Complete Semialgebraic Invariant Synthesis for the Kannan-Lipton Orbit Problem

Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, James Worrell

► To cite this version:

Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, James Worrell. Complete Semialgebraic Invariant Synthesis for the Kannan-Lipton Orbit Problem. *Theory of Computing Systems*, 2019, 63 (5), pp.1027-1048. 10.1007/s00224-019-09913-3 . hal-02503357

HAL Id: hal-02503357

<https://hal.science/hal-02503357>

Submitted on 18 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Complete Semialgebraic Invariant Synthesis for the Kannan-Lipton Orbit Problem

Nathanaël Fijalkow · Pierre Ohlmann ·
Joël Ouaknine · Amaury Pouly · James
Worrell

Received: date / Accepted: date

Abstract The *Orbit Problem* consists of determining, given a matrix 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 but 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 succinct invariants of polynomial size.

Nathanaël Fijalkow was supported by the Alan Turing Institute under EPSRC grant EP/N510129/1, Nathanaël Fijalkow, Pierre Ohlmann, and Amaury Pouly were supported by the CODYS project ANR-18-CE40-0007, Joël Ouaknine was supported by ERC grant AVS-ISS (648701) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) — Projektnummer 389792660 — TRR 248, and James Worrell was supported by EPSRC Fellowship EP/N008197/1.

Nathanaël Fijalkow
CNRS, LaBRI, Bordeaux, France, and
The Alan Turing Institute of data science and artificial intelligence, London, UK

Pierre Ohlmann
IRIF, Université Paris Diderot - Paris 7, France

Joël Ouaknine
Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany, and
Department of Computer Science, Oxford University, UK

Amaury Pouly
Max Planck Institute for Software Systems (MPI-SWS), Saarland Informatics Campus, Germany

James Worrell
Department of Computer Science, Oxford University, UK

Our results imply that the class of closed semialgebraic invariants is *closure-complete*: there exists a closed semialgebraic invariant if and only if y is not in the topological closure of the orbit of x under A .

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

1 Introduction

The *Orbit Problem* was introduced by Kannan and Lipton in the seminal papers [KL80, KL86], 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 [Har69]. 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, an interesting and natural 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 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). In the context of imperative programs there has been work on synthesizing affine and algebraic invariants [MS04a, MS04b, Col07], as well as polyhedral invariants [CH78], but as far we know no work on the more general class of semialgebraic invariants.

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 $y' \notin \mathcal{P}$.

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 [KL86], 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 [Tao08].

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:

- *Completeness*: in instances of non-reachability, does a suitable invariant in $\mathbf{D}$ *always* exist?
- *Effectiveness*: if not, 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?

The completeness question can be further refined when considering *topologically closed* invariants for the Euclidian topology. Indeed, if y is in the topological closure of the orbit, then there cannot exist a closed invariant. The converse is a completeness property:

- *Closure-Completeness*: if the target vector is not in the topological closure of the orbit, does there exist a closed invariant in $\mathbf{D}$?

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 show that whether a suitable semialgebraic invariant exists or not is decidable in polynomial space, and moreover in positive instances we show how to synthesise a suitable invariant of polynomial size. Further, checking whether a semialgebraic set is an invariant can be done in polynomial space. Both results hold for both semialgebraic and closed semialgebraic invariants.

- We provide a simple characterisation of instances of non-reachability for which there does not exist a suitable semialgebraic invariant,
- We obtain that the class of closed semialgebraic invariants is closure-complete, *i.e.*, there exists a closed semialgebraic invariant if and only if the target vector is not in the topological closure of the orbit.

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 [KL80, KL86]. 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

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

The 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 $\overline{\mathcal{O}}$ of the orbit $\mathcal{O} = \{x, Ax, A^2x, \dots\}$ 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 $\mathbb{R}^2$, see the remark following the examples.) More general orthogonal matrices can be handled along similar lines, but the analysis becomes substantially more involved. In general, the only cases in which $y \notin \mathcal{O}$ but there is no 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

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

The 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 $x, Ax, A^2x, \dots$, 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 however 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

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

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 $x, Ax, A^2x, \dots$ 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” $f : \mathcal{Q} \rightarrow \mathbb{R}$ that increases under application of A . Fixing a constant $c > 0$, such a set is

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

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 $x, Ax, A^2x, \dots$ 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 initial segment of the orbit $x, Ax, A^2x, \dots$.

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, 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

Semialgebraic sets

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.

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 [BCR98] for more details on the matter.

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 $\pi : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^n$, where $\pi(x_1, \dots, x_{n+1}) = (x_1, \dots, x_n)$, 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.

Succinct semialgebraic sets

The standard representation of semialgebraic sets involves representing polynomials as an array of coefficients, or equivalently with exponents written in unary. By contrast, we will need to represent polynomials with few terms but large exponents, such as $f(x) = x^{2^{100}}$. Therefore we cannot use the standard representation of semialgebraic set and maintain good complexity bounds. Furthermore, it will be important to allow for composition without an exponential blowup: for example $f(x+y) = (x+y)^{2^{100}}$ has $2^{100} + 1$ monomials over x and y . This rules out the usual *sparse* representation where polynomials are given by a list of coefficients. We introduce a *succinct* representation where polynomials are represented as terms generated by the following grammar:

$$p ::= x \mid v \mid p + p \mid p \cdot p$$

where x denotes a variable and $v \in \mathbb{A}$ a (real) algebraic number. The size of such a term is the number of subterms (in other words we see a term as a DAG), noting that we use the standard representation of algebraic numbers (in particular, integers coefficients are written in binary). For example, f above has size approximately 100 since f can be defined by taking x and squaring it 100 times. Composing two such polynomials amounts to a substitution of variables and only increases the size linearly, so $f(x+y)$ also has size roughly 100. We can also have a small representation for the real and imaginary parts of $f(x+iy)$. As an example, consider r_n the real part of $(x+iy)^{2^n}$ and i_n its imaginary part. Then $r_{n+1} = r_n^2 - i_n^2$ and $i_{n+1} = 2r_n i_n$. This yields a representation for r_{n+1} and i_{n+1} of size 6 plus the sizes of r_n and i_n , thus linear in n . More generally, the size of the representation of the real and imaginary parts of $f(x+iy)$ is linear in the size of the representation of f .

