such that $P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u} \cap Y = \emptyset$.*

PROOF. Working along the lines of the proof of Lemma 6.5, we define

$$U = \{(z_1, \dots, z_d)^\top \in \mathbb{C}^d : \forall (p_1, \dots, p_d) \in \mathbb{T}, P(z_1 p_1, \dots, z_d p_d)^\top \in \mathbb{R}^d \setminus Y\},$$

and let φ be a quantifier-free formula defining U . As in the proof of Lemma 6.5, let $\tau(t) = (t^{b_1} Q_{1,1}(\log t), \dots, t^{b_k} Q_{k,d_k}(\log t))$, then for the purpose of evaluating $\varphi(\tau(t))$, we can assume φ is of the form $\bigwedge_{l=1}^m R_l(u_1, \dots, u_d, v_1, \dots, v_d) \sim_l 0$. Since $P \cdot C_{t_0} \cap Y = \emptyset$, then $R_l(\tau(t)) \sim_l 0$ for every l and every $t \geq t_0$. Writing

$$R_l(\tau(t)) = \sum_i t^{n_{i,1} \cdot b_1 + \dots + n_{i,k} \cdot b_k} \cdot f_i(\log t), \quad (4)$$

we show, first, how to replace the exponents vector b by any exponents vector in $\mathbb{S} \cap \text{Box}(\ell, u)$ for appropriate ℓ, u , and second, how to replace $\log t$ by s for $\mu \leq s \leq t^\epsilon$ for some appropriate μ and ϵ , while maintaining the inequality or equality prescribed by $\sim_l$.

Denote by N the set of vectors $\mathbf{n}_i = (n_{i,1}, \dots, n_{i,k})$ of exponents in Equation (4). As in the proof of Lemma 6.5, we can compute $\delta > 0$ such that for every $\mathbf{n}, \mathbf{n}' \in N$, if $\mathbf{b} \cdot (\mathbf{n} - \mathbf{n}') \neq 0$ then $|\mathbf{b} \cdot (\mathbf{n} - \mathbf{n}')| > \delta$. Let $M = \max_{\mathbf{n}, \mathbf{n}' \in N} \|\mathbf{n} - \mathbf{n}'\|$ (where $\|\cdot\|$ is the Euclidean norm in $\mathbb{R}^k$).

CLAIM 4. *Let $\mathbf{c} \in \mathbb{R}^d$ be such that $\|\mathbf{b} - \mathbf{c}\| \leq \frac{\delta}{2M}$, then, for all $\mathbf{n}, \mathbf{n}' \in N$, if $\mathbf{b} \cdot (\mathbf{n} - \mathbf{n}') > 0$ then $\mathbf{c} \cdot (\mathbf{n} - \mathbf{n}') > \frac{\delta}{2}$.*

PROOF OF CLAIM 4. Suppose that $\mathbf{b} \cdot (\mathbf{n} - \mathbf{n}') > 0$, then by the above we have $\mathbf{b} \cdot (\mathbf{n} - \mathbf{n}') > \delta$, and hence,

$$\mathbf{c} \cdot (\mathbf{n} - \mathbf{n}') = \mathbf{b} \cdot (\mathbf{n} - \mathbf{n}') + (\mathbf{c} - \mathbf{b}) \cdot (\mathbf{n} - \mathbf{n}') \geq \delta - \|\mathbf{c} - \mathbf{b}\| \cdot \|\mathbf{n} - \mathbf{n}'\| \geq \delta - \frac{\delta}{2M} M = \frac{\delta}{2}. \quad \square$$

We can now choose ℓ and $\mathbf{u}$ such that $u_i - \ell_i \leq \frac{\delta}{2M\sqrt{k}}$ so that for all $\mathbf{c} \in \text{Box}(\ell, \mathbf{u})$ we have

$$\|\mathbf{b} - \mathbf{c}\| \leq \sqrt{\sum_{i=1}^k (u_i - \ell_i)^2} \leq \sqrt{\frac{\delta^2}{(2M)^2}} = \frac{\delta}{2M}.$$

It follows from Claim 4 and from the definition of $\mathbb{S}$ that, intuitively, every $\mathbf{c} \in \text{Box}(\ell, \mathbf{u})$ maintains the order of magnitude of the monomials $t^{n_{i,1} \cdot b_1 + \dots + n_{i,k} \cdot b_k}$ in $R_l(\tau(t))$.

