O-Minimal Invariants for Linear Loops*

Shaull Almagor[1], Dmitry Chistikov[2], Joël Ouaknine[3], and James Worrell[4]

1 **Department of Computer Science, Oxford University, UK**
shaull.almagor@cs.ox.ac.uk
2 **Centre for Discrete Mathematics and its Applications (DIMAP) & Department of Computer Science, University of Warwick, UK**
d.chistikov@warwick.ac.uk
3 **Max Planck Institute for Software Systems, Germany & Department of Computer Science, Oxford University, UK**
joel@mpi-sws.org
4 **Department of Computer Science, Oxford University, UK**
jbw@cs.ox.ac.uk

Abstract

The termination analysis of linear loops plays a key rôle in several areas of computer science, including program verification and abstract interpretation. Such deceptively simple questions also relate to a number of deep open problems, 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 paper, 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.

1998 ACM Subject Classification F.3.1 Specifying and Verifying and Reasoning about Programs

Keywords and phrases Invariants, linear loops, linear dynamical systems, non-termination, o-minimality

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.

1 Introduction

This paper is concerned with the existence and algorithmic synthesis of suitable *invariants* for linear loops, or equivalently for discrete-time linear dynamical systems. 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., [17], and particularly Sec. 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 following form:

$$P: \quad x \leftarrow s; \ \ \texttt{while}\ x \notin F\ \texttt{do}\ x \leftarrow Ax\,, \tag{1}$$

* Joël Ouaknine was supported by ERC grant AVS-ISS (648701), and James Worrell was supported by EPSRC Fellowship EP/N008197/1.

45th International Colloquium on Automata, Languages, and Programming (ICALP 2018).
Editors: Ioannis Chatzigiannakis, Christos Kaklamanis, Daniel Marx, and Don Sannella;
Article No. ; pp. :1–:15

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

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., [3, 2, 24]. 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 [15, 16]. This work was subsequently extended to instances in which F is a low-dimensional vector space [6, 8] or a low-dimensional polyhedron [7]. 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 [28, §3.9], although once again positive results are known in low dimensions [19, 31]. 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 [22, 21].

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 $S = \mathbb{R}^d$ and $S = \mathbb{Q}^d$ are known [30, 4]; see also [20]. In the vast majority of cases, however, termination is a hard problem (and often undecidable [33]), 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 (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 [14], 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., [11, 9, 27, 25, 10, 26, 13].

In order to make this notion more precise, let us associate with our loop P a *discrete-time linear dynamical system* (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., $\mathcal{I} \cap F = \emptyset$.

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

In instances of non-termination, one notes that the orbit $\mathcal{O}$ 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 [11], algebraic sets [26], and semi-algebraic sets [13].

Main contributions. We focus on loops of the form given in (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. We also consider two large classes of halting sets, namely semi-algebraic sets, as well as sets definable in the first-order theory of the reals with exponentiation, denoted $\mathfrak{R}_{\exp}$. 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 $\mathcal{I}$ containing s and disjoint from F, and moreover in positive instances such an invariant can be defined explicitly in $\mathfrak{R}_{\exp}$; for the more general case in which F is $\mathfrak{R}_{\exp}$-definable, the same holds assuming Schanuel's conjecture.

We illustrate below some of the key ideas from our approach. Consider a linear dynamical system (A, s) with $A \in \mathbb{Q}^{3\times 3}$ whose orbit $\mathcal{O}$ is depicted in Figure 1. In our example, $\mathcal{O}$ 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 $\mathcal{O}$, 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 $\mathfrak{R}_{\exp}$-definable o-minimal invariant by simply adjoining a finite number of orbit points to the truncated trajectory cone, as depicted in Figure 4.

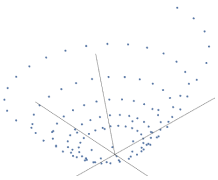

Figure 1 The orbit $\mathcal{O}$ of (A, s)

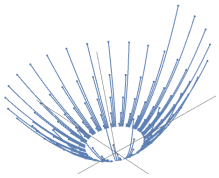

Figure 2 Trajectory rays of $\mathcal{O}$.

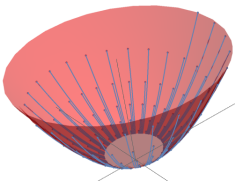

Figure 3 Trajectory cone for $\mathcal{O}$.

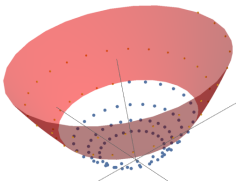

Figure 4 Invariant for $\mathcal{O}$.

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 paper.

The work that is closest to ours in the literature is [13], which considers 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 a procedure for deciding the existence of semi-algebraic invariants. The present paper 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 paper correspondingly makes heavy use of model-theoretic and number-theoretic tools that are entirely absent from [13]. It is interesting to note, however, that the question of the existence of semi-algebraic (rather than o-minimal) invariants in the present setting appears to be a challenging open problem.

2 Preliminaries

The *first-order theory of the reals*, denoted $\mathfrak{R}_0$, is the collection of true sentences in the first-order logic of the structure $\langle \mathbb{R}, 0, 1, +, \cdot, < \rangle$. Sentences in $\mathfrak{R}_0$ are quantified Boolean combinations of atomic propositions of the form $P(x_1, \ldots, x_n) > 0$ where P is a polynomial with integer coefficients, and $x_1, \ldots, x_n$ are variables. Tarski famously showed that this theory admits quantifier elimination [29] and is therefore decidable. In addition to $\mathfrak{R}_0$, we also consider the *first-order theory of the reals with exponentiation*, denoted $\mathfrak{R}_{\exp}$, which augments $\mathfrak{R}_0$ with the exponentiation function $x \mapsto e^x$.