We extend this succinct representation to formulas by

$$\varphi ::= (p = 0) \mid (p > 0) \mid \varphi \wedge \varphi \mid \varphi \vee \varphi$$

and we define the size of a formula in the usual way, with the size of atoms using the succinct representation above. The other comparisons $\geq, <, \leq$ and $\neq$ can be rewritten using $=$ and $>$ only with a linear blowup in size: for example $(p \neq 0) \iff (p < 0) \vee (p > 0)$. Similarly, negation is not included because it can always be pushed down to the comparisons operators; this will be important for the following lemma.

A *succinct* representation for a semialgebraic set $S \subseteq \mathbb{R}^n$ is a succinct formula φ such that $S = \{x \in \mathbb{R}^n : \varphi(x)\}$. The size of this representation is the size of φ .

The usual complexity results on semialgebraic sets and the existential theory of reals assume the standard encoding. It is therefore important to point out that we can translate our succinct encoding into a standard one by introducing existential quantifiers. This translation shows that some *decision problems* on succinct semialgebraic sets are decidable.

Lemma 1 *Let $S \subseteq \mathbb{R}^n$ be a semialgebraic set with succinct representation φ . Then there exists $k \in \mathbb{N}$ and a formula ϕ , using the standard representation, such that*

$$S = \{x \in \mathbb{R}^n : \exists y \in \mathbb{R}^k, \phi(x, y)\}$$

and the size of (the standard representation of) ϕ is polynomial in the size of (the succinct representation of) φ . Furthermore, there is a PTIME algorithm to compute k and ϕ from φ .

Proof It is sufficient to show the result for $\varphi \equiv (p = 0)$ or $\varphi \equiv (p > 0)$. We introduce one variable for each subterm of p , and mimic the construction of p . For instance, if $p = p_1 \cdot p_2$, then we add $y = y_1 \cdot y_2$ in the formula ϕ , where the variables y, y_1, y_2 correspond to the terms p, p_1, p_2 .

In particular, membership and semialgebraic set inclusion (with a twist), remain decidable in polynomial space:

Lemma 2 *The following problems can be solved in polynomial space, where S is a semialgebraic set given by a succinct representation:*

- compute a succinct representation for S^c , the complement of S ,
- compute a succinct representation for AS where A is an invertible matrix,
- decide if $x \in S$,
- decide if $AS \subseteq T$ where A is a matrix and S is a semialgebraic sets,
- compute a succinct representation for $\{A^k x\}$ where A is a matrix and k is given in binary.

In all problems, the coefficients of A and x are algebraic numbers.

Proof The first item is clear: the complement of S is define by the negation of its succinct formula. By pushing the negation to the polynomials, we can obtain a new succinct formula with a linear blowup in size.

The second item is also clear: first compute $B = A^{-1}$, which has algebraic coefficients and can be computed in polynomial time. Write $S = \{x : \varphi(x)\}$, then $AS = \{x : \varphi(Bx)\} = \{x : \varphi'(x)\}$ where $\varphi' = \varphi[x \leftarrow Bx]$ is a substitution and thus only incurs a polynomial blowup in size.

The third item follows from Lemma 1: compute a formula ϕ such that $S = \{x : \exists y, \phi(x, y)\}$. Then we have

$$x \in S \iff \exists y. \phi(x, y)$$

which is a first order formula in the existential theory of the reals, and is thus decidable in polynomial space. Note that it is crucial that the formula ϕ uses the standard representation.

The fourth item is a bit more involved since we cannot easily compute a succinct representation for AS in the case where A is not invertible. By the first item and Lemma 1, we can compute a formula ψ such that $T^c = \{x : \exists y. \psi(x, y)\}$. By Lemma 1 we can compute a formula ϕ such that $S = \{x : \exists y, \phi(x, y)\}$. But then

$$\begin{aligned} AS \subseteq T &\iff AS \cap T^c = \emptyset \\ &\iff \{z : \exists x \in S. z = Ax\} \cap T^c = \emptyset \\ &\iff \{z : \exists x, y. \phi(x, y) \wedge z = Ax\} \cap T^c = \emptyset \\ &\iff \{z : \exists x, y. \phi(x, y) \wedge z = Ax \wedge \exists y. \psi(z, y)\} = \emptyset \\ &\iff \neg \exists z. (\exists x, y. \phi(x, y) \wedge z = Ax \wedge \exists y. \psi(z, y)) \end{aligned}$$

which is decidable in polynomial space by the existential theory of the reals.