More precisely, let $\tau'(t) = (t^{c_1} Q_{1,1}(\log t), \dots, t^{c_k} Q_{k,d_k}(\log t))$ for some $\mathbf{c} \in \text{Box}(\ell, \mathbf{u})$. Then the exponent of the ratio of every two monomials in $R_l(\tau'(t))$ has the same (constant) sign as the corresponding exponent in $R_l(\tau(t))$. Moreover, the exponents of distinct monomials in $R_l(\tau(t))$ differ by at least $\frac{\delta}{2}$ in $R_l(\tau'(t))$.

We are now ready to handle the $\log t$ factor. First, assume t_0 is large enough that $f_i(\log t)$ has constant sign for every $t \geq t_0$ (otherwise increase t_0 accordingly). We can now let μ be large enough that for every $s \geq \mu$, the sign of $f_i(\log t)$ coincides with the sign of $f_i(s)$ for every $t \geq t_0$. It remains to give an upper bound on s of the form t^ϵ such that substituting $f_i(s)$ instead of $f_i(\log t)$ does not change the ordering of the terms (by their magnitude) in $R_l(\tau'(t))$. Let D be the maximum degree of all polynomials f_i in Equation (4), and define $\epsilon = \frac{\delta}{3D}$ (in fact, any $\epsilon < \frac{\delta}{2D}$ would suffice), then we have that, for $t \geq t_0$, $f_i(s)$ has the same sign as $f_i(\log t)$ for every $\mu \leq s \leq t^\epsilon$ (by our choice of μ), and guarantees that substituting t^ϵ instead of t does not change the ordering of the terms (by their magnitude) in R_l . Since the exponents of the monomials in $R_l(\tau'(t))$ differ by at least $\frac{\delta}{2}$, it follows that their order is maintained when replacing $\log t$ by s .

Let $\tau''(t) = (t^{c_1} Q_{1,1}(s), \dots, t^{c_k} Q_{k,d_k}(s))$ for some $\mathbf{c} \in \text{Box}(\ell, \mathbf{u})$ and $\mu \leq s \leq t^\epsilon$, then by our choice of ϵ , the dominant term in $R_l(\tau''(t))$ is the same as that in $R_l(\tau(t))$. Therefore, for large enough t , the signs of $R_l(\tau''(t))$ and $R_l(\tau(t))$ are the same.

By repeating this argument for each R_l , we can compute $t_1 \in \mathbb{R}$, $\epsilon > 0$, $\mu \in \mathbb{R}$, and $\ell, \mathbf{u} \in \mathbb{Q}^k$ such that $P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, \mathbf{u}} \cap Y = \emptyset$, and we are done. $\square$

We are now ready to complete the proof of Theorem 6.3.

PROOF OF THEOREM 6.3. By Theorem 6.4, we can decide whether there exists an o-minimal invariant for the LDS (A, s) that avoids a semi-algebraic target Y . Moreover, in positive instances, we can synthesize such an invariant of the form $\mathcal{I} = P \cdot C_{t_0} \cup \{A^n s \mid 0 \leq n \leq \log t_0\}$ for some $t_0 \geq 1$.

In particular, $P \cdot C_{t_0} \cap Y = \emptyset$, so by Lemma 6.8, there exist $\mu, t_1 \in \mathbb{R}$, $\epsilon > 0$, and $\ell, \mathbf{u} \in \mathbb{Q}^d$ such that $P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, \mathbf{u}} \cap Y = \emptyset$ and $P \cdot \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, \mathbf{u}}$ is invariant under A . Furthermore, by Lemma 6.7, $\mathcal{F}_{t_1, \epsilon, \mu}^{\ell, \mathbf{u}}$ is semi-algebraic.

Naively, one might think that $\mathcal{F}_{t_1, \epsilon, \mu}^{\ell, \mathbf{u}}$ with the addition of the “finite tail” $\{A^n s \mid 0 \leq n \leq \log t_1\}$ would make up a semi-algebraic invariant for (A, s) . This, however, may not be the case. Recall

that

$$\mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u} = \left\{ \begin{pmatrix} r_1 p_1 Q_{1,1}(s) \\ \vdots \\ r_k p_k Q_{k,d_k}(s) \end{pmatrix} : (p_1, \dots, p_d) \in \mathbb{T}, t \geq t_1, \mu \leq s \leq t^\epsilon, \mathbf{r} \in \mathbb{S}' \cap \text{Box}(t^\ell, t^u) \right\},$$

and observe that if μ is large and ϵ is small, then for small enough $t \geq t_1$, it may be the case that $t^\epsilon < \mu$, so there does not exist an appropriate s . Thus, we may need to extend the finite tail to capture this.

Let $t_2 \geq t_1$ such that for every $n \geq \log t_2$ we have that $A^n s \in \mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u}$ (note that we can compute t_2 by iterating over the orbit, until the first point in $\mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u}$ is found), then $\mathcal{F}_{t_1, \epsilon, \mu}^{\ell, u} \cup \{A^n s \mid 0 \leq n \leq \log t_2\}$ is a semi-algebraic invariant for (A, s) that avoids Y , which concludes the proof. $\square$

6.3 A Note on Complexity

The proof of Theorem 6.3 does not offer a complexity bound. We now provide a brief analysis of the complexity. We recall that the algorithmic bottlenecks in the proof of Theorem 6.3 are the following:

- (1) Computing a basis for the group of multiplicative relations of the eigenvalues, hence obtaining a description of $\mathbb{T}$.
- (2) Computing the cell decomposition of the set U as per Section 6.2.
- (3) Computing $\epsilon > 0$ to separate the exponents in the description of $R_l(\tau(t))$, as above.
- (4) Concluding whether there exists $t_0 > 0$ such that $P \cdot C_{t_0} \cap F = \emptyset$.
- (5) Checking whether $A^n s \in F$ for some $n \leq \log t_0$.

In the following, we use $\|\cdot\|$ to denote encoding length, when the type of encoding is understood from context.

By Masser's Theorem (Theorem 3.1), computing the basis for the group of multiplicative relations can be done in polynomial space in $\|A\|$, and ends up with a basis whose description length is polynomial in $\|A\|$.

Next, by Theorem 1.2 of [31], computing the cell decomposition of U takes double-exponential time in $\|F, A\|$ (where F is the halting set) and the resulting expression consists of an expression of the form $\bigvee_{i=1}^I \bigwedge_{j=1}^J R_{i,j} \sim_{i,j} 0$, as in Section 6.2, where both I and J are double-exponential in $\|F, A\|$, as well as the degrees of the polynomials $R_{i,j}$, and the description length of their coefficients. In particular, this means that there are double-exponential disjuncts of the form $\bigwedge R_l \sim_l 0$ to reason about in the proof of Lemma 6.5.

In order to substitute the latter into the bound given by Baker's Theorem (Theorem 6.6), we actually need a more explicit form of the bound H^{-C} therein (see [1]). Specifically, in the setting of Theorem 6.6, we have

$$|\Lambda| > \exp\left(-(16mD)^{2(m+2)} \log H \log B\right),$$

where B is an upper bound on the integer coefficients, and D is the degree of the extension field $\mathbb{Q}(\alpha_1, \dots, \alpha_m)$ over $\mathbb{Q}$.

Recalling that in our case $D \leq k!$, and B has double-exponential description length (i.e., $\log B$ is double exponential in $\|F, A\|$), we get that ϵ itself has double-exponential description length in $\|F, A\|$. Observe that ϵ is efficiently computable since all the functions involved are efficiently computable (in particular, the logarithm of an algebraic number is efficiently computable by [17]). Thus, we can determine the dominant exponent of each R_l in double-exponential time in $\|F, A\|$.

We remark that tighter bounds are possible if the halting set, F , is given as a first-order formula with bounded alternation depth.

Recall that now, determining the existence of t_0 such that $P \cdot C_{t_0} \cap F = \emptyset$ amounts to checking the sign of the leading coefficient of f_i , which is simple. However, in order to obtain a concrete bound on t_0 , we must determine when the leading monomial in $R_I(\tau(t))$ actually dominates the rest of the terms.

Recall that the exponents of the monomials in $R_I(\tau(t))$ differ from each other by at least ϵ , and that the coefficients of the monomials are polylogarithmic in t . Thus, an upper bound on t_0 is obtained in the worst case when

$$R_I(\tau(t)) = t^n - \beta \cdot t^{n-\epsilon} \cdot \log^\kappa t,$$

where β has double-exponential description length in $\|F, A\|$ and $|\kappa|$ is double exponential in $\|F, A\|$. Indeed, this implies that the leading monomial has the slowest growth, and the second-highest monomial has the fastest.

Thus, we can take t_0 to be the maximal solution of $t^n = \beta \cdot t^{n-\epsilon} \cdot \log^\kappa t$. By dividing by $t^{n-\epsilon}$, this is equivalent to $t^\epsilon = \beta \log^\kappa t$, and by setting $t = x^{\kappa/\epsilon}$ this is equivalent in turn to $x = c \log x$ with $c = \kappa \cdot \beta^{1/\kappa}/\epsilon$. Let $\psi(x) = x/\log x$, then we wish to bound the maximal solution to $\psi(x) = c$. We show that the maximal solution is at most $2c \log c$ (this is also proved in [35, Lemma A.1]).

Recall that $\psi(x)$ is increasing for $x \geq e$. We can assume w.l.o.g. that $c > e$. Indeed, if $c \leq e = \psi(e)$ then the maximal solution for $\psi(x) = c$ is with $x \leq e$ and we are done.

Observe that $\psi(2c \log c) = c \cdot 2 \log c / (\log c + \log \log c + \log 2) > c$ as long as $\log c > \log \log c + \log 2$, which is the case for $c \geq e$, in particular. So the Equation $\psi(x) = c$ has no solutions with $x \geq 2c \log c$. Therefore, $t_0 \leq (2c \log c)^{\kappa/\epsilon}$. In particular, c and t_0 are at most triply exponential in $\|F, A\|$, with doubly exponential description lengths in $\|F, A\|$.

Finally, recall that the last algorithmic step is to check whether $A^n s \in F$ for $n \leq \log t_0$. Each such check may potentially take double-exponential time in $\|F, A\|$, and there are double-exponential checks to take.

Thus, overall, the entire algorithm can be implemented in double-exponential time.

REFERENCES

- [1] Alan Baker and Gisbert Wüstholz. 1993. Logarithmic forms and group varieties. *Journal für die reine und Angewandte Mathematik* 442, 3 (1993), 19–62.
- [2] Amir M. Ben-Amram and Samir Genaim. 2014. Ranking functions for linear-constraint loops. *Journal of the ACM* 61, 4 (2014), 26:1–26:55.
- [3] Amir M. Ben-Amram, Samir Genaim, and Abu Naser Masud. 2012. On the termination of integer loops. *ACM Transactions on Programming Languages and Systems* 34, 4 (2012), 16:1–16:24.
- [4] Mark Braverman. 2006. Termination of integer linear programs. In *Proceedings of the 18th International Conference on Computer Aided Verification*. 372–385.
- [5] Jin-yi Cai, Richard J. Lipton, and Yechezkel Zalcstein. 2000. The complexity of the A B C problem. *SIAM Journal on Computing* 29, 6 (2000), 1878–1888.
- [6] John W. S. Cassels. 1965. *An Introduction to Diophantine Approximation*. Cambridge University Press.
- [7] Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2013. The orbit problem in higher dimensions. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*. 941–950.
- [8] Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2015. The polyhedron-hitting problem. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*. 940–956.
- [9] Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2016. On the complexity of the orbit problem. *Journal of the ACM* 63, 3 (2016), 23:1–23:18.
- [10] Michael Colón, Sriram Sankaranarayanan, and Henny Sipma. 2003. Linear invariant generation using non-linear constraint solving. In *Proceedings of the 15th International Conference on Computer Aided Verification*. 420–432.

- [11] Patrick Cousot. 2005. Proving program invariance and termination by parametric abstraction, lagrangian relaxation and semidefinite programming. In *Proceedings of the 6th International Conference on Verification, Model Checking, and Abstract Interpretation*, 1–24.
- [12] Patrick Cousot and Nicolas Halbwachs. 1978. Automatic discovery of linear restraints among variables of a program. In *Proceedings of the 5th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. 84–96.
- [13] L. P. D. van den Dries. 1998. *Tame Topology and O-minimal Structures*. Cambridge University Press.
- [14] Nathanaël Fijalkow, Engel Lefauchaux, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2019. On the monniaux problem in abstract interpretation. In *Proceedings of the 26th International Symposium on Static Analysis*.
- [15] Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2017. Semialgebraic invariant synthesis for the kannan-lipton orbit problem. In *Proceedings of the 34th Symposium on Theoretical Aspects of Computer Science*. 29:1–29:13.
- [16] Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2019. Complete semialgebraic invariant synthesis for the kannan-lipton orbit problem. *Theory of Computing Systems* 63, 5 (2019), 1027–1048.
- [17] Esther Galby, Joël Ouaknine, and James Worrell. 2015. On matrix powering in low dimensions. In *Proceedings of the 32nd International Symposium on Theoretical Aspects of Computer Science*.
- [18] Ashutosh Gupta, Thomas A. Henzinger, Rupak Majumdar, Andrey Rybalchenko, and Ru-Gang Xu. 2008. Proving non-termination. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 147–158.
- [19] Mehran Hosseini, Joël Ouaknine, and James Worrell. 2019. Termination of linear loops over the integers. In *Proceedings of the 46th International Colloquium on Automata, Languages, and Programming*, 118:1–118:13.
- [20] Ravindran Kannan and Richard J. Lipton. 1980. The orbit problem is decidable. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing*. 252–261.
- [21] Ravindran Kannan and Richard J. Lipton. 1986. Polynomial-time algorithm for the orbit problem. *Journal of the ACM* 33, 4 (1986), 808–821.
- [22] Zachary Kincaid, John Cyphert, Jason Breck, and Thomas W. Reps. 2018. Non-linear reasoning for invariant synthesis. *Proceedings of the ACM on Programming Languages* 2, POPL (2018), 54:1–54:33.
- [23] Angus Macintyre and Alex J. Wilkie. 1996. On the decidability of the real exponential field. In *Proceedings of the Kreiseliana. About and Around Georg Kreisel*, Piergiorgio Odifreddi (Ed.), A K Peters, 441–467.
- [24] David W. Masser. 1988. Linear relations on algebraic groups. *New Advances in Transcendence Theory*. Cambridge University Press.
- [25] M. Mignotte, T. Shorey, and R. Tijdeman. 1984. The distance between terms of an algebraic recurrence sequence. *Journal für Die Reine und Angewandte Mathematik* 1984, 349 (1984), 63–76.
- [26] Joël Ouaknine, João Sousa Pinto, and James Worrell. 2015. On termination of integer linear loops. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms*, 957–969.
- [27] Joël Ouaknine and James Worrell. 2014. On the positivity problem for simple linear recurrence sequences. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming*. 318–329.
- [28] Joël Ouaknine and James Worrell. 2014. Positivity problems for low-order linear recurrence sequences. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*. 366–379.
- [29] Joël Ouaknine and James Worrell. 2014. Ultimate positivity is decidable for simple linear recurrence sequences. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming*. 330–341.
- [30] Joël Ouaknine and James Worrell. 2015. On linear recurrence sequences and loop termination. *SIGLOG News* 2, 2 (2015), 4–13.
- [31] James Renegar. 1992. On the computational complexity and geometry of the first-order theory of the reals. Part I: Introduction. Preliminaries. The geometry of semi-algebraic sets. The decision problem for the existential theory of the reals. *Journal of Symbolic Computation* 13, 3 (1992), 255–299.
- [32] Enric Rodríguez-Carbonell and Deepak Kapur. 2004. An abstract interpretation approach for automatic generation of polynomial invariants. In *Proceedings of the 11th International Symposium on Static Analysis*. 280–295.
- [33] Enric Rodríguez-Carbonell and Deepak Kapur. 2007. Generating all polynomial invariants in simple loops. *Journal of Symbolic Computation* 42, 4 (2007), 443–476.
- [34] Sriram Sankaranarayanan, Henny Sipma, and Zohar Manna. 2004. Non-linear loop invariant generation using Gröbner bases. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 318–329.
- [35] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- [36] T. Tao. 2008. *Structure and Randomness: Pages from Year One of a Mathematical Blog*. American Mathematical Society.
- [37] Alfred Tarski. 1951. *A Decision Method for Elementary Algebra and Geometry*. RAND Corporation.

- [38] Ashish Tiwari. 2004. Termination of linear programs. In *Proceedings of the 16th International Conference on Computer Aided Verification*. 70–82.
- [39] N. K. Vereshchagin. 1985. The problem of appearance of a zero in a linear recurrence sequence (in Russian). *Matematicheskie Zametki* 38, 2 (1985), 609–615.
- [40] A. J. Wilkie. 1996. Model completeness results for expansions of the ordered field of real numbers by restricted pfaffian functions and the exponential function. *Journal of the American Mathematical Society* 9, 4 (1996), 1051–1094.
- [41] Bican Xia and Zhihai Zhang. 2010. Termination of linear programs with nonlinear constraints. *Journal of Symbolic Computation* 45, 11 (2010), 1234–1249.

Received June 2020; revised June 2021; accepted November 2021