A set $S \subseteq \mathbb{R}^d$ is *definable* in a theory $\mathfrak{R}$ if there exists a formula $\varphi(x_1, \ldots, x_d)$ in $\mathfrak{R}$ with free variables $x_1, \ldots, x_d$ such that $S = \{(c_1, \ldots, c_d) \in \mathbb{R}^d \mid \varphi(c_1, \ldots, c_d) \text{ is true}\}$. A function $f \colon B \to \mathbb{R}^m$ 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 first-order theory of the reals, $\mathfrak{R}_0$-definable sets (resp. functions) are known as *semi-algebraic* sets (resp. functions).

A theory $\mathfrak{R}$ is said to be *o-minimal* if every $\mathfrak{R}$-definable subset of the reals $S \subseteq \mathbb{R}$ is a finite union of points and (possibly unbounded) intervals.

▶ **Definition 1.** A set $S \subseteq \mathbb{R}^d$ is *o-minimal* if it is definable in some o-minimal theory that extends $\mathfrak{R}_{\exp}$.

Tarski's result on quantifier elimination [29] also implies that $\mathfrak{R}_0$ is o-minimal. The o-minimlity of $\mathfrak{R}_{\exp}$, on the other hand, is due to Wilkie [32]. O-minimal theories enjoy many useful properties, some of which we list below, referring the reader to [12] for precise definitions and proofs. In what follows, $\mathfrak{R}$ is a fixed o-minimal theory.

1. For an $\mathfrak{R}$-definable set $S \subseteq \mathbb{R}^d$, its topological closure $\overline{S}$ is also $\mathfrak{R}$-definable.
2. For an $\mathfrak{R}$-definable function $f \colon S \to \mathbb{R}$, the number $\inf\{f(x) \mid x \in S\}$ is $\mathfrak{R}$-definable (as a singleton set).
3. O-minimal theories 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 $m \in \{0, 1, \ldots, d\}$. 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 $f \colon S \to \mathbb{R}^m$, 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-algberaic set, we can effectively compute its cell decomposition. Unfortunately, few more expressive theories are known to be decidable. The theory $\mathfrak{R}_{\exp}$ is decidable provided that Schanuel's conjecture, an assertion in transcendental number theory, holds [18]. Our decidability result in Theorem 11 is subject to Schanuel's conjecture; somewhat surprisingly, however, we exhibit in Theorem 12 an unconditional decidability result.

▶ Remark. While all our $\mathfrak{R}$-definable sets live in $\mathbb{R}^d$, it is often convenient or necessary to consider sets in $\mathbb{C}^d$. To this end, by identifying $\mathbb{C}$ with $\mathbb{R}^2$, we define a set $S \subseteq \mathbb{C}^d$ to be $\mathfrak{R}$-definable if the set $\{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid x + iy \in S\}$ in $\mathbb{R}^{2d}$ is $\mathfrak{R}$-definable.

A *discrete-time linear dynamical system* (LDS) consists of a pair (A, x), where $A \in \mathbb{Q}^{d \times d}$ and $x \in \mathbb{Q}^d$. Its *orbit* $\mathcal{O}$ is the set $\{A^n x \mid x \in \mathbb{N}\}$. An *invariant* for (A, x) is a set $\mathcal{I} \subseteq \mathbb{R}^d$ that contains x and is stable under applications of A, i.e., $A\mathcal{I} \subseteq \mathcal{I}$. Given a set $F \subseteq \mathbb{R}^d$, we say that the invariant $\mathcal{I}$ *avoids* F if the two sets are disjoint.

3 From the Orbit to Trajectory Cones and Rays

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

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

Here $\rho_1, \ldots, \rho_k \in \mathbb{R}_{\ge 0}$, $\lambda_1, \ldots, \lambda_k \in \mathbb{A}$ are of modulus 1, and $\sum_{i=1}^{k} d_i = d$. To reflect the block structure of J, we often range over $\{1, \ldots, d\}$ via a pair (i, j), with $1 \le i \le k$ and $1 \le j \le d_i$, which denotes the index corresponding to row j in block i; we refer to this notation as *block-row indexing*.

▶ Remark. Henceforth, we assume that for all $1 \le i \le k$ we have that $\rho_i > 0$ (i.e., that the matrices A and J are invertible). Indeed, if $\rho_i = 0$, 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 $\mathbb{R}^{d-d_i}$.

Observe that now, for every set $F \subseteq \mathbb{R}^d$, we have that $A^n x \in F$ iff $J^n x' \in P^{-1}F$ where $x' = P^{-1}x$.

For every $n > d$, $J^n = \mathrm{diag}(B_1^n, \ldots, B_k^n)$ with

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

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

Let $R = \mathrm{diag}(\rho_1, \ldots, \rho_k)$ and $L = \mathrm{diag}(\lambda_1, \ldots, \lambda_k)$. We define $\mathbb{T}$ to be the subgroup of the torus in $\mathbb{C}^k$ generated by the multiplicative relations of the normalised eigenvalues $\lambda_1, \ldots, \lambda_k$. That is, consider the subgroup $G = \left\{v = (v_1, \ldots, v_k) \in \mathbb{Z}^k \mid \lambda_1^{v_1} \cdots \lambda_k^{v_k} = 1\right\}$ of $\mathbb{Z}^k$, and let

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

Using Kronecker's theorem on inhomogeneous simultaneous Diophantine approximation [5] it is shown in [23] that $\{L^n \mid n \in \mathbb{N}\}$ is a dense subset of $\{\mathrm{diag}(\alpha_1, \ldots, \alpha_k) \mid (\alpha_1, \ldots, \alpha_k) \in \mathbb{T}\}$.

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

$$J^n x' \in \left\{ \begin{pmatrix} \rho_1^n p_1 Q_{1,1}(n) \\ \vdots \\ \rho_k^n p_k Q_{k,d_k}(n) \end{pmatrix} \ : \ (p_1, \ldots, p_k) \in \mathbb{T} \right\}.$$

We now define a continuous over-approximation of the expressions ρ_i^n. To this end, if there exists some modulus ρ_i larger than 1 (in which case, without loss of generality, assume

that $\rho_k > 1$), then for every $1 \leq i \leq k$ let $b_i = \log_{\rho_k} \rho_i$, and observe that $\rho_i^n = (\rho_k^n)^{b_i}$. We then replace the expression ρ_k^n with a continuous variable t, so that ρ_i^n becomes t^{b_i}, and n is replaced by $\log_{\rho_k} t$. If all moduli are at most 1 and some are strictly smaller than 1 (in which case, without loss of generality, $\rho_k < 1$), then replace the expression ρ_k^n with $1/t$. Note that in both cases, t grows unboundedly large as n tends to infinity. In Appendix A.1 we handle the special (and simpler) case in which all eigenvalues have modulus exactly 1. Henceforth, we assume that $\rho_k > 1$. If $\rho_k < 1$ the proofs are carried out *mutatis mutandis*.

This over-approximation leads to the following definition, which is central to our approach.

▶ **Definition 2.** For every $t_0 \geq 1$, we define the *trajectory cone*[1] for the orbit $\mathcal{O}$ as

$$\mathcal{C}_{t_0} = \left\{ \begin{pmatrix} t^{b_1} p_1 Q_{1,1}(\log_{\rho_k} t) \\ \vdots \\ t^{b_k} p_k Q_{k,d_k}(\log_{\rho_k} t) \end{pmatrix} \;:\; (p_1, \ldots, p_d) \in \mathbb{T},\ t \geq t_0 \right\}.$$

In particular, we have that $J^n x' \in \mathcal{C}_1$.

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.** For every $p = (p_1, \ldots, p_k) \in \mathbb{T}$ and every $t_0 \geq 1$, we define the *ray*[2] $\mathsf{r}(p, t_0) = \left\{ \left(t^{b_1} p_1 Q_{1,1}(\log_{\rho_k} t), \ldots, t^{b_k} p_k Q_{k,d_k}(\log_{\rho_k} t)\right)^\top \mid t \geq t_0 \right\}$.

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

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, x) with $A = PJP^{-1}$, the image under P of every trajectory cone $\mathcal{C}_{t_0}$, augmented with finitely many points from $\mathcal{O}$, is an invariant. Moreover, we show that such invariants are $\mathfrak{R}_{\exp}$-definable, and hence o-minimal. Complementing this, we show in Section 5 that *any* o-minimal invariant must contain some trajectory cone.

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

▶ **Theorem 4.** *For every $t_0 \geq 1$, the set $P\mathcal{C}_{t_0} \cup \{A^n x \mid n < \log_{\rho_k} t_0\}$ is an $\mathfrak{R}_{\exp}$-definable invariant for the LDS (A, x).*

The intuition behind Theorem 4 is as follows. Clearly, the orbit $\mathcal{O}$ itself is always an invariant for (A, x). However, it is generally not definable in any o-minimal theory (in particular, since it has infinitely many connected components). In order to recover definability in $\mathfrak{R}_{\exp}$ while maintaining stability under A, the invariants constructed in Theorem 4 over-approximate the orbit by the image of the trajectory cone $\mathcal{C}_{t_0}$ 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.

---

[1] These 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 $\mathcal{C}_{t_0}$ 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.

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

The proof of Theorem 4 has several parts. First, recall that the trajectory cone itself, $\mathcal{C}_{t_0}$, is an over-approximation of the set $\{J^n P^{-1} x \mid n \in \mathbb{N}\}$. As such, clearly $\mathcal{C}_{t_0} \subseteq \mathbb{C}^d$. In comparison, the orbit can be written as $\mathcal{O} = \{PJ^n P^{-1} x \mid n \in \mathbb{N}\} \subseteq \mathbb{R}^d$. We prove in Appendix A.2 the following lemma, from which it follows that the entire set $P\mathcal{C}_{t_0}$ is also a subset of $\mathbb{R}^d$.

▶ **Lemma 5.** *For every $p \in \mathbb{T}$ and $t_0 \geq 1$, we have $P\mathsf{r}(p, t_0) \subseteq \mathbb{R}^d$.*

Let us simply remark here that by analysing the structure of the matrices involved in defining $P\mathsf{r}(p, t_0)$, and using the facts that the columns of P are generalised eigenvectors of A, and that conjugate pairs of eigenvalues correspond to conjugate pairs of generalised eigenvectors, it is not hard to see that the above product does indeed yield only real values. However, a formal proof of this involves fairly tedious calculations. We invoke instead an analytic argument in Appendix A.2.

In the second part of the proof of Theorem 4, we show that $P\mathcal{C}_{t_0}$ is stable under A. The key ingredient is the following lemma, which characterises the action of J on rays, and is proved in Section 4.1.

▶ **Lemma 6.** *For every $p = (p_1, \ldots, p_k) \in \mathbb{T}$ and $t_0 \geq 1$, we have $J\mathsf{r}(p, t_0) = \mathsf{r}(Lp, \rho_k t_0)$.*

The next lemma then lifts Lemma 6 to the entire trajectory cone.

▶ **Lemma 7.** *For every $t_0 \geq 1$, we have $J\mathcal{C}_{t_0} \subseteq \mathcal{C}_{t_0}$.*

Proof. Recall that $\mathcal{C}_{t_0} = \bigcup_{p \in \mathbb{T}} \mathsf{r}(p, t_0)$. By Lemma 6 we have that $J\mathcal{C}_{t_0} = \bigcup_{p \in \mathbb{T}} \mathsf{r}(Lp, \rho_k t_0)$. Since $\rho_k > 1$, it follows that $\rho_k t_0 \geq t_0$. In addition, $p \in \mathbb{T}$ iff $Lp \in \mathbb{T}$. Hence we have that $\mathsf{r}(Lp, \rho_k t_0) \subseteq \mathsf{r}(Lp, t_0)$, from which we conclude that $J\mathcal{C}_{t_0} \subseteq \bigcup_{p \in \mathbb{T}} \mathsf{r}(Lp, t_0) = \bigcup_{p \in \mathbb{T}} \mathsf{r}(p, t_0) = \mathcal{C}_{t_0}$. ◀

The proof of Theorem 4 combines all these ingredients together and is given in subsection 4.2.

4.1 Proof of Lemma 6

Let $y = \begin{pmatrix} t^{b_1} p_1 Q_{1,1}(\log_{\rho_k} t) \\ \vdots \\ t^{b_k} p_k Q_{k,d_k}(\log_{\rho_k} t) \end{pmatrix} \in \mathsf{r}(p, t_0)$. We claim that $Jy = \begin{pmatrix} (\rho_k t)^{b_1} \lambda_1 p_1 Q_{1,1}(\log_{\rho_k}(\rho_k t)) \\ \vdots \\ (\rho_k t)^{b_k} \lambda_k p_k Q_{k,d_k}(\log_{\rho_k}(\rho_k t)) \end{pmatrix}$.
Note that since $Lp = (\lambda_1 p_1, \ldots, \lambda_k p_k)$, 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 λ, ρ, and d instead of λ_i, ρ_i, and d_i, respectively. Then we have

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

Recall that

$$Q_{i,j}(\log_{\rho_k} t) = \sum_{c=0}^{d-j} \frac{\binom{\log_{\rho_k} t}{c}}{(\rho\lambda)^c} x'_{i,j+c}\,,$$

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

$$(Jy)_m = \lambda\rho t^{b_i} p_i \sum_{c=0}^{d-j} \frac{\binom{\log_{\rho_k} t}{c}}{(\rho\lambda)^c} x'_{i,j+c} + t^{b_i} p_i \sum_{c=0}^{d-j-1} \frac{\binom{\log_{\rho_k} t}{c}}{(\rho\lambda)^c} x'_{i,j+c+1}\,. \tag{2}$$

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

$$(\rho_k t)^{b_i}\lambda p_i Q_{i,j}(\log_{\rho_k}(\rho_k t)) = (\rho_k t)^{b_i}\lambda p_i \sum_{c=0}^{d-j} \frac{\binom{\log_{\rho_k}(\rho_k t)}{c}}{(\rho\lambda)^c} x'_{i,j+c}\,. \tag{3}$$

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

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

$$\lambda\rho t^{b_i} p_i \frac{\binom{\log_{\rho_k} t}{c}}{(\rho\lambda)^c} + t^{b_i} p_i \frac{\binom{\log_{\rho_k} t}{c-1}}{(\rho\lambda)^{c-1}} = \frac{t^{b_i}\rho\lambda p_i}{(\rho\lambda)^c}\left(\binom{\log_{\rho_k} t}{c} + \binom{\log_{\rho_k} t}{c-1}\right) = \frac{t^{b_i}\lambda\rho p_i}{\lambda^c}\binom{\log_{\rho_k} t + 1}{c}$$

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

4.2 Proof of Theorem 4

Let $t_0 \geq 1$. By applying Lemma 5 to every $p \in \mathbb{T}$, we conclude that $P\mathcal{C}_{t_0} \subseteq \mathbb{R}^d$. It is then easy to see that $P\mathcal{C}_{t_0}$ is definable in $\mathfrak{R}_{\exp}$ (note that the only reason the set $\mathcal{C}_{t_0}$ might fail to be $\mathfrak{R}_{\exp}$-definable is that the underlying domain should be $\mathbb{R}$ and not $\mathbb{C}$).

Next, by Lemma 7 we have that $J\mathcal{C}_{t_0} \subseteq \mathcal{C}_{t_0}$. Applying P from the left, we get $PJ\mathcal{C}_{t_0} \subseteq P\mathcal{C}_{t_0}$. Thus, we have $AP\mathcal{C}_{t_0} = PJP^{-1}P\mathcal{C}_{t_0} = PJ\mathcal{C}_{t_0} \subseteq P\mathcal{C}_{t_0}$.

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

5 O-Minimal Invariants Must Contain Trajectory Cones

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

▶ **Theorem 8.** *Consider an $\mathfrak{R}$-definable invariant $\mathcal{I}$ for the LDS (A, x). Then there exists $t_0 \geq 1$ such that $P\mathcal{C}_{t_0} \subseteq \mathcal{I}$.*

To prove Theorem 8, we begin by making following claims of increasing strength:

▶ Claim 1. For every $p \in \mathbb{T}$ there exists $t_0 \geq 1$ such that $Pr(p, t_0) \subseteq \mathcal{I}$ or $Pr(p, t_0) \cap \mathcal{I} = \emptyset$.

▶ Claim 2. For every $p \in \mathbb{T}$ there exists $t_0 \geq 1$ such that $Pr(p, t_0) \subseteq \mathcal{I}$.

▶ Claim 3. There exists $t_0 \geq 1$ such that for every $p \in \mathbb{T}$ we have $Pr(p, t_0) \subseteq \mathcal{I}$.

Proof of Claim 1: Fix $p \in \mathbb{T}$. Observe that by Lemma 5, $Pr(p, 1)$ is $\mathfrak{R}$-definable. Further note that $Pr(p, 1)$ is of dimension 1 (as it is homeomorphic to $[1, \infty)$). Thus, the dimension of $Pr(p, 1) \cap \mathcal{I}$ is at most 1, so its cell decomposition contains finitely many connected components of dimensions 0 or 1. In particular, either one component is unbounded, in which case there exists a t_0 such that $Pr(p, t_0) \subseteq \mathcal{I}$, or all the components are bounded, in which case there exists a t_0 such that $Pr(p, t_0) \cap \mathcal{I} = \emptyset$. ◀

Before proceeding to Claim 2, we prove an auxiliary lemma, which is an adaptation of a similar result from [13]. 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 9.** *Let $S, F \subseteq \mathbb{T}$ be $\mathfrak{R}$-definable[3] sets such that $\overline{S} = \overline{F} = \mathbb{T}$. Then $F \cap S \neq \emptyset$.*

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 $\dim(X) = \dim(\overline{X})$ [12, 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 non-empty interior in Y [12, Chapter 4, Corollary 1.9]. In the situation at hand, since $\dim(F) = \dim(\overline{F})$, it follows that F has 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 non-empty interior in $\overline{S}$, and thus $S \cap F \neq \emptyset$. ◀

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 $Pr(p, t_0) \cap \mathcal{I} = \emptyset$, and consider $J^{-1}\mathsf{r}(p, t_0)$. Let $q \in \mathbb{T}$ be $L^{-1}p = (\frac{p_1}{\lambda_1}, \ldots, \frac{p_k}{\lambda_k})$ and let $t_1 = \frac{t_0}{\rho_k}$. Then $p = Lq$ and $t_0 = \rho_k t_1$ and, by Lemma 6, $J\mathsf{r}(q, t_1) = \mathsf{r}(Lq, \rho_k t_1) = \mathsf{r}(p, t_0)$. Since J is invertible, we conclude that $J^{-1}\mathsf{r}(p, t_0) = \mathsf{r}(q, t_1)$.

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

Recall that, following the discussion in section 3, we have $\rho_k > 1$. This implies $t_1 \leq t_0$ and $\mathsf{r}(q, t_0) \subseteq \mathsf{r}(q, t_1)$, so in particular $Pr(q, t_0) \cap \mathcal{I} = \emptyset$. Thus, assuming $Pr(p, t_0) \cap \mathcal{I} = \emptyset$, we have just proved that $Pr(L^{-1}p, t_0) \cap \mathcal{I} = \emptyset$; repeating this argument, we get that for every $n \in \mathbb{N}$, the point $s = L^{-n}p$ satisfies $Pr(s, t_0) \cap \mathcal{I} = \emptyset$.

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 $S' = \{s \in \mathbb{T} \mid Pr(s, t_0) \cap \mathcal{I} = \emptyset\}$. 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 $Pr(q, t_0) \cap \mathcal{I} \neq \emptyset$. It follows that for every $n \in \mathbb{N}$, the point $q' = L^n q$ also satisfies $Pr(q', t_0) \cap \mathcal{I} \neq \emptyset$. Define $F = \{L^n q \mid n \in \mathbb{N}\}$, then F is dense in $\mathbb{T}$. But then the set $F' = \{q \in \mathbb{T} \mid Pr(q, t_0) \cap \mathcal{I} \neq \emptyset\}$ 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 9 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 $P\mathcal{C}_{t_0} \cap \mathcal{I} = \emptyset$, which is again a contradiction, since $P\mathcal{C}_{t_0} \cap \mathcal{O} \neq \emptyset$ and $\mathcal{O} \subseteq \mathcal{I}$, so we are done. ◀

Proof of Claim 3: Consider the function $f\colon \mathbb{T} \to \mathbb{R}$ defined by $f(p) = \inf\{t \in \mathbb{R} \mid Pr(p, t) \subseteq \mathcal{I}\}$. By Claim 2 this function is well-defined. Since $Pr(p, t)$ 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[4]. 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

[3] Recall that in order to reason about $\mathbb{T} \subseteq \mathbb{C}^k$ in $\mathfrak{R}$ we identify $\mathbb{C}$ with $\mathbb{R}^2$.
[4] In case $m = 0$, the proof actually follows immediately from Claim 2, since $\mathbb{T}$ is finite.

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 $\{L^{n_1}K'', \ldots, L^{n_m}K''\}$. Since $K'' \subseteq K'$, it follows that $\{L^{n_1}K', \ldots, L^{n_m}K'\}$ is a finite closed cover of $\mathbb{T}$.

We now show that, for all $p \in \mathbb{T}$, we have $f(Lp) \leq \rho_k \cdot f(p)$. Indeed, consider any $p \in \mathbb{T}$ and $t > 0$ such that $Pr(p,t) \subseteq \mathcal{I}$. Applying $A = PJP^{-1}$, we get $PJr(p,t) \subseteq A\mathcal{I} \subseteq \mathcal{I}$. By Lemma 6, $Jr(p,t) = r(Lp, \rho_k t)$, so we can conclude that $Pr(Lp, \rho_k t) \subseteq \mathcal{I}$. This means that $Pr(p,t) \subseteq \mathcal{I}$ implies $Pr(Lp, \rho_k t) \subseteq \mathcal{I}$; therefore, $f(Lp) \leq \rho_k \cdot f(p)$.

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

Finally, we conclude from Claim 3 that there exists $t_0 \geq 1$ such that $P\mathcal{C}_{t_0} \subseteq \mathcal{I}$. This completes the proof of Theorem 8.

6 Deciding the Existence of O-Minimal Invariants

We now turn to the algorithmic aspects of invariants and present our two main results, Theorems 11 and 12.

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

The following is a consequence of Theorems 4 and 8.

▶ **Lemma 10.** *Let (A, x) and $\mathfrak{R}$ be as above, and let F be $\mathfrak{R}$-definable. Then there exists an o-minimal invariant $\mathcal{I}$ for (A, x) that avoids F iff there is some $t_0 \geq 1$ such that $P\mathcal{C}_{t_0} \cap F = \emptyset$ and such that $A^n x \notin F$ for every $0 \leq n \leq \log_{\rho_k} t_0$.*

Proof. By Theorem 8, if an o-minimal invariant $\mathcal{I}$ for (A, x) exists, then there exists $t_0 \geq 1$ such that $P\mathcal{C}_{t_0} \subseteq \mathcal{I}$. Moreover, $\mathcal{I} \cap F = \emptyset$ implies $\mathcal{O} \cap F = \emptyset$, so that $A^n x \notin F$ for every $n \in \mathbb{N}$, and in particular for $0 \leq n \leq \log_{\rho_k} t_0$.

Conversely, let there be $t_0 \geq 1$ such that $P\mathcal{C}_{t_0} \cap F = \emptyset$ and, for every $0 \leq n \leq \log_{\rho_d} t_0$, it holds that $A^n x \notin F$. Let $t_0' \in \mathbb{Q}$ be such that $t_0' \geq t_0$. By Theorem 4, the set $\mathcal{I} = P\mathcal{C}_{t_0'} \cup \left\{A^n x \mid 0 \leq n \leq \log_{\rho_d} t_0'\right\}$ is an $\mathfrak{R}_{\exp}$-definable invariant that avoids F. ◀

Observe that the formula $\exists t_0 \geq 1 : P\mathcal{C}_{t_0} \cap F = \emptyset$ is a sentence in $\mathfrak{R}_{\exp}$, and by Lemma 10, deciding the existence of an invariant amounts to determining the truth value of this sentence.

Decidability for $\mathfrak{R}_{\exp}$-definable halting sets assuming Schanuel's conjecture. Applying Theorem 4, we note that an invariant for (A, x) that avoids F—if one exists—can always be defined in $\mathfrak{R}_{\exp}$.

▶ **Theorem 11.** *The O-Minimal Invariant Synthesis Problem for $\mathfrak{R}_{\exp}$-Definable Halting Sets is decidable, assuming Schanuel's conjecture.*

Proof. Assume Schanuel's conjecture. Then by [18], we have that $\mathfrak{R}_{\exp}$ is decidable. Thus we can decide whether there exists $t_0 \geq 1$ such that $P\mathcal{C}_{t_0} \cap F = \emptyset$. If the sentence is false, then by Lemma 10 there is no invariant, and we are done. If the sentence is true, however, it still remains to check whether $A^n x \notin F$ for every $0 \leq n \leq \log_{\rho_k} t_0$. While we can decide whether $A^n x \in F$ for a fixed n, observe that we do not have an *a priori* bound on t_0. Hence we proceed as follows: For every $n \in 1, 2, \ldots$, check both whether $A^n x \in F$ and, for $t_0 = \rho_k^n$, whether $PC_{t_0} \cap F = \emptyset$. In case $A^n x \in F$, then clearly there is no invariant, since $\mathcal{O} \cap F \neq \emptyset$,

and we are done. On the other hand, if $PC_{t_0} \cap F = \emptyset$, then return the semi-algebraic invariant as per Lemma 10.

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

▶ Remark. 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 worse that could happen is that decidability of $\mathfrak{R}_{\exp}$ fails in that the putative algorithm of [18] simply never halts, so no verdict is ever returned.

Unconditional decidability for semi-algebraic halting sets.

▶ **Theorem 12.** *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* $\mathfrak{R}_{\exp}$.

By Lemma 10, in order to prove Theorem 12 it is enough to decide the truth value of the $\mathfrak{R}_{\exp}$ sentence $\exists t_0 \geq 1 : PC_{t_0} \cap F = \emptyset$. Indeed, as $A^n x \in \mathbb{Q}^d$, one can always check unconditionally whether for a given n the vector $A^n x$ belongs to the semi-algebraic set F. The algorithm is then otherwise the same as that presented in the proof of Theorem 11. The proof of Theorem 12 therefore boils down to the following lemma.

▶ **Lemma 13.** *For F a semi-algebraic set, it is decidable whether there exists $t_0 \geq 1$ such that $PC_{t_0} \cap F = \emptyset$.*

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

▶ **Theorem 14** (Baker's theorem [1]). *Let $\alpha_1, \ldots, \alpha_m \in \mathbb{C}$ be algebraic numbers different from 0 and let $b_1, \ldots, b_m \in \mathbb{Z}$ be integers. Write $\Lambda = b_1 \log \alpha_1 + \ldots + b_m \log \alpha_m$. There exists a number C effectively computable from $b_1, \ldots, b_m, \alpha_1, \ldots, \alpha_m$ such that if $\Lambda \neq 0$ then $|\Lambda| > H^{-C}$, where H is the maximum height of $\alpha_1, \ldots, \alpha_m$.*

As in Section 3, we assume that $\rho_d > 1$ (with the cases of $\rho_d < 1$ and $\rho_d = 1$ being analogous and easier, respectively). 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 $\vec{\tau}(t) = (t^{b_1} Q_{1,1}(\log_{\rho_k} t), \ldots, t^{b_k} Q_{k,d_k}(\log_{\rho_k} t))$, and consider the set

$$U = \left\{ \begin{pmatrix} z_1 \\ \vdots \\ z_d \end{pmatrix} \in \mathbb{C}^d : \forall (p_1, \ldots, p_d) \in \mathbb{T},\ P \begin{pmatrix} z_1 p_1 \\ \vdots \\ z_d p_d \end{pmatrix} \in \mathbb{R}^d \setminus F \right\}.$$

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

Observe that U is a semi-algebraic set. (Recall that we identify $\mathbb{C}$ with $\mathbb{R}^2$; see Remark on page 4 in the Preliminaries.) Using cell decomposition, describe U as a finite union of connected components, each of which is given by a conjunction of the form $\bigwedge_{l=1}^{m} R_l(u_1, \ldots, u_d, v_1, \ldots, v_d) \sim_l 0$. Here, for every $1 \leq l \leq m$, $\sim_l \in \{>, =\}$ and R_l is a polynomial with integer coefficients in variables $u_1, \ldots, u_d, v_1, \ldots, v_d$; for each i, the variables u_i and v_i represent $\operatorname{Re} z_i$ and $\operatorname{Im} z_i$, the real and imaginary parts of z_i, respectively.

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

get a constraint on t expressible in $\mathfrak{R}_{\exp}$. By o-minimality of $\mathfrak{R}_{\exp}$, 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 $\vec{\tau}(t) \in U$, 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 $R_l(\vec{\tau}(t)) \sim_l 0$ 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

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

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

$$a_i \cdot t^{(n'_{i,1}+n''_{i,1})\cdot b_1+\ldots+(n'_{i,d}+n''_{i,d})\cdot b_k} \cdot (\operatorname{Re} Q_{1,1}(\log_{\rho_k} t))^{n'_{i,1}} \cdot \ldots \cdot (\operatorname{Re} Q_{k,d_k}(\log_{\rho_k} t))^{n'_{i,k}} \cdot (\operatorname{Im} Q_{1,1}(\log_{\rho_k} t))^{n''_{i,1}} \cdot \ldots \cdot (\operatorname{Im} Q_{k,d_k}(\log_{\rho_k} t))^{n''_{i,k}}$$

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 $\operatorname{Re} Q_{i,j}$ and $\operatorname{Im} Q_{i,j}$; these are polynomials in one real variable with real algebraic coefficients. We rewrite $R_l(\vec{\tau}(t))$ in the form

$$\sum_i t^{n_{i,1}\cdot b_1+\ldots+n_{i,k}\cdot b_k} \cdot f_i(\log_{\rho_k} t)$$

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

Observe that $R_l(\vec{\tau}(t))$ is a function of a single variable $t > 0$. In order to reason about the sign of this expression as $t \to \infty$, 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 $n_{i,1}b_1 + \ldots + n_{i,k}b_k$ and $n_{j,1}b_1 + \ldots + n_{j,k}b_k$ 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's enough to decide whether $n_1b_1 + \ldots + n_kb_k > 0$ for some $n_1, \ldots, n_k \in \mathbb{Z}$. Recall that $b_j = \log_{\rho_k} \rho_j = \log \rho_j / \log \rho_k$ where log denotes the natural logarithm. By Baker's theorem, there exists an effectively computable $\epsilon > 0$ such that either $n_1b_1 + \ldots + n_kb_k = 0$, or $|n_1b_1 + \ldots + n_kb_k| > \epsilon$.

We now proceed by computing an approximation Δ of $n_1b_1 + \ldots + n_kb_k$ 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 $n_1b_1 + \ldots + n_kb_k = 0$, and otherwise we have $\operatorname{sign}(\Delta) = \operatorname{sign}(n_1b_1 + \ldots + n_kb_k)$. Thus we can sort the exponents $n_{i,1} \cdot b_1 + \ldots + n_{i,k} \cdot b_k$ 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 $t^m \cdot f_i(\log_{\rho_k} t)$. Then the sign of $R_l(\vec{\tau}(t))$ as $t \to \infty$ 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 $R_l(\vec{\tau}(t))$ be 0 for all sufficiently large t.

The argument above shows that we can compute the leading terms of the expressions $R_l(\vec{\tau}(t))$ 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.

References

1 Alan Baker and Gisbert Wüstholz. Logarithmic forms and group varieties. *J. reine angew. Math*, 442(19-62):3, 1993.

2 Amir M. Ben-Amram and Samir Genaim. Ranking functions for linear-constraint loops. *J. ACM*, 61(4):26:1–26:55, 2014.

3 Amir M. Ben-Amram, Samir Genaim, and Abu Naser Masud. On the termination of integer loops. *ACM Trans. Program. Lang. Syst.*, 34(4):16:1–16:24, 2012.

4 Mark Braverman. Termination of integer linear programs. In *Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings*, pages 372–385, 2006.

5 John W.S. Cassels. *An Introduction to Diophantine Approximation*. Cambridge University Press, 1965.

6 Ventsislav Chonev, Joël Ouaknine, and James Worrell. The Orbit Problem in higher dimensions. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 941–950, 2013.

7 Ventsislav Chonev, Joël Ouaknine, and James Worrell. The polyhedron-hitting problem. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 940–956, 2015.

8 Ventsislav Chonev, Joël Ouaknine, and James Worrell. On the complexity of the orbit problem. *J. ACM*, 63(3):23:1–23:18, 2016.

9 Michael Colón, Sriram Sankaranarayanan, and Henny Sipma. Linear invariant generation using non-linear constraint solving. In *Computer Aided Verification, 15th International Conference, CAV 2003, Boulder, CO, USA, July 8-12, 2003, Proceedings*, pages 420–432, 2003.

10 Patrick Cousot. Proving program invariance and termination by parametric abstraction, lagrangian relaxation and semidefinite programming. In *Verification, Model Checking, and Abstract Interpretation, 6th International Conference, VMCAI 2005, Paris, France, January 17-19, 2005, Proceedings*, pages 1–24, 2005.

11 Patrick Cousot and Nicolas Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, Tucson, Arizona, USA, January 1978*, pages 84–96, 1978.

12 L. P. D. van den Dries. *Tame Topology and O-minimal Structures*. London Mathematical Society Lecture Note Series. Cambridge University Press, 1998.

13 Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. Semialgebraic invariant synthesis for the kannan-lipton orbit problem. In *34th Symposium on Theoretical Aspects of Computer Science, STACS 2017, March 8-11, 2017, Hannover, Germany*, pages 29:1–29:13, 2017.

14 Ashutosh Gupta, Thomas A. Henzinger, Rupak Majumdar, Andrey Rybalchenko, and Ru-Gang Xu. Proving non-termination. In *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2008, San Francisco, California, USA, January 7-12, 2008*, pages 147–158, 2008.

15 Ravindran Kannan and Richard J. Lipton. The orbit problem is decidable. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing, April 28-30, 1980, Los Angeles, California, USA*, pages 252–261, 1980.

16 Ravindran Kannan and Richard J. Lipton. Polynomial-time algorithm for the orbit problem. *J. ACM*, 33(4):808–821, 1986.

17 Zachary Kincaid, John Cyphert, Jason Breck, and Thomas W. Reps. Non-linear reasoning for invariant synthesis. *PACMPL*, 2(POPL):54:1–54:33, 2018.

18 Angus Macintyre and Alex J. Wilkie. On the decidability of the real exponential field. In Piergiorgio Odifreddi, editor, *Kreiseliana. About and Around Georg Kreisel*, pages 441–467. A K Peters, 1996.

19 M. Mignotte, T. Shorey, and R. Tijdeman. The distance between terms of an algebraic recurrence sequence. *J. für die reine und angewandte Math.*, 349, 1984.

20 Joël Ouaknine, João Sousa Pinto, and James Worrell. On termination of integer linear loops. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 957–969, 2015.

21 Joël Ouaknine and James Worrell. On the positivity problem for simple linear recurrence sequences,. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*, pages 318–329, 2014.

22 Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 366–379, 2014.

23 Joël Ouaknine and James Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In *International Colloquium on Automata, Languages, and Programming*, pages 330–341. Springer, 2014.

24 Joël Ouaknine and James Worrell. On linear recurrence sequences and loop termination. *SIGLOG News*, 2(2):4–13, 2015.

25 Enric Rodríguez-Carbonell and Deepak Kapur. An abstract interpretation approach for automatic generation of polynomial invariants. In *Static Analysis, 11th International Symposium, SAS 2004, Verona, Italy, August 26-28, 2004, Proceedings*, pages 280–295, 2004.

26 Enric Rodríguez-Carbonell and Deepak Kapur. Generating all polynomial invariants in simple loops. *J. Symb. Comput.*, 42(4):443–476, 2007.

27 Sriram Sankaranarayanan, Henny Sipma, and Zohar Manna. Non-linear loop invariant generation using gröbner bases. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2004, Venice, Italy, January 14-16, 2004*, pages 318–329, 2004.

28 T. Tao. *Structure and randomness: pages from year one of a mathematical blog*. American Mathematical Society, 2008.

29 Alfred Tarski. A decision method for elementary algebra and geometry. *Bulletin of the American Mathematical Society*, 59, 1951.

30 Ashish Tiwari. Termination of linear programs. In *Computer Aided Verification, 16th International Conference, CAV 2004, Boston, MA, USA, July 13-17, 2004, Proceedings*, pages 70–82, 2004.

31 N. K. Vereshchagin. The problem of appearance of a zero in a linear recurrence sequence (in Russian). *Mat. Zametki*, 38(2), 1985.

32 A. J. Wilkie. 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):1051–1094, 1996.

33 Bican Xia and Zhihai Zhang. Termination of linear programs with nonlinear constraints. *J. Symb. Comput.*, 45(11):1234–1249, 2010.

A Proofs

A.1 All Eigenvalues Have Modulus 1

In this section we adapt the proof of Section 3 to the situation in which all eigenvalues have modulus 1. Observe that in this case, every coordinate of $J^n x'$ is of the form $\lambda_i^n Q_{i,j}(n)$ with

$|\lambda_i| = 1$. Then, to define the trajectory cones, we over-approximate n by t, and thus we define $\mathcal{C}_{t_0} = \left\{ \begin{pmatrix} p_1 Q_{1,1}(t) \\ \vdots \\ p_k Q_{k,d_k}(t) \end{pmatrix} \;:\; (p_1, \ldots, p_d) \in \mathbb{T},\ t \geq t_0 \right\}$.

The rest of the proof can be followed *mutatis mutandis* and is in fact much simpler, as the entries are now polynomials in t, rather than exponential polynomials.

A.2 Proof of Lemma 5

Let $p \in \mathbb{T}$ and $t_0 \geq 1$. Consider $\begin{pmatrix} t^{b_1} p_1 Q_{1,1}(\log_{\rho_k} t) \\ \vdots \\ t^{b_k} p_k Q_{k,d_k}(\log_{\rho_k} t) \end{pmatrix} \in \mathsf{r}(p, t_0)$ and observe that for every $1 \leq i \leq k$ and $1 \leq j \leq d_i$ we have

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

with $(i, j+m)$ being a block-row index (see Section 3).[5]

For $t \geq t_0$, let $M(t)$ denote image of the above vector under P. Then

$$M(t) = P \begin{pmatrix} t^{b_1} p_1 Q_{1,1}(\log_{\rho_k} t) \\ \vdots \\ t^{b_k} p_k Q_{k,d_k}(\log_{\rho_k} t) \end{pmatrix} = P \begin{pmatrix} t^{b_1} p_1 C_1 & & \\ & \ddots & \\ & & t^{b_k} p_k C_k \end{pmatrix} P^{-1} x$$

where for $1 \leq i \leq k$, $C_i \in \mathbb{C}^{d_i \times d_i}$ is the matrix $C_i = \begin{pmatrix} 1 & \frac{\log_{\rho_k} t}{\rho_i \lambda_i} & \cdots & \frac{\binom{\log_{\rho_k} t}{d_i - 1}}{(\rho_i \lambda_i)^{d_i - 1}} \\ & 1 & \cdots & \frac{\binom{\log_{\rho_k} t}{d_i - 2}}{(\rho_i \lambda_i)^{d_i - 2}} \\ & & \ddots & \vdots \\ & & & 1 \end{pmatrix}$. Recall that $A = PJP^{-1}$ where $A \in \mathbb{R}^{d \times d}$, and observe that for $t = \rho_k^n$ with $n \in \mathbb{N}$ and $n > d$, we have that $M(t) = M(\rho_k^n) = PJ^n P^{-1} x = A^n x \in \mathbb{R}^d$.

For every $1 \leq i \leq d$, consider coordinate $M(t)_i$ as a function of t. That is, $M(t)_i = \sum_{1 \leq j \leq k} \sum_{1 \leq j' \leq d_i} P_{(j,j'),i} t^{b_j} p_j Q_{j,j'}(\log_{\rho_k} t)$ (where (j, j') is a block-row index).

A priori, we have $M(\cdot)_i \colon \mathbb{R} \to \mathbb{C}$. However, by the above we see that its imaginary part $\mathrm{Im}(M(t)_i) = 0$ for infinitely many values of t, namely ρ_k^n for every $n > d$. We now show that this implies $\mathrm{Im}(M(t)_i)$ is identically 0.

Consider $|\mathrm{Im}(M(t)_i)|$. Observe that each term in the sum above is of the form $t^b \mathrm{poly}(\log t)$ for some $b \in \mathbb{R}$. Thus, $|\mathrm{Im}(M(t)_i)|$ can be bounded from below by a function of this form, and in particular its roots are bounded, unless it is identically 0. Indeed, clearly if $b_j > 0$ for some j then this function tends to ∞, but even if $b_j \leq 0$ for every $1 \leq j \leq k$, we can bound this from below by an eventually-positive (possibly decreasing) function. However, since the roots are not bounded, we conclude that $|\mathrm{Im}(M(t)_i)|$ is identically 0. It follows that $M(t) \in \mathbb{R}^d$ for every $t \in \mathbb{R}$, so $Pr(p, t_0) \subseteq \mathbb{R}^d$.

[5] Here, 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.