To show the fifth item, one can compute a Jordan normal form PJP^{-1} for A in polynomial time, then observe that $A^k x = PJ^k P^{-1} x$, thus it is enough to compute a succinct representation for the entries of J^k . This is done block-wise, using the following formula for a Jordan block of size d :

$$\begin{bmatrix} \lambda & 1 & & \\ & \ddots & \ddots & \\ & & \lambda & 1 \\ & & & \lambda \end{bmatrix} = \begin{bmatrix} \lambda^k & \binom{k}{1} \lambda^{k-1} & \dots & \binom{k}{d} \lambda^{k-1} \\ & \lambda^k & \binom{k}{1} \lambda^{k-1} & \vdots \\ & & \lambda^k & \binom{k}{1} \lambda^{k-1} \\ & & & \lambda^k \end{bmatrix}.$$

Indeed, λ^k can be written succinctly by repeatedly squaring λ . Similarly, $\binom{k}{i}$ is an integer so it can be obtained by repeatedly squaring 2 since it has polynomial size (in the size of k) when written in binary (since i is bounded by the dimension of the matrix).

The Orbit Problem

An *instance of the Orbit Problem*, or *Orbit instance* for short, is given by a square matrix $A \in \mathbb{Q}^{d \times d}$ and two vectors $x, y \in \mathbb{Q}^d$. 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. It was shown by Kannan and Lipton [KL80, KL86] that the Orbit problem is decidable.

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.

Representation of semialgebraic invariants

As it is the case in Examples 2 and Examples 3, many of the invariants we synthesise are comprised of an initial segment of the orbit $\{x, Ax, \dots, A^{k-1}x\}$ together with a semialgebraic invariant Q for the instance $(A, A^k x, y)$. Since the smallest such k may be exponential and polynomials describing Q might involve large exponents we choose to represent such an invariant by the pair (k, φ) where k is written in binary and φ is a succinct formula. More precisely, (k, φ) represents the semialgebraic set

$$\{x, Ax, \dots, A^{k-1}x\} \cup \{x : x \text{ satisfies } \varphi\}.$$

By abuse of notation, we will sometimes write (k, φ) to denote the above set. This representation is more succinct than a more classical representation given by a single boolean combination of polynomial inequalities using a standard representation. However, this does not affect complexity as stated in the following lemma.

We say that (k, Q) is an *eventual invariant* if $y \notin \{x, Ax, \dots, A^{k-1}x\}$ and Q is an invariant for the instance $(A, A^k x, y)$. Or equivalently if

- $A^k x \in Q := \{z : \varphi(z)\}$,
- $AQ \subseteq Q$,
- $y \notin \{x, Ax, \dots, A^{k-1}x\} \cup Q$.

One checks that any eventual invariant is invariant but the converse is not true. Indeed, one can build an invariant (k, Q) such that $AQ \not\subseteq Q$. However, if (k, Q) is invariant, then $(0, \{x, \dots, A^k x\} \cup Q)$ is trivially an eventual invariant, but its description can be exponential bigger, even using the succinct representation for semialgebraic sets.

Lemma 3 *For the representation defined above, the following two problems can be decided in polynomial space:*

- *decide if (k, Q) is an eventual invariant,*
- *assume (k, Q) is an eventual invariant, decide if $z \in (k, Q)$,*

Proof To check that (k, Q) is an eventual invariant, we need to check that:

- $A^k x \in Q$: we cannot write down $A^k x$ since k could be exponential but thanks to Lemma 2 we can express $\{A^k x\}$ succinctly and check if $\{A^k x\} \subseteq Q$.
- $AQ \subseteq Q$: this is an instance of Lemma 2.
- $y \notin \mathcal{P} := \{x, Ax, \dots, A^{k-1}x\} \cup Q$: since $A^k x \in Q$, it follows that the entire orbit of x is contained in $\mathcal{P}$. Thus $y \notin \mathcal{P}$ if and only if y is not in the orbit of x and $y \notin Q$. The former is decidable in polynomial time [KL80] and the latter in polynomial space using Lemma 2.

Now given (k, Q) an eventual invariant, and using the same remark as above, checking if $z \in (k, Q)$ is equivalent to checking if z is in the orbit of x and if $z \in Q$.

Note that these two complexity bounds match the complexity for the more classical representation. However, it is unclear if the same lemma holds for succinct invariants that are not eventual invariants. Indeed, checking that $AQ \subseteq (k, Q)$ (as opposed to $AQ \subseteq Q$) appears nontrivial. A consequence of our main result is that whenever an invariant exists, a succinct eventual invariant of small size always exists.

Throughout this paper we assume that semialgebraic sets are given by the representation (k, Q) , above, and the size of a semialgebraic set refers to the number of bits required for its representation.

3 Semialgebraic Invariants

Our main result is the following.

Theorem 1 *There exist PSPACE algorithms for the following two problems: given an Orbit instance,*

- *does it admit a semialgebraic invariant?*
- *does it admit a closed semialgebraic invariant?*

Furthermore,

- *both algorithms construct such an invariant when it exists, and the invariant produced has polynomial-size description,*
- *the invariants produced by the algorithms are always eventual invariants,*
- *there exists a closed semialgebraic invariant if and only if the target vector is not in the topological closure of the orbit.*

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.

The remainder of the paper is devoted to proving Theorem 1. To this end, let $\ell = (A, x, y)$ be a non-reachability Orbit instance in dimension d .

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 4 *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 . Moreover, the size of $Q^{-1}\mathcal{P}$ is at most the sum of the size of Q and of $\mathcal{P}$. Finally, if $\mathcal{P}$ is an eventual invariant then $Q^{-1}\mathcal{P}$ is also an eventual invariant.*

Proof First of all, $Q^{-1}\mathcal{P}$ is semialgebraic if and only if $\mathcal{P}$ is semialgebraic and we can build a succinct representation of $Q^{-1}\mathcal{P}$ from that of $\mathcal{P}$ using Lemma 2 since Q^{-1} is invertible. 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, the last point being clear. One easily checks that this transformation also preserves eventual invariants.

Thanks to Lemma 4, 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 [Cai00, CLZ00].

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 $k \in \{1, \dots, d(J)\}$ where $d(J)$ 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 let v_S denote the projected vector on $\mathbb{C}^S$. We extend it to matrices: for $A \in \mathbb{C}^d \times \mathbb{C}^d$, we define $A_S \in \mathbb{C}^S \times \mathbb{C}^S$. The notation $v_{J, > k}$ is for the projection on the coordinates of the Jordan

block J whose indices are greater than k . We let S^c denote the complement of S in $\{1, \dots, d\}$.

There are a few degenerate cases that 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 , x_J is 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 5 *The existence of (closed) semialgebraic eventual invariants for Orbit instances reduces in polynomial 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 $y_J \neq 0$, then consider $\mathcal{P} = (k, Q)$ where $\{z \in \mathbb{C}^d : z_J = 0\}$ is semialgebraic, which we claim is an eventual 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 $(A^n x)_J = 0$, i.e. $A^n x \in \mathcal{P}$. If $z_J = 0$ then $(Az)_J = 0$, i.e. $AQ \subseteq Q$. 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, let J be the dimensions corresponding to Jordan blocks associated with the eigenvalue 0, we have that $y_J = 0$. Consider the Orbit instance $\ell_{J^c} = (A_{J^c}, (A^d x)_{J^c}, y_{J^c})$. We claim that ℓ admits a semialgebraic eventual invariant if and only if ℓ_{J^c} does.

Let $\mathcal{P} = (k, Q)$ be an eventual semialgebraic invariant for ℓ and construct $Q' = \{z \in \mathbb{C}^{J^c} : (z, 0) \in Q\}$. We argue that $\mathcal{P}_{J^c} = (d, Q')$ is an eventual invariant for ℓ_{J^c} . Indeed, $A_{J^c} Q' \subseteq Q'$ since $AQ \subseteq Q$ and $A(z, 0) = (A_{J^c} z, 0)$ for any z . Then $A_{J^c}^k (A^d x)_{J^c} = (A^{k+d} x)_{J^c} \in Q'$ since $A^k \in Q$, Q is stable under A and $(A^{k+d} x)_J = 0$. Finally $y_{J^c} \notin \mathcal{P}_{J^c}$ because $y_J = 0$, so $y_{J^c} \in \mathcal{P}_{J^c}$ would imply $y \in \mathcal{P}$.

Conversely, let $\mathcal{P}_{J^c} = (k, Q')$ be a semialgebraic invariant for ℓ_{J^c} , construct $\mathcal{P} = (k, Q')$ where $Q' = \{(z, 0) \in \mathbb{C}^{J^c} \times \mathbb{C}^J : z \in Q'\}$. One easily checks that it is an eventual invariant for ℓ using a similar reasoning.

We reduced the existence of semialgebraic invariants from ℓ to ℓ_{J^c} , with the additional property that the matrix is invertible.
- If A contains a Jordan block J such that $x_J = 0$, we distinguish two cases.
 - If for some Jordan block J we have $x_J = 0$ and $y_J \neq 0$, then $\mathcal{P} = (0, \{z \in \mathbb{C}^d : z_J = 0\})$ is a semialgebraic eventual invariant for ℓ .
 - Otherwise, let J be the dimensions corresponding to Jordan blocks for which $x_J = y_J = 0$. Consider the Orbit instance $\ell_{J^c} = (A_{J^c}, x_{J^c}, y_{J^c})$, we claim that ℓ admits an eventual invariant if and only if ℓ_{J^c} does.

Let $\mathcal{P} = (k, Q)$ be a semialgebraic eventual invariant for ℓ , construct $\mathcal{P}_{J^c} = (k, Q')$ where $Q' = \{z \in \mathbb{C}^{J^c} : (z, 0) \in Q\}$. We easily see that $\mathcal{P}_{J^c}$ is a semialgebraic eventual invariant for ℓ_{J^c} .

Conversely, let $\mathcal{P}_{J^c} = (k, Q)$ be a semialgebraic eventual invariant for ℓ_{J^c} , construct $\mathcal{P} = (k, Q')$ where $Q' = \{(z, 0) \in \mathbb{C}^{J^c} \times \mathbb{C}^J : z \in Q\}$. We easily see that $\mathcal{P}$ is a semialgebraic invariant for ℓ .

We reduced the existence of semialgebraic eventual invariants from ℓ to ℓ_{J^c} , with the additional property that for each Jordan block J , $x_J \neq 0$.

- If A contains a non-diagonal Jordan block J such that $x_{J, >1} = 0$, we distinguish two cases.

- If for some non-diagonal Jordan block J we have that $x_{J, >1} = 0$ and $y_{J, >1} \neq 0$, then $\mathcal{P} = (0, \{z \in \mathbb{C}^d : z_{J, >1} = 0\})$ is a semialgebraic eventual invariant for ℓ .

- Otherwise, let J be the dimensions corresponding to non-diagonal Jordan blocks for which $x_{J, >1} = y_{J, >1} = 0$. Let $S = J^c \cup \bigcup_J (J, 1)$, *i.e.*, the dimensions outside J plus the first dimension of each block in J . Consider the Orbit instance $\ell_S = (A_S, x_S, y_S)$, we claim that ℓ admits a semialgebraic eventual invariant if and only if ℓ_S does.

Let $\mathcal{P} = (k, Q)$ be a semialgebraic invariant for ℓ , construct $\mathcal{P}_S = (k, Q')$ where $Q' = \{z \in \mathbb{C}^S : (z, 0) \in \mathcal{P}\}$. We easily see that $\mathcal{P}_S$ is a semialgebraic eventual invariant for ℓ_S .

Conversely, let $\mathcal{P}_S = (k, Q)$ be a semialgebraic invariant for ℓ_S , construct $\mathcal{P} = (k, Q')$ where $Q' = \{(z, 0) \in \mathbb{C}^S \times \mathbb{C}^{S^c} : z \in \mathcal{P}_S\}$. We easily see that $\mathcal{P}$ is a semialgebraic eventual 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 , $x_{J, >1} \neq 0$.

This concludes the proof. Note that in all constructions closed invariants were constructed, hence the proof also yields a reduction for the existence of closed invariants.

3.1 A useful lemma

We give a number theoretic lemma which we will need on several occasions. It concerns eigenvalues of A of modulus different than 1 and is an easy consequence of a separation bound by Mignotte [Mig82] asserting that for two roots $\alpha \neq \beta$ of a polynomial $P \in \mathbb{Z}[x]$, we have

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

where d and H are respectively the degree and height of the polynomial P .

Lemma 6 *Let λ be a nonzero eigenvalue of a rational matrix A . If $|\lambda| \neq 1$, then $\frac{1}{|\log(|\lambda|)|}$ is bounded by an exponential in the size of A .*

Proof (Sketch) Let p be the minimal polynomial of A , then p has degree at most d , the dimension, and height H exponential in the size of A . Let q be an integer polynomial such that the roots of q are the products of the roots of p . There are standard constructions using the resultant, one can build such a q in a way that only increases the degree and height by a polynomial factor. Since A is a real matrix, λ and $\bar{\lambda}$ are eigenvalues of A , thus roots of p , thus $|\lambda|^2 = \lambda\bar{\lambda}$ is a root of q . Without loss of generality, we can assume that 1 is a root of q , by considering $r(x) = q(x)(x-1)$, which incurs only a polynomial blowup in the degree and height. Finally, 1 and $|\lambda|^2$ are roots of r , thus by the Mignotte bound, and since $1 \neq |\lambda|^2$, $||\lambda|^2 - 1| > \sqrt{6}/M$ where $M = e^{(e+1)/2}G^{d-1}$ where e is polynomial in d and G is polynomial in H and d . If $|\lambda| > 1$, we then have, using $\log(1+x) \geq \frac{x}{x+1}$, that

$$(\log(|\lambda|^2))^{-1} \leq \frac{|\lambda|^2}{|\lambda|^2 - 1} \leq \sqrt{(eG)^2 6} e^{(e+1)/2} G^{d-1}$$

which is exponential in G and e and thus exponential in the size of A at most.

3.2 Some eigenvalue has modulus greater than one

Lemma 7 *Let $\ell = (J, x, y)$ be a non-reachability instance, where J is a Jordan block whose modulus is greater than 1 and $x \neq 0$, then there exists a closed semialgebraic eventual invariant for ℓ with size polynomial in the size of ℓ . This invariant can be computed in polynomial time.*

Proof Some coordinate of $(J^n x)_{n \in \mathbb{N}}$ 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 $x \neq 0$, let k such that $x_k \neq 0$ and $x_{>k} = 0$. For all $n \in \mathbb{N}$, we have $(J^n x)_k = \lambda^n x_k$, so $|(J^n x)_k|$ diverges to infinity. It follows that there exists $n_0 \in \mathbb{N}$ such that $|(J^{n_0} x)_k| > |y_k|$. Let $\mathcal{P} = (n_0, Q)$ where

$$Q = \{z \in \mathbb{C}^d : |z_k| \geq |(J^{n_0} x)_k| \text{ and } z_{>k} = 0\}.$$

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

To finish the proof we argue that $\mathcal{P}$ indeed has polynomial size. Indeed one can choose $n_0 = \left\lceil \frac{\log(|y_k|) - \log(|x_k|)}{\log(|\lambda|)} \right\rceil$, and using Lemma 6, we see that n_0 is bounded by an exponential in the size of ℓ and we can use Lemma 2 to write $(J^{n_0} x)_k$ succinctly.

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

3.3 Some eigenvalue has modulus smaller than one

Lemma 8 *Let $\ell = (J, x, y)$ be a non-reachability instance, where J is a Jordan block whose modulus is smaller than 1 and $x \neq 0$, then there exists a semialgebraic eventual invariant for ℓ . Further, if $y \neq 0$ then there exists a closed semialgebraic eventual invariant for ℓ . Moreover, these invariants have size polynomial in the size of ℓ and can be computed in polynomial time.*

Proof If $y = 0$ then $\mathcal{P} = (0, \{z \in \mathbb{C}^d : z \neq 0\})$ is clearly a semialgebraic eventual invariant, but it is not closed.

For the remainder of the proof we assume that $y \neq 0$ and construct a closed semialgebraic eventual invariant. The situation is similar to Lemma 7, 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 $x_k \neq 0$ and $x_{>k} = 0$, we may have that $y_k = 0$, implying that $((J^n x)_k)_{n \in \mathbb{N}}$ does not become smaller than y_k . Working on another dimension entails giving up the following diagonal behaviour: $(J^n x)_k = \lambda^n x_k$, making it hard to find a stable set under J . To overcome this problem, the invariant we define depends upon all the coordinates of J .

We let d denote the dimension of J . We have that $(J^n x)_{n \in \mathbb{N}}$ converges to 0. It follows that there exists $n_0 \in \mathbb{N}$ such that for each dimension k of J , i.e., for $k \in \{1, \dots, d\}$, we have $|(J^{n_0} x)_k| \leq (1 - |\lambda|)^k \cdot \|y\|_\infty$. Let $\mathcal{P} = (n_0, Q)$ where

$$Q = \{z \in \mathbb{C}^d : \forall k \in \{1, \dots, d\}, |z_k| \leq (1 - |\lambda|)^k \cdot \|y\|_\infty\}.$$

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

$$\begin{aligned} |(Jz)_k| &= |\lambda z_k + z_{k+1}| \\ &\leq |\lambda| |z_k| + |z_{k+1}| \\ &\leq |\lambda| (1 - |\lambda|)^k \cdot \|y\|_\infty + (1 - |\lambda|)^{k+1} \cdot \|y\|_\infty \\ &= (|\lambda| + (1 - |\lambda|))(1 - |\lambda|)^k \cdot \|y\|_\infty \\ &= (1 - |\lambda|)^k \cdot \|y\|_\infty. \end{aligned}$$

The case $k = d$ is similar but easier.

It remains to see that $\mathcal{P}$ has a polynomial representation. We show that n_0 can be chosen exponential in the size of ℓ . We will consider the case where J is not diagonal, which is a bit harder and easily adapted to the diagonal case. Note that

$$|(J^n x)_k| = |\lambda^n x_k + n\lambda^{n-1}x_{k+1} + \dots + \binom{n}{d-k}\lambda^{n-d+k}x_d| \leq d|\lambda|^{n-d}n^d\|x\|_\infty.$$

It follows that n_0 can be chosen to be

$$\left\lceil \frac{\log(\|y\|_\infty) - \log(d) - \log(\|x\|_\infty)}{\log(|\lambda|)} + d \right\rceil$$

which is exponential in the size of ℓ thanks to Lemma 6.

3.4 Some eigenvalue has modulus one and corresponds to a non-diagonal Jordan block

Lemma 9 *Let $\ell = (J, x, y)$ be a non-reachability instance, where J is a non-diagonal Jordan block whose modulus is equal to 1 and $x_{>1} \neq 0$, then there exists a closed semialgebraic eventual invariant for ℓ of size polynomial in the size of ℓ . This invariant can be computed in polynomial time.*

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

Proof Let k such that $x_k \neq 0$ and $x_{>k} = 0$, we have $k \geq 1$ and

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

so $(|(J^n x)_{k-1}|)_{n \in \mathbb{N}}$ diverges to infinity since $|\lambda| = 1$. It follows that there exists $n_0 \in \mathbb{N}$ such that $|(J^{n_0} x)_{k-1}| > |y_{k-1}|$. Without loss of generality we assume $n_0 \geq -\frac{\langle \lambda x_{k-1}, x_k \rangle}{|x_k|^2}$. 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 $\text{Re}(u\bar{v})$.

This quantity will appear later; note that it only depends on x and J . Let $\mathcal{P} = (n_0, Q)$ where

$$Q = \{z \in \mathbb{C}^d : |z_{k-1}| \geq |(J^{n_0}x)_{k-1}|, \text{ and } \langle \lambda z_{k-1}, z_k \rangle \geq 0, \text{ and } z_{>k} = 0\}.$$

We argue that $\mathcal{P}$ is a semialgebraic eventual invariant for ℓ . It is a semialgebraic set: the condition $\langle \lambda z_{k-1}, z_k \rangle \geq 0$ 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 Q under J .

First, $J^{n_0}x \in Q$. Indeed, using $|\lambda| = 1$ and the assumption on n_0 ,

$$\begin{aligned} \langle \lambda(J^{n_0}x)_{k-1}, (J^{n_0}x)_k \rangle &= \langle \lambda \cdot (\lambda^{n_0}x_{k-1} + n_0\lambda^{n_0-1}x_k), \lambda^{n_0}x_k \rangle \\ &= |\lambda^{n_0}|^2 \langle \lambda x_{k-1}, x_k \rangle + n_0 |\lambda^{n_0}x_k|^2 \\ &= \langle \lambda x_{k-1}, x_k \rangle + n_0 |x_k|^2 \\ &\geq 0. \end{aligned}$$

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

$$\begin{aligned} |(Jz)_{k-1}|^2 &= |\lambda z_{k-1} + z_k|^2 \\ &= |z_{k-1}|^2 + 2\langle \lambda z_{k-1}, z_k \rangle + |z_k|^2 \\ &\geq |z_{k-1}|^2 \\ &\geq |(J^{n_0}x)_{k-1}|^2, \end{aligned}$$

and

$$\begin{aligned} \langle \lambda(Jz)_{k-1}, (Jz)_k \rangle &= \langle \lambda(\lambda z_{k-1} + z_k), \lambda z_k \rangle \\ &= |\lambda|^2 \langle \lambda z_{k-1} + z_k, z_k \rangle \\ &= \langle \lambda z_{k-1}, z_k \rangle + |z_k|^2 \\ &\geq 0. \end{aligned}$$

Hence $Jz \in Q$, and $\mathcal{P}$ is a semialgebraic eventual invariant for ℓ .

To conclude we show that $\mathcal{P}$ has size polynomial in the size of ℓ . Indeed, it is enough to have $n_0 \geq \left\lceil \frac{|y_{k-1}| - |x_{k-1}|}{|x_k|} \right\rceil$ and we can use Lemma 2.

Example 4 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

$$A^n x = \begin{bmatrix} e^{in\theta} + ne^{i(n-1)\theta} & e^{in\theta} \end{bmatrix},$$

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 :

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

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 $\langle e^{i\theta}z_1, z_2 \rangle \geq 0$.

3.5 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 is not reachable and yet no closed semialgebraic invariant exists. (Recall Example 1 from the Introduction.) Write $\mathcal{O} = \{A^n x : n \in \mathbb{N}\}$ for the orbit of x under A , and $\overline{\mathcal{O}}$ for its topological closure. Using results from Diophantine approximation and algebraic number theory, we show that $\overline{\mathcal{O}}$ 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 $y \notin \overline{\mathcal{O}}$.

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., [Dri98, Chapter 4]). However to establish the lemma it suffices to note the following two properties of the dimension. First, for any semialgebraic set $X \subseteq \mathbb{R}^n$ we have $\dim(X) = \dim(\overline{X})$ [Dri98, Chapter 4, Theorem 1.8]. Secondly, if $X \subseteq Y$ are semialgebraic subsets of $\mathbb{R}^n$ that have the same dimension, then X has non-empty interior in Y [Dri98, 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 $A = \text{diag}(\lambda_1, \dots, \lambda_d)$ is a diagonal $d \times d$ matrix with entries $\lambda_1, \dots, \lambda_d \in \mathbb{C}$ all having modulus one. Then*

- *The set $\overline{\mathcal{O}}$ is a semialgebraic set that is computable from ℓ in polynomial space and has polynomial size using the succinct representation.*
- *Any semialgebraic 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

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

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

$$T_A = \{\text{diag}(\mu_1, \dots, \mu_d) : \mu \in \mathbb{T}^d \text{ and } \forall v \in L_A, (\mu_1^{v_1} \cdots \mu_d^{v_d} = 1)\}$$

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 [Cas65], it is shown in [OW14, Proposition 3.5] that

$$\{A^n : n \in \mathbb{N}\}$$

is a dense subset of T_A . This immediately gives

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

We now show that $\overline{\mathcal{O}}$ is semialgebraic. Observe that L_A is finitely generated, being a subgroup of the free finitely generated abelian group $\mathbb{Z}^d$. Moreover, if $B \subseteq L_A$ is a basis of L_A then we can write

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

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

From an upper bound on the length of B due to Masser [Mas88], it can be shown that one can compute a basis for L_A of polynomial size, in polynomial space in the description of A (see [OW14, Corollary 3.3]) and thereby compute a description of T_A as a semialgebraic set, also in polynomial space in the description of A . More precisely, the resulting basis B has at most d elements and each vector in it has polynomial size (when writing integers in binary). Using the succinct representation, we can thus write $\mu_1^{v_1} \cdots \mu_d^{v_d}$ using polynomial size only by repeated squaring. Doing so for each of the (at most d) vectors of the basis yields a polynomial size formula.

Now we move to the second item in the statement of the lemma. Let $\mathcal{P}$ be a semialgebraic 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 semialgebraic 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 $\mathcal{P} = \overline{\mathcal{O}}$.

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 $A^{-1} = \text{diag}(\lambda_1^{-1}, \dots, \lambda_d^{-1})$ 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., $\lambda_1, \dots, \lambda_d$ satisfy exactly the same multiplicative relations as $\lambda_1^{-1}, \dots, \lambda_d^{-1}$), and hence also $T_A = T_{A^{-1}}$.

Now we claim that the following chain of equalities holds:

$$\overline{\{A^{-n}z : n \in \mathbb{N}\}} = \{Mz : M \in T_{A^{-1}}\} \quad (2)$$

$$= \{Mz : M \in T_A\} \quad (3)$$

$$= \{Mx : M \in T_A\} \quad (4)$$

$$= \overline{\mathcal{O}} = \overline{\mathcal{P}}.$$

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_0x 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 1 *Let the Orbit instance ℓ be as described in Lemma 11. Then ℓ admits a semialgebraic eventual invariant if and only if $y \notin \overline{\mathcal{O}}$. Furthermore, when it exists, such an eventual can be computed in polynomial space and has size polynomial in the size of ℓ .*

Proof If $y \notin \overline{\mathcal{O}}$, then $\overline{\mathcal{O}}$ is a semialgebraic invariant for ℓ by the first item in Lemma 11 and thus $(0, \overline{\mathcal{O}})$ is an eventual invariant. Conversely, if there exists a semialgebraic invariant $\mathcal{P}$ for ℓ , then $\overline{\mathcal{O}} \subseteq \mathcal{P}$ by the second item in Lemma 11, implying that $y \notin \overline{\mathcal{O}}$.

3.6 Proof of Theorem 1

We now draw together the results of the previous sections to prove our main result, Theorem 1, 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.

Lemma 12 *Let $\ell = (A, x, y)$ be a non-trivial Orbit instance in Jordan normal form. Assume that there exists a Jordan block J such that either the eigenvalue has modulus different from 1, or J is a non-diagonal Jordan block. Then there exists n_0 at most exponential such that if $A^n x = y$, then $n < n_0$.*

Proof Clear by looking at the modulus of $A^n x$.

We first consider the case of (non necessarily closed) semialgebraic invariants, and divide into four cases.

1. Either some eigenvalue of A has modulus greater than 1.

2. Or some eigenvalue of A has modulus smaller than 1.
3. Or some eigenvalue of A has modulus 1 and a non-diagonal Jordan block.
4. Otherwise, all eigenvalues have modulus 1 and the matrix is diagonalisable, so thanks to Corollary 1 there exists a semialgebraic invariant if and only if the topological closure of the orbit $\overline{\mathcal{O}}$ is such an invariant, which holds if and only if the closure does not contain y .

In the first three cases, let J be the corresponding Jordan block. Thanks to Lemma 12, we can see that $(J, J^{n_0}x_J, y_J)$ is a non-reachability Orbit instance. Now thanks to either Lemma 7, Lemma 8, or Lemma 9, we obtain a semi-algebraic invariant, which then easily induces a semialgebraic invariant for ℓ .

In all cases, when an invariant exists we can compute an eventual invariant of polynomial size, and do so in polynomial space (and in fact polynomial time, except for the last case).

We now consider closed semialgebraic invariants. Lemma 7, Lemma 9, and Corollary 1 construct closed semialgebraic invariants, but Lemma 8 applies only if the Jordan block J whose eigenvalue is smaller than 1 satisfies $y_J \neq 0$. Hence we have to change the case distinction above. The first three cases are identical, the last case becomes: for all Jordan blocks J , the associated eigenvalue is either

- of modulus 1 and the block is diagonalisable,
- of modulus smaller than 1 and $y_J = 0$.

In the last case, we conclude thanks to the following lemma.

Lemma 13 *Let $\ell = (A, x, y)$ be a non-trivial Orbit instance in Jordan normal form. Assume that for each Jordan block J , either the eigenvalue has modulus 1 and J is diagonalisable, or the eigenvalue has modulus smaller than 1 and $y_J = 0$. Then there exists a closed semialgebraic invariant for ℓ if and only if $y \notin \overline{\mathcal{O}}$.*

Proof Let U denote the dimensions corresponding to eigenvalues of modulus 1. Observe that y is in the closure of the orbit of A under x if and only if y_U is in the closure of the orbit of A_U under x_U , because the action of $A_{\overline{U}}$ is to converge towards $y_{\overline{U}} = 0$ since the associated eigenvalues have modulus smaller than 1. Thanks to Corollary 1, this implies the claim.

Again, in all cases, when a closed invariant exists we can compute a closed eventual invariant of polynomial size, and do so in polynomial space.

Thus we obtain an effective characterisation of the class of Orbit instances for which there exists a semialgebraic invariant, and similarly for a closed semialgebraic invariant. Note that non-reachability Orbit instances for which there do not exist semialgebraic invariants are extremely sparse. In those cases in which there exists an invariant we have shown how to compute such an invariant of polynomial size in polynomial space.

Moreover, we obtain the completeness of the class of closed semialgebraic invariants: there exists a closed semialgebraic invariant if and only if the target vector is not in the topological closure of the orbit.

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 algorithm is the problem of computing a basis of the group of all multiplicative relations among a given collection of algebraic numbers $\alpha_1, \dots, \alpha_d$, 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 $\alpha_1^{n_1} \dots \alpha_d^{n_d} \stackrel{?}{=} 1$. Assuming that $\alpha_1, \dots, \alpha_d$ are represented as elements of an explicitly given finite-dimensional algebra K over $\mathbb{Q}$, Ge [Ge93] gave a polynomial-time algorithm for verifying multiplicative relations. In our setting, however, where $\alpha_1, \dots, \alpha_d$ are roots of the characteristic polynomial of matrix A , the dimension of K may be exponential in d . Note that in order to obtain an invariant of polynomial size, we had to introduce a succinct representation and the notion of eventual invariant. This representation is succinct in two orthogonal ways: we need to encode a prefix of the orbit succinctly because it can be exponentially long, and we need succinct representation of the semialgebraic “eventual” part. succinctly encoding the prefix of the orbit seems necessary in virtually all cases, whereas succinctly encoding the semialgebraic part is necessary because encoding the relations among the eigenvalues can involve very high exponents.

Acknowledgements We would like to thank the reviewers for their very detailed and helpful comments, in particular pointing out a flaw in the complexity analysis in the first version of the paper.

References

- [BCR98] Jacek Bochnak, Michel Coste, and Marie-Françoise Roy. *Real Algebraic Geometry*, volume 36 of *A Series of Modern Surveys in Mathematics*. Springer-Verlag Berlin Heidelberg, 1998.

- [Cai00] Jin-Yi Cai. Computing Jordan normal forms exactly for commuting matrices in polynomial time. Technical report, SUNY at Buffalo, 2000.
- [Cas65] John W. S. Cassels. *An introduction to Diophantine approximation*. Cambridge University Press, 1965.
- [CH78] Patrick Cousot and Nicolas Halbwachs. Automatic discovery of linear restraints among variables of a program. In *POPL*, pages 84–96. ACM Press, 1978.
- [CLZ00] 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.
- [Col07] Michael Colón. Polynomial approximations of the relational semantics of imperative programs. *Science of Computer Programming*, 64(1):76–96, 2007.
- [Dri98] Laurentius Petrus Dignus van den Dries. *Tame Topology and O-minimal Structures*. London Mathematical Society Lecture Note Series. Cambridge University Press, May 1998.
- [Ge93] Guoqiang Ge. Testing equalities of multiplicative representations in polynomial time. In *SFCS*, pages 422–426. IEEE Computer Society, 1993.
- [Har69] Michael A. Harrison. *Lectures on linear sequential machines*. New York-Londres, Academic Press, 1969.
- [KL80] Ravindran Kannan and Richard J. Lipton. The Orbit Problem is decidable. In *STOC*, pages 252–261, 1980.
- [KL86] Ravindran Kannan and Richard J. Lipton. Polynomial-time algorithm for the Orbit Problem. *Journal of the ACM*, 33(4):808–821, 1986.
- [Mas88] David W. Masser. Linear relations on algