Semialgebraic Invariant Synthesis for the Kannan-Lipton Orbit Problem

Nathanaël Fijalkow[1], Pierre Ohlmann[2], Joël Ouaknine[3], Amaury Pouly[4], and James Worrell[5]

1 **Department of Computer Science, Oxford University, Oxford, UK**
2 **École Normale Supérieure de Lyon, Lyon, France**
3 **Department of Computer Science, Oxford University, Oxford, UK; and Max Planck Institute for Software Systems (MPI-SWS), Saarland Informatics Campus, Saarbrücken, Germany**
4 **Department of Computer Science, Oxford University, Oxford, UK**
5 **Department of Computer Science, Oxford University, Oxford, UK**

Abstract

The *Orbit Problem* consists of determining, given a linear transformation A on $\mathbb{Q}^d$, together with vectors x and y, whether the orbit of x under repeated applications of A can ever reach y. This problem was famously shown to be decidable by Kannan and Lipton in the 1980s.

In this paper, we are concerned with the problem of synthesising suitable *invariants* $\mathcal{P} \subseteq \mathbb{R}^d$, *i.e.*, sets that are stable under A and contain x and not y, thereby providing compact and versatile certificates of non-reachability. We show that whether a given instance of the Orbit Problem admits a semialgebraic invariant is decidable, and moreover in positive instances we provide an algorithm to synthesise suitable invariants of polynomial size.

It is worth noting that the existence of *semilinear* invariants, on the other hand, is (to the best of our knowledge) not known to be decidable.

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

Keywords and phrases Verification, algebraic computation, Skolem Problem, Orbit Problem, invariants

Digital Object Identifier 10.4230/LIPIcs.STACS.2017.29

1 Introduction

The *Orbit Problem* was introduced by Kannan and Lipton in the seminal papers [8, 9], and shown there to be decidable in polynomial time, answering in the process a decade-old open problem of Harrison on accessibility for linear sequential machines [7]. The Orbit Problem can be stated as follows:

> Given a square matrix $A \in \mathbb{Q}^{d\times d}$ together with vectors $x, y \in \mathbb{Q}^d$, decide whether there exists a non-negative integer n such that $A^n x = y$.

In other words, if one considers the discrete 'orbit' of the vector x under repeated applications of the linear transformation A, does the orbit ever hit the target y? Although it is not *a priori* obvious that this problem is even decidable, Kannan and Lipton showed that it can in fact be solved in polynomial time, by making use of spectral techniques as well as some sophisticated results from algebraic number theory.

In instances of non-reachability, a natural and interesting question is whether one can produce a suitable *invariant* as certificate, *i.e.*, a set $\mathcal{P} \subseteq \mathbb{R}^d$ that is stable under A (in the

34th Symposium on Theoretical Aspects of Computer Science (STACS 2017).
Editors: Heribert Vollmer and Brigitte Vallée; Article No. 29; pp. 29:1–29:13
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

sense that $A\mathcal{P} \subseteq \mathcal{P}$) and such that $x \in \mathcal{P}$ and $y \notin \mathcal{P}$. The existence of such an invariant then immediately entails by induction that the orbit of x does indeed avoid y.

Invariants appear in a wide range of contexts, from gauge theory, dynamical systems, and control theory in physics, mathematics, and engineering to program verification, static analysis, abstract interpretation, and programming language semantics (among others) in computer science. Automated invariant synthesis is a topic of active current research, particularly in the fields of theorem proving and program verification; in the latter, for example, one might imagine that y corresponds to a faulty or undesirable program state, and an invariant $\mathcal{P}$ as described above amounts to a succinct 'safety' certificate (here the program or procedure in question corresponds to a simple WHILE loop with linear updates).

The widespread use of invariants should not come as a surprise. In addition to their obvious advantage in constituting easily understandable safety certificates, their inductive nature makes them ideally suited to modular reasoning, often allowing one to analyse complex systems by breaking them down into simpler parts, each of which can then be handled in isolation. Invariants, viewed as safety certificates, also enable one to reason over large sets of program states rather than individual instances: in the context of the Orbit Problem, for example, an invariant $\mathcal{P} \subseteq \mathbb{R}^d$ such that $x \in \mathcal{P}$ and $y \notin \mathcal{P}$ doesn't merely certify that y is not reachable from x, but in fact guarantees that from *any* starting point $x' \in \mathcal{P}$, it is impossible to reach *any* of the points $y' \in \mathbb{R}^d \setminus \mathcal{P}$.

In general, when searching for invariants, one almost always fixes ahead of time a class of suitable potential candidates. Indeed, absent such a restriction, one would point out that the orbit $\mathcal{O}(x) = \{A^n x : n \geq 0\}$ is always by definition stable under A, and in instances of non-reachability will therefore always constitute a safety invariant. Such an invariant will however often not be of much use, as it will usually lack good algorithmic properties; for example, as observed in [9], in dimension $d = 5$ and higher, the question of whether the orbit $\mathcal{O}(x)$ reaches a given $(d-1)$-dimensional hyperplane corresponds precisely to the famous *Skolem Problem* (of whether an order-d linear recurrence sequence over the integers has a zero), whose decidability has been open for over 80 years [12].

Thus let us assume that we are given a domain $\mathbf{D} \subseteq 2^{\mathbb{R}^d}$ of suitable potential invariants. At a minimum, one would require that the relevant stability and safety conditions (*i.e.*, for any $\mathcal{P} \in \mathbf{D}$, whether $A\mathcal{P} \subseteq \mathcal{P}$, $x \in \mathcal{P}$, and $y \notin \mathcal{P}$) be algorithmically checkable (with reasonable complexity). The following natural questions then arise:

1. In instances of non-reachability, does a suitable invariant in $\mathbf{D}$ *always* exist?
2. If not, can we characterise the exceptional instances in some way?
3. In instances of non-reachability, can we algorithmically determine whether a suitable invariant in $\mathbf{D}$ exists, and when this is the case can we moreover synthesise such an invariant?

1. and **3.** are usually referred to as *completeness* and *relative completeness* respectively, whereas **2.** attempts to measure the extent to which completeness fails.

Main results. The main results of this paper concern the synthesis of semialgebraic invariants for non-reachability instances of the Kannan-Lipton Orbit Problem, where the input is provided as a triple (A, x, y) with all entries rational, and can be summarised as follows:

- We prove that whether a suitable semialgebraic[1] invariant exists or not is decidable in

[1] A semialgebraic set is the set of solutions of a Boolean combination of polynomial inequalities, with the polynomials in question having integer coefficients.

polynomial space, and moreover in positive instances we show how to synthesise a suitable invariant of polynomial size in polynomial space.

- We provide a simple characterisation of instances of non-reachability for which there does not exist a suitable semialgebraic invariant, and show that such instances are very 'rare', in a measure-theoretic sense.

Since the existence of suitable semialgebraic invariants for the Orbit Problem does not coincide precisely with non-reachability, our proof necessarily departs substantially from that given by Kannan and Lipton in [8, 9]. In particular, handling negative instances relies upon certain topological and geometrical insights into the structure of semialgebraic sets, and positive instances require the explicit construction of suitable semialgebraic invariants of polynomial size. We achieve this by making use of techniques from algebraic number theory such as Kronecker's Theorem on inhomogeneous simultaneous Diophantine approximation, and Masser's deep results on multiplicative relations among algebraic numbers.

The following three examples illustrate a range of phenomena that arise in searching for semialgebraic invariants.

▶ **Example 1.** Consider the matrix

$$A = \tfrac{1}{5}\begin{pmatrix} 4 & -3 \\ 3 & 4 \end{pmatrix} .$$

Matrix A defines a counterclockwise rotation around the origin by angle $\arctan(3/5)$, which is an irrational multiple of π. Thus the topological closure of the orbit $\mathcal{O} = \{x, Ax, A^2x, \ldots\}$ is a circle in $\mathbb{R}^2$. If $y \notin \overline{\mathcal{O}}$ then $\overline{\mathcal{O}}$ itself is clearly a suitable semialgebraic invariant. On other hand, it can be shown that if $y \in \overline{\mathcal{O}} \setminus \mathcal{O}$ then there does not exist a suitable semialgebraic invariant. (In passing, it is also not difficult to see that the only polygons $\mathcal{P}$ that are invariant under A are $\emptyset$, $\{(0,0)\}$, and $\mathbb{R}^2$.) More general orthogonal matrices can be handled along similar lines to the present case, but the analysis is substantially more involved. In general, the only cases in which $y \notin \mathcal{O}$ but there need not be a semialgebraic invariant are when the matrix A is diagonalisable and all eigenvalues have modulus one, as in the case at hand.

▶ **Example 2.** Consider the matrix

$$A = \frac{4}{25}\begin{pmatrix} 4 & -3 & 4 & -3 \\ 3 & 4 & 3 & 4 \\ 0 & 0 & 4 & -3 \\ 0 & 0 & 3 & 4 \end{pmatrix}$$

Matrix A has spectral radius $\frac{4}{5}$ and so $A^n x$ converges to 0 for any initial vector $x \in \mathbb{Q}^4$. Given a non-zero target $y \in \mathbb{Q}^4$ that does not lie in the orbit $x, Ax, A^2x, \ldots$, a natural candidate for an invariant is an initial segment of the orbit, together with some neighbourhood $\mathcal{N}$ of the origin in $\mathbb{R}^4$ that excludes y and is invariant under A. Note though that A is not contractive with respect to either the 1-norm or the 2-norm, so we cannot simply take $\mathcal{N}$ to be a ball of suitably small radius with respect to either of these norms. However, for $\varepsilon > 0$, the set

$$\mathcal{N}_\varepsilon = \left\{u \in \mathbb{R}^4 : u_1^2 + u_2^2 \le \varepsilon^2 \wedge u_3^2 + u_4^2 \le \tfrac{1}{16}\varepsilon^2\right\}$$

is invariant under A. Thus we obtain a semialgebraic invariant as the union of $\mathcal{N}_\varepsilon$, where ε is chosen sufficiently small such that $y \notin \mathcal{N}_\varepsilon$, together with an (easily computable) initial segment of the orbit $x, Ax, A^2x, \ldots$ comprising all points in the orbit that lie outside $\mathcal{N}_\varepsilon$.

▶ **Example 3.** Consider the following scaled version of the matrix from the previous example:

$$A = \frac{1}{5}\begin{pmatrix} 4 & -3 & 4 & -3 \\ 3 & 4 & 3 & 4 \\ 0 & 0 & 4 & -3 \\ 0 & 0 & 3 & 4 \end{pmatrix}.$$

Note that A is a non-diagonalisable matrix with spectral radius 1. Example 1 concerned an orthogonal matrix, while the matrix in Example 2 was (morally speaking, if not literally) length-decreasing. Here, by contrast, the idea is to identify a subset $\mathcal{Q} \subseteq \mathbb{R}^4$ that is invariant under A, together with a "length measure" $f : \mathcal{Q} \to \mathbb{R}$ that increases under application of A. Fixing a constant $c > 0$, such a set is

$$\mathcal{Q} = \left\{u \in \mathbb{R}^4 : u_1^2 + u_2^2 \geq c \wedge u_1u_3 + u_2u_4 \geq 0\right\}$$

with length measure $f(u) = u_1^2 + u_2^2$. A key property of $\mathcal{Q}$ is that for any vector $x \in \mathbb{R}^4$ such that $x_3 \neq 0$ or $x_4 \neq 0$, the orbit $x, Ax, A^2x, \ldots$ eventually enters $\mathcal{Q}$. By choosing c suitably large, we can exclude y from $\mathcal{Q}$. Thus we obtain an invariant as the union of $\mathcal{Q}$ and an appropriate finite intitial segment of the orbit $x, Ax, A^2x, \ldots$.

We would like to draw the reader's attention to the critical role played by the underlying domain $\mathbf{D}$ of potential invariants. In the examples above as well as the rest of this paper, we focus exclusively on the domain of semialgebraic sets. However one might naturally consider instead the domain of *semilinear* sets, *i.e.*, sets defined by Boolean combinations of linear inequalities with integer coefficients, or equivalently consisting of finite unions of (bounded or unbounded) rational polytopes. As pointed out above, in Example 1 no non-trivial instance admits a semilinear invariant, whereas one can show that in Example 2 semilinear invariants can always be found. Interestingly, the question of relative completeness (*i.e.*, determining in general whether or not a suitable semilinear invariant exists in non-reachability instances) is not known to be decidable, and appears to be a challenging problem.

2 Preliminaries

It is convenient in this paper to work over the field of (complex) algebraic numbers, denoted $\mathbb{A}$. All standard algebraic operations, such as sums, products, root-finding of polynomials and computing Jordan normal forms of matrices with algebraic entries can be performed effectively; we refer the reader to [4] for more details on the matter.

An *instance of the Orbit Problem*, or *Orbit instance* for short, is given by a square matrix $A \in \mathbb{A}^{d\times d}$ and two vectors $x, y \in \mathbb{A}^d$. The triple (A, x, y) is a *reachability* instance if there is $n \in \mathbb{N}$ such that $A^n x = y$, and otherwise is a *non-reachability* instance.

We are interested in non-reachability certificates given as invariants. Formally, given an Orbit instance (A, x, y) in dimension d, a set $\mathcal{P} \subseteq \mathbb{C}^d$ is a *non-reachability invariant* if $A\mathcal{P} \subseteq \mathcal{P}$, $x \in \mathcal{P}$, and $y \notin \mathcal{P}$.

For the remainder of this paper, we focus on *semialgebraic* invariants. Identifying $\mathbb{C}^d$ with $\mathbb{R}^{2d}$, a set $\mathcal{P}$ is semialgebraic if it is the set of real solutions of some Boolean combination of polynomial inequalities with integer coefficients.

A central result about semialgebraic sets is the Tarski-Seidenberg Theorem: if $S \subseteq \mathbb{R}^{n+1}$ is semialgebraic then the image $\pi(S)$ under the projection $\pi : \mathbb{R}^{n+1} \to \mathbb{R}^n$, where $\pi(x_1, \ldots, x_{n+1}) = (x_1, \ldots, x_n)$, is also semialgebraic. Among the consequences of this result is the fact that the topological closure of a semialgebraic set (in either $\mathbb{R}^n$ or $\mathbb{C}^n$) is again semialgebraic.

3 Semialgebraic Invariants

Our main result is the following.

▶ **Theorem 4.** *It is decidable whether an Orbit instance admits a semialgebraic invariant. Furthermore, there exists an algorithm which constructs such an invariant when it exists, and the invariant produced has polynomial-size description.*

The remainder of the paper is devoted to proving Theorem 4. To this end, let $\ell = (A, x, y)$ be a non-reachability Orbit instance in dimension d.[2]

As a first step, recall that every matrix A can be written in the form $A = Q^{-1}JQ$, where Q is invertible and J is in Jordan normal form. The following lemma transfers semialgebraic invariants through the change-of-basis matrix Q.

▶ **Lemma 5.** *Let $\ell = (A, x, y)$ be an Orbit instance, and Q an invertible matrix in $\mathbb{A}^{d\times d}$. Construct the Orbit instance $\ell_Q = (QAQ^{-1}, Qx, Qy)$. Then $\mathcal{P}$ is a semialgebraic invariant for ℓ_Q if, and only if, $Q^{-1}\mathcal{P}$ is a semialgebraic invariant for ℓ.*

Proof. First of all, $Q^{-1}\mathcal{P}$ is semialgebraic if, and only if, $\mathcal{P}$ is semialgebraic. We have:

- $QAQ^{-1}\mathcal{P} \subseteq \mathcal{P}$ if, and only if, $AQ^{-1}\mathcal{P} \subseteq Q^{-1}\mathcal{P}$,
- $Qx \in \mathcal{P}$ if, and only if, $x \in Q^{-1}\mathcal{P}$,
- $Qy \notin \mathcal{P}$, if, and only if, $y \notin Q^{-1}\mathcal{P}$.

This concludes the proof. ◀

Thanks to Lemma 5, we can reduce the problem of the existence of semialgebraic invariants for Orbit instances to cases in which the matrix is in Jordan normal form, *i.e.*, is a diagonal block matrix, where the blocks (called Jordan blocks) are of the form:

$$\begin{bmatrix} \lambda & 1 & & \\ & \lambda & \ddots & \\ & & \ddots & 1 \\ & & & \lambda \end{bmatrix}$$

Note that this transformation can be achieved in polynomial time [1, 2].

Formally, a Jordan block is a matrix $\lambda I + N$ with $\lambda \in \mathbb{A}$, I the identity matrix and N the matrix with 1's on the upper diagonal, and 0's everywhere else. The number λ is an eigenvalue of A. A Jordan block of dimension one is called diagonal, and A is diagonalisable if, and only if, all Jordan blocks are diagonal.

The d dimensions of the matrix A are indexed by pairs (J, k), where J ranges over the Jordan blocks and $k \in \{1, \dots, \delta\}$ where δ is the dimension of the Jordan block J. For instance, if the matrix A has two Jordan blocks, J_1 of dimension 1 and J_2 of dimension 2, then the three dimensions of A are $(J_1, 1)$ (corresponding to the Jordan block J_1) and $(J_2, 1), (J_2, 2)$ (corresponding to the Jordan block J_2).

For a vector v and a subset S of $\{1, \dots, d\}$, we denote v_S the projection vector of v on the dimensions in S, and extend this notation to matrices. As a special case, $v_{J,>k}$ denotes the vector restricted to the coordinates of the Jordan block J whose index is greater than k. We denote $\overline{S}$ the complement of S in $\{1, \dots, d\}$.

[2] Kannan and Lipton showed the decidability of reachability for Orbit instances over rational numbers; their proof carries over to instances with algebraic entries, however without the polynomial-time complexity.

There are a few degenerate cases which we handle now. We say that an Orbit instance $\ell = (A, x, y)$ in Jordan normal form is non-trivial if:

- There is no Jordan block associated with the value 0, or equivalently A is invertible,
- For each Jordan block J, both x_J and y_J are not the zero vector,
- For each non-diagonal Jordan block J, the vector x_J has at least a non-zero coordinate other than the first one, *i.e.*, $x_{J,>1}$ is not the zero vector.

▶ **Lemma 6.** *The existence of semialgebraic invariants for Orbit instances reduces in poylnomial time to the same problem for non-trivial Orbit instances in Jordan normal form.*

Proof. Let $\ell = (A, x, y)$ be an Orbit instance in Jordan normal form.

- If A is not invertible, we distinguish two cases.
  - If for some Jordan block J associated with the eigenvalue 0, we have that y is not the zero vector, *i.e.*, $y_J \neq 0$, then consider $\mathcal{P} = \{x, Ax, \dots, A^{d-1}x\} \cup \{z \in \mathbb{C}^d \mid z_J = 0\}$ is a semialgebraic invariant. Indeed, the Jordan block J is nilpotent, so for any vector u and $n \geq d$, we have that $J^n u = 0$, so in particular $(A^n x)_J = 0$. Moreover, since by assumption y is not reachable, it is not one of $A^n x$ for $n < d$, and $y_J \neq 0$, so $y \notin \mathcal{P}$.
  - Otherwise, denote J the dimensions corresponding to Jordan blocks associated with the eigenvalue 0, we have that $y_J = 0$. Consider the Orbit instance $\ell_J = (A_{\overline{J}}, (A^d x)_{\overline{J}}, y_{\overline{J}})$. We claim that ℓ admits a semialgebraic invariant if, and only if, ℓ_J does.

    Let $\mathcal{P}$ be a semialgebraic invariant for ℓ. Construct $\mathcal{P}_J$ the set of vectors z in $\mathbb{C}^{\overline{J}}$ such that z augmented with 0's in the J dimensions yields a vector in $\mathcal{P}$, we argue that $\mathcal{P}_J$ is a semialgebraic invariant for ℓ_J. Indeed, $(A^d x)_{\overline{J}} \in \mathcal{P}_J$ since $A^d x \in \mathcal{P}$ and $(A^d x)_J = 0$, because the Jordan block J is nilpotent. The stability of $\mathcal{P}_J$ under $A_{\overline{J}}$ is clear, and $y_{\overline{J}} \notin \mathcal{P}_J$ because $y_J = 0$, so $y_{\overline{J}} \in \mathcal{P}_J$ would imply $y \in \mathcal{P}$.

    Conversely, let $\mathcal{P}_J$ be a semialgebraic invariant for ℓ_J, extend it to $\mathcal{P} \subseteq \mathbb{C}^d$ by allowing any complex numbers in the J dimensions, then $\{x, Ax, \dots, A^{d-1}x\} \cup \mathcal{P}$ is a semialgebraic invariant for ℓ.

    We reduced the existence of semialgebraic invariants from ℓ to ℓ_J, with the additional property that the matrix is invertible.
- Suppose A contains a Jordan block J such that either $x_J = 0$ or $y_J = 0$. We distinguish three cases.
  - If for some Jordan block J we have $x_J = 0$ and $y_J \neq 0$, then $\mathcal{P} = \{z \in \mathbb{C}^d \mid z_J = 0\}$ is a semialgebraic invariant for ℓ.
  - If for some Jordan block J we have $x_J \neq 0$ and $y_J = 0$, let k such that $x_{J,k} \neq 0$ and $x_{J,>k} = 0$, then $\mathcal{P} = \{z \in \mathbb{C}^d \mid z_{J,k} \neq 0 \text{ and } z_{J,>k} = 0\}$ is a semialgebraic invariant for ℓ.
  - Otherwise, denote J the dimensions corresponding to Jordan blocks for which $x_J = y_J = 0$. Consider the Orbit instance $\ell_J = (A_{\overline{J}}, x_{\overline{J}}, y_{\overline{J}})$, we claim that ℓ admits a semialgebraic invariant if, and only if, ℓ_J does.

    Let $\mathcal{P}$ be a semialgebraic invariant for ℓ. Construct $\mathcal{P}_J$ the set of vectors z in $\mathbb{C}^{\overline{J}}$ such that z augmented with 0 in the J dimensions yields a vector in $\mathcal{P}$, then $\mathcal{P}_J$ is a semialgebraic invariant for ℓ_J.

    Conversely, let $\mathcal{P}_J$ be a semialgebraic invariant for ℓ_J, extend it to $\mathcal{P} \subseteq \mathbb{C}^d$ by allowing only 0 in the J dimensions, then $\mathcal{P}$ is a semialgebraic invariant for ℓ.

    We reduced the existence of semialgebraic invariants from ℓ to ℓ_J, with the additional property that for each Jordan block J, both x_J and y_J are not the zero vector.
- If A contains a non-diagonal Jordan block J such that the vector x_J is zero except on the first coordinate $(J, 1)$, we distinguish two cases.

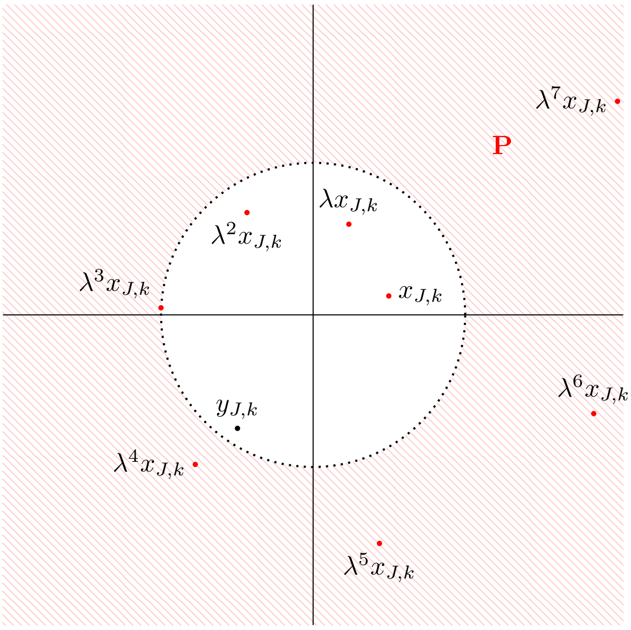

Figure 1 Case $|\lambda| > 1$. This figure represents the complex plane, which is the projection on the coordinate (J, k).

- If for some non-diagonal Jordan block J we have that $x_{J,>1} = 0$ and $y_{J,>1} \neq 0$, then $\mathcal{P} = \left\{z \in \mathbb{C}^d \mid z_{J,>1} = 0\right\}$ is a semialgebraic invariant for ℓ.
- Otherwise, denote J the dimensions corresponding to non-diagonal Jordan blocks for which $x_{J,>1} = y_{J,>1} = 0$. Let $S = \overline{J} \cup \bigcup_J (J, 1)$, *i.e.*, the dimensions outside J plus the first dimensions of each block in J. Consider the Orbit instance $\ell_S = (A_S, x_S, y_S)$, we claim that ℓ admits a semialgebraic invariant if, and only if, ℓ_S does.

  Let $\mathcal{P}$ be a semialgebraic invariant for ℓ. Construct $\mathcal{P}_S$ the set of vectors z in $\mathbb{C}^S$ such that z augmented with 0 in the $\overline{S}$ dimensions yields a vector in $\mathcal{P}$, then $\mathcal{P}_S$ is a semialgebraic invariant for ℓ_S.

  Conversely, let $\mathcal{P}_S$ be a semialgebraic invariant for ℓ_S, extend it to $\mathcal{P} \subseteq \mathbb{C}^d$ by allowing only 0 in the $\overline{S}$ dimensions, then $\mathcal{P}$ is a semialgebraic invariant for ℓ.

  We reduced the existence of semialgebraic invariants from ℓ to ℓ_S, with the additional property that for each non-diagonal Jordan block J, $x_{J,>1}$ is not the zero vector.

This concludes the proof. ◀

3.1 Some eigenvalue has modulus different from one

▶ **Lemma 7.** *Let $\ell = (A, x, y)$ be a non-trivial Orbit instance in Jordan normal form. Assume that ℓ is a non-reachability instance. If the matrix A has an eigenvalue whose modulus is not equal to* 1*, then there exists a semialgebraic invariant for ℓ.*

Proof. We distinguish two cases according to whether there exists an eigenvalue of modulus strictly more than 1 or en eigenvalue of modulus strictly less than 1.

- Suppose that A contains a Jordan block J associated with an eigenvalue λ with $|\lambda| > 1$. In this case, some coordinate of $(A^n x)_{n\in\mathbb{N}}$ diverges to infinity, so eventually gets larger in modulus than the corresponding coordinate in y. This allows us to construct a

semialgebraic invariant for ℓ by taking the first points and then all points having a large coordinate in the diverging dimension. This case is illustrated in Figure 1.
By assumption x_J is non-zero, let k such that $x_{J,k} \neq 0$ and $x_{J,>k} = 0$ (Note that if the Jordan block J is diagonal, $k = 1$). For all $n \in \mathbb{N}$, we have $(A^n x)_{J,k} = \lambda^n x_{J,k}$, so $|(A^n x)_{J,k}|$ diverges to infinity. It follows that there exists $n_0 \in \mathbb{N}$ such that $|(A^{n_0} x)|_{J,k} > |y_{J,k}|$. Let

$$\mathcal{P} = \left\{x, Ax, \ldots, A^{n_0-1}x\right\} \cup \left\{z \in \mathbb{C}^d \mid |z_{J,k}| \geq |(A^{n_0}x)_{J,k}| \text{ and } z_{J,>k} = 0\right\}.$$

We argue that $\mathcal{P}$ is a semialgebraic invariant for ℓ. The non-trivial point is that $\mathcal{P}$ is stable under A. Note that $(A^{n_0}x)_{J,>k} = 0$, so $A^{n_0}x \in \mathcal{P}$. Let $z \in \mathbb{C}^d$ such that $|z_{J,k}| \geq |(A^{n_0}x)_{J,k}|$ and $z_{J,>k} = 0$. Then $(Az)_{J,k} = \lambda z_{J,k}$ and $(Az)_{J,>k} = 0$, so $Az \in \mathcal{P}$.

- If A contains a Jordan block J associated with an eigenvalue λ with $|\lambda| < 1$.
The situation is similar to the former, except that the convergence is towards the origin. The construction of the semialgebraic invariant is much more subtle though, for the following reason: for k such that $x_{J,k} \neq 0$ and $x_{J,>k} = 0$, we may have that $y_{J,k} = 0$, implying that $((A^n x)_{J,k})_{n\in\mathbb{N}}$ does not become smaller than $y_{J,k}$. Working on another dimension implies to give up the following diagonal behaviour: $(A^n x)_{J,k} = \lambda^n x_{J,k}$, making it hard to find a stable set under A. To overcome this problem, the invariant we define depends upon all the coordinates of the Jordan block J.
Denote $d(J)$ the dimension of the Jordan block J. We have that $((A^n x)_J)_{n\in\mathbb{N}}$ converges to 0. It follows that there exists $n_0 \in \mathbb{N}$ such that for each dimension k of the Jordan block J, *i.e.*, for $k \in \{1, \ldots, d(J)\}$, we have $|(A^{n_0}x)_{J,k}| \leq (1 - |\lambda|)^k \cdot ||y_J||_\infty$.
Let

$$\mathcal{P} = \left\{x, Ax, \ldots, A^{n_0-1}x\right\} \cup \left\{z \in \mathbb{C}^d \mid \forall k \in \{1, \ldots, d(J)\}, |z_{J,k}| \leq (1 - |\lambda|)^k \cdot ||y_J||_\infty\right\}.$$

We argue that $\mathcal{P}$ is a semialgebraic invariant for ℓ. Note that $y \notin \mathcal{P}$ since for k such that $||y_J||_\infty = |y_{J,k}|$, this would imply $||y_J||_\infty \leq (1 - |\lambda|)^k \cdot ||y_J||_\infty$, which cannot be since $k \geq 1$, $y_J \neq 0$ and $|\lambda| < 1$. We examine the stability of $\mathcal{P}$ under A. Let $z \in \mathbb{C}^d$ such that for each dimension $k \in \{1, \ldots, d(J)\}$, we have $|z_{J,k}| \leq (1 - |\lambda|)^k \cdot ||y_J||_\infty$. Let $k < d(J)$, then

$$\begin{array}{rcl} |(Az)_{J,k}| = |\lambda z_{J,k} + z_{J,k+1}| & \leq & |\lambda||z_{J,k}| + |z_{J,k+1}| \\ & \leq & |\lambda|(1 - |\lambda|)^k \cdot ||y_J||_\infty + (1 - |\lambda|)^{k+1} \cdot ||y_J||_\infty \\ & = & (|\lambda| + (1 - |\lambda|))(1 - |\lambda|)^k \cdot ||y_J||_\infty \\ & = & (1 - |\lambda|)^k \cdot ||y_J||_\infty. \end{array}$$

The case $k = d(J)$ is similar but easier.

This concludes the proof. ◀

3.2 All eigenvalues have modulus one and the matrix is not diagonalisable

▶ **Lemma 8.** *Let $\ell = (A, x, y)$ be a non-trivial Orbit instance in Jordan normal form and assume that ℓ is a non-reachability instance. If all the eigenvalues of the matrix A have modulus 1 and A is not diagonalisable, then there exists a semialgebraic invariant for ℓ.*

We illustrate the construction of the semialgebraic invariant in an example following the proof. (See also Example 3 from the Introduction.)

Proof. By assumption, there exists a non-diagonal Jordan block J. Since ℓ is non-trivial, x has a non-zero coordinate in J which is not the first one. Let k such that $x_{J,k} \neq 0$ and $x_{J,>k} = 0$, we have $k \geq 2$ and

$$(A^n x)_{J,k-1} = \lambda^n x_{J,k-1} + n\lambda^{n-1} x_{J,k},$$

so $(|(A^n x)_{J,k-1}|)_{n\in\mathbb{N}}$ diverges to infinity since $|\lambda| = 1$. It follows that there exists $n_0 \in \mathbb{N}$ such that $|(A^{n_0} x)_{J,k-1}| > |y_{J,k-1}|$. Without loss of generality we assume $n_0 \geq -\frac{\langle \lambda x_{J,k-1}, x_{J,k}\rangle}{|x_{J,k}|^2}$. The notation $\langle u, v\rangle$ designates the scalar product of the complex numbers u and v viewed as vectors in $\mathbb{R}^2$, defined by $\operatorname{Re}(u\overline{v})$. This quantity will appear later; note that it only depends on x and A.

Let

$$\mathcal{P} = \left\{x, Ax, \dots, A^{n_0-1}x\right\} \cup \left\{ z \in \mathbb{C}^d \;\middle|\; \begin{array}{l} |z_{J,k-1}| \geq |(A^{n_0}x)_{J,k-1}|, \text{ and} \\ \langle \lambda z_{J,k-1}, z_{J,k}\rangle \geq 0, \text{ and } z_{J,>k} = 0 \end{array} \right\}.$$

We argue that $\mathcal{P}$ is a semialgebraic invariant for ℓ. It is a semialgebraic set: the condition $\langle \lambda z_{J,k-1}, z_{J,k}\rangle \geq 0$ is of the form $P(z) \geq 0$ for a polynomial P with algebraic coefficients, where z is seen as a vector in $\mathbb{R}^{2d}$. The part to be looked at closely is the stability of $\mathcal{P}$ under A.

First, $A^{n_0}x \in \mathcal{P}$. Indeed, using $|\lambda| = 1$ and the assumption on n_0:

$$\begin{array}{rcl} \langle \lambda (A^{n_0}x)_{J,k-1}, (A^{n_0}x)_{J,k}\rangle & = & \langle \lambda \cdot \left(\lambda^{n_0} x_{J,k-1} + n_0 \lambda^{n_0-1} x_{J,k}\right), \lambda^{n_0} x_{J,k}\rangle \\ & = & |\lambda^{n_0}|^2 \langle \lambda x_{J,k-1}, x_{J,k}\rangle + n_0 |\lambda^{n_0} x_{J,k}|^2 \\ & = & \langle \lambda x_{J,k-1}, x_{J,k}\rangle + n_0 |x_{J,k}|^2 \\ & \geq & 0. \end{array}$$

Now, let $z \in \mathbb{C}^d$ such that $|z_{J,k-1}| \geq |(A^{n_0}x)_{J,k-1}|$, $\langle \lambda z_{J,k-1}, z_{J,k}\rangle \geq 0$ and $z_{J,>k} = 0$. We have $(Az)_{J,k-1} = \lambda z_{J,k-1} + z_{J,k}$, $(Az)_{J,k} = \lambda z_{J,k}$ and $(Az)_{J,>k} = 0$. It follows that:

$$\begin{array}{rcl} |(Az)_{J,k-1}|^2 & = & |\lambda z_{J,k-1} + z_{J,k}|^2 \\ & = & |z_{J,k-1}|^2 + 2\langle \lambda z_{J,k-1}, z_{J,k}\rangle + |z_{J,k}|^2 \\ & \geq & |z_{J,k-1}|^2 \\ & \geq & |(A^{n_0}x)_{J,k-1}|^2, \end{array}$$

and:

$$\begin{array}{rcl} \langle \lambda (Az)_{J,k-1}, (Az)_{J,k}\rangle & = & \langle \lambda(\lambda z_{J,k-1} + z_{J,k}), \lambda z_{J,k}\rangle \\ & = & |\lambda|^2 \langle \lambda z_{J,k-1} + z_{J,k}, z_{J,k}\rangle \\ & = & \langle \lambda z_{J,k-1}, z_{J,k}\rangle + |z_{J,k}|^2 \\ & \geq & 0. \end{array}$$

Hence $Az \in \mathcal{P}$, and $\mathcal{P}$ is a semialgebraic invariant for ℓ. ◀

▶ **Example 9.** Consider the following matrix:

$$A = \begin{bmatrix} e^{i\theta} & 1 \\ 0 & e^{i\theta} \end{bmatrix},$$

where $\theta \in \mathbb{R}$ is an angle such that $\frac{\theta}{\pi} \notin \mathbb{Q}$. We start from the vector $x = [1,\ 1]^T$. We have

$$A^n x = \left[e^{in\theta} + ne^{i(n-1)\theta},\ e^{in\theta}\right],$$

so the projection on the second coordinate is a dense subset of the unit circle, and the projection on the first coordinate describes a growing spiral (similar to that shown in Figure 1). A tentative invariant for excluding some vector y is the complement of a circle on the first coordinate, large enough not to include y. However, this set is not a priori invariant. Geometrically, the action of A on a vector $[z_1,\ z_2]$ is to rotate both z_1 and z_2 by an angle of θ, and to push the first coordinate in the direction of z_2:

$$A\,[z_1,\ z_2] = \left[e^{i\theta} z_1 + z_2,\ e^{i\theta} z_2\right].$$

A natural way to restrict the above set to make it invariant is to ensure that z_2 pushes away from the origin, *i.e.*, that the norm of $(Az)_1$ increases. This is achieved by requiring that $\langle e^{i\theta} z_1, z_2\rangle \geq 0$.

3.3 All eigenvalues have modulus one and the matrix is diagonalisable

This case is the most involved and is the only one in which it might hold that y not be reachable and yet no semialgebraic invariant exists. (Recall Example 1 from the Introduction.) Using results from Diophantine approximation and algebraic number theory, we show that the topological closure of the orbit $\overline{\{A^n x : x \in \mathbb{N}\}}$ is (effectively) semialgebraic. Furthermore, using topological properties of semialgebraic sets we show that any semialgebraic invariant must contain the closure of the orbit. It follows that there exists a semialgebraic invariant just in case $y \notin \overline{\{A^n x : x \in \mathbb{N}\}}$.

We start with the following topological fact about semialgebraic sets.

▶ **Lemma 10.** *Let $E, F \subseteq \mathbb{R}^n$ be two sets such that $\overline{E} = \overline{F}$ and F is semialgebraic. Then $E \cap F \neq \emptyset$.*

Proof. The proof uses the notion of the dimension of a semialgebraic set. The formal definition of dimension uses the cell-decomposition theorem (see, e.g., [5, Chapter 4]). However to establish the lemma it suffices to note the following two properties of the dimension. First, for any semi-algebraic set $X \subseteq \mathbb{R}^n$ set we have $\dim(X) = \dim(\overline{X})$ [5, Chapter 4, Theorem 1.8]. Secondly, if $X \subseteq Y$ are semi-algebraic subsets of $\mathbb{R}^n$ that have the same dimension, then X has non-empty interior in Y [5, 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) in $\overline{F} = \overline{E}$. But then E is dense in $\overline{E}$ while F has non-empty interior in $\overline{E}$, and thus E and F meet. ◀

▶ **Lemma 11.** *Let $\ell = (A, x, y)$ be an Orbit instance, where $A = \operatorname{diag}(\lambda_1, \ldots, \lambda_d)$ is a diagonal $d \times d$ matrix with entries $\lambda_1, \ldots, \lambda_d \in \mathbb{C}$ all having modulus one. Write $\mathcal{O} = \{A^n x : n \in \mathbb{N}\}$ for the orbit of x under A. Then*

- *The topological closure of $\mathcal{O}$ in $\mathbb{C}^d$ is a semi-algebraic set that is computable from ℓ in polynomial space.*
- *Any semi-algebraic invariant for ℓ contains $\overline{\mathcal{O}}$.*

Proof. We start by proving the first item.

Write $\mathbb{T}$ for the unit circle in $\mathbb{C}$. Let

$$L_A = \left\{v \in \mathbb{Z}^d \mid \lambda_1^{v_1} \cdots \lambda_d^{v_d} = 1\right\}$$

be the set of all multiplicative relations holding among $\lambda_1, \ldots, \lambda_d$. Notice that L_A is an additive subgroup of $\mathbb{Z}^d$. Consider the set of diagonal $d \times d$ matrices

$$T_A = \left\{\operatorname{diag}(\mu_1, \ldots, \mu_d) \mid \mu \in \mathbb{T}^d \text{ and } \forall v \in L_A\, (\mu_1^{v_1} \cdots \mu_d^{v_d} = 1)\right\}$$

whose diagonal entries satisfy the multiplicative relations in L_A. Notice that T_A forms a group under matrix multiplication that is also a closed subset of $\mathbb{C}^{d\times d}$.

Using Kronecker's Theorem on inhomogeneous simultaneous Diophantine approximation [3], it is shown in [11, Proposition 3.5] that $\{A^n : n \in \mathbb{N}\}$ is a dense subset of T_A. This immediately gives

$$\overline{\mathcal{O}} = \overline{\{A^n x : n \in \mathbb{N}\}} = \{Mx : M \in T_A\}\,. \tag{1}$$

We now show that $\overline{\mathcal{O}}$ is semi-algebraic. Observe that L_A is finitely generated, being a subgroup of a finitely generated group. Moreover, if $B \subseteq L_A$ is a basis of L_A then we can write

$$T_A = \left\{\operatorname{diag}(\mu_1, \ldots, \mu_d) \mid \mu \in \mathbb{T}^d \text{ and } \forall v \in B\, (\mu_1^{v_1} \cdots \mu_d^{v_d} = 1)\right\}\,.$$

It follows that T_A is a semi-algebraic subset of $\mathbb{C}^{d\times d}$ and thus from (1) that $\overline{\mathcal{O}}$ is a semi-algebraic set.

From an upper bound on the length of B due to Masser [10], it can be shown that one can compute a basis for L_A in polynomial space in the description of A (see [11, Corollary 3.3]) and thereby compute a description of T_A as a semi-algebraic set, also in polynomial space in the description of A.

Now we move to the second item in the statement of the lemma. Let $\mathcal{P}$ be a semi-algebraic invariant for ℓ. Our goal is to show that $\overline{\mathcal{O}} \subseteq \mathcal{P}$. To show this we can, without loss of generality, replace $\mathcal{P}$ by $\mathcal{P} \cap \overline{\mathcal{O}}$, since the latter is also a semi-algebraic invariant. Moreover, since any invariant necessarily contains the orbit $\mathcal{O}$, we may suppose that $\mathcal{O} \subseteq \mathcal{P} \subseteq \overline{\mathcal{O}}$, and hence $\overline{\mathcal{P}} = \overline{\mathcal{O}}$.

We now prove that $\overline{\mathcal{O}} \subseteq \mathcal{P}$, that is, we pick an arbitrary element $z \in \overline{\mathcal{O}}$ and show that $z \in \mathcal{P}$. To this end, consider the orbit of z under the matrix A^{-1}. Now $A^{-1} = \operatorname{diag}(\lambda_1^{-1}, \ldots, \lambda_d^{-d})$ and we may define groups $L_{A^{-1}}$ and $T_{A^{-1}}$ analogously with L_A and T_A. In fact it is clear that L_A and $L_{A^{-1}}$ coincide (i.e., $\lambda_1, \ldots, \lambda_d$ satisfy exactly the same multiplicative relations as $\lambda_1^{-1}, \ldots, \lambda_d^{-1}$), and hence also $T_A = T_{A^{-1}}$.

Now we claim that the following chain of equalities holds:

$$\begin{aligned}
\overline{\{A^{-n}z : n \in \mathbb{N}\}} &= \{Mz : M \in T_{A^{-1}}\} && (2)\\
&= \{Mz : M \in T_A\} && (3)\\
&= \{Mx : M \in T_A\} && (4)\\
&= \overline{\mathcal{O}} \;=\; \overline{\mathcal{P}}\,.
\end{aligned}$$

Indeed, Equation (2) is an instance of (1), but with A^{-1} and z in place of A and x. Equation (3) follows from the fact that $T_A = T_{A^{-1}}$. To see Equation (4), observe from (1) that z has the form $M_0 x$ for some $M_0 \in T_A$. But $\{MM_0x : M \in T_A\} = \{Mx : M \in T_A\}$ since T_A, being a group, contains M_0^{-1}.

Now we have established that

$$\overline{\{A^{-n}z : n \in \mathbb{N}\}} = \overline{\mathcal{P}}\,.$$

Then by Lemma 10 we have that $A^{-n}z$ lies in $\mathcal{P}$ for some $n \in \mathbb{N}$. But since $\mathcal{P}$ is invariant under A we have $z \in \mathcal{P}$. ◀

▶ **Corollary 12.** *Let the Orbit instance ℓ be as described in Lemma 11. Then ℓ admits a semi-algebraic invariant if and only if $y \notin \overline{\mathcal{O}}$.*

Proof. If $y \notin \overline{\mathcal{O}}$, then $\overline{\mathcal{O}}$ is a semi-algebraic invariant for ℓ by the first item in Lemma 11. Conversely, if there exists a semi-algebraic invariant $\mathcal{P}$ for ℓ, then $\overline{\mathcal{O}} \subseteq \mathcal{P}$ by the second item in Lemma 11, implying that $y \notin \overline{\mathcal{O}}$. ◀

3.4 Proof of Theorem 4

We now draw together the results of the previous sections to prove our main result, Theorem 4, giving an effective characterisation of the existence of semialgebraic invariants and a procedure to compute such an invariant when it exists.

Let $\ell = (A, x, y)$ be a non-reachability Orbit instance. First we put A in Jordan normal form and simplify ℓ to obtain a non-trivial Orbit instance. We then divide into three cases.

- If some eigenvalue of A has modulus different from 1 then there is a semialgebraic invariant (see Section 3.1).
- If all eigenvalues have modulus 1 and the matrix is not diagonalisable then there is a semialgebraic invariant (see Section 3.2).
- If all eigenvalues have modulus 1 and the matrix is diagonalisable, then there exists a semialgebraic invariant if and only if the topological closure of the orbit $\overline{\{A^n x : n \in \mathbb{N}\}}$ is such an invariant, which holds if and only if the closure does not contain y (see Section 3.3). Note therefore that non-reachability Orbit instances for which there do not exist semialgebraic invariants are extremely sparse.

Thus we obtain an effective characterisation of the class of Orbit instances for which there exists a semialgebraic invariant. Moreover in those cases in which there exists an invariant we have shown how to compute such an invariant in polynomial space.

4 Conclusions

This paper is a first step towards the study of invariants for discrete linear dynamical systems. At present, the question of the existence and of the algorithmic synthesis of suitable invariants for higher-dimensional versions of the Orbit Problem (i.e., when the 'target' y to be avoided consists of either a vector space, a polytope, or some other higher-dimensional object) is completely open. Given, as pointed out earlier, that reachability questions with high-dimensional targets appear themselves to be very difficult, one does not expect the corresponding invariant synthesis problems to be easy, yet this approach might prove a tractable alternative well worth exploring.

Our main result is a polynomial-space procedure for deciding existence and computing semialgebraic invariants in instances of the Orbit Problem. The only obstacle to obtaining a polynomial-time bound is the problem of computing a basis of the group of all multiplicative relations among a given collection of algebraic numbers $\alpha_1, \ldots, \alpha_d$, which is not known to be solvable in polynomial time. Less ambitiously one can ask for a polynomial-time procedure to verify a putative relation $\alpha_1^{n_1} \ldots \alpha_d^{n_d} \stackrel{?}{=} 1$. Assuming that $\alpha_1, \ldots, \alpha_d$ are represented as elements of an explicitly given finite-dimensional algebra K over $\mathbb{Q}$, Ge [6] gave a polynomial-time algorithm for verifying multiplicative relations. In our setting, however, where $\alpha_1, \ldots, \alpha_d$ are roots of the characteristic polynomial of matrix A, the dimension of K may be exponential in d.

References

1 Jin-Yi Cai. Computing Jordan normal forms exactly for commuting matrices in polynomial time. Technical report, SUNY at Buffalo, 2000.

2 Jin-Yi Cai, Richard J. Lipton, and Yechezkel Zalcstein. The complexity of the A B C problem. *SIAM J. Comput.*, 29(6):1878–1888, 2000. doi:10.1137/S0097539794276853.

3 John W. S. Cassels. *An introduction to Diophantine approximation*. Cambridge University Press, 1965.

4 Ventsislav Chonev, Joël Ouaknine, and James Worrell. On the complexity of the Orbit Problem. *Journal of the ACM*, 63(3):23, 2016. doi:10.1145/2857050.

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

6 G. Ge. Testing equalities of multiplicative representations in polynomial time. In *Proceedings of SFCS*, pages 422–426. IEEE Computer Society, 1993.

7 Michael A. Harrison. *Lectures on linear sequential machines*. New York-Londres, Academic Press, 1969.

8 Ravindran Kannan and Richard J. Lipton. The Orbit Problem is decidable. In *Proceedings of STOC*, pages 252–261, 1980. doi:10.1145/800141.804673.

9 Ravindran Kannan and Richard J. Lipton. Polynomial-time algorithm for the Orbit Problem. *Journal of the ACM*, 33(4):808–821, 1986. doi:10.1145/6490.6496.

10 David W. Masser. Linear relations on algebraic groups. In Alan Baker, editor, *New Advances in Transcendence Theory*, pages 248–262. Cambridge University Press, 1988. doi:10.1017/CBO9780511897184.016.

11 Joël Ouaknine and James Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In *Proceedings of ICALP*, pages 330–341, 2014. doi:10.1007/978-3-662-43951-7_28.

12 Terence Tao. *Structure and Randomness*. AMS, 2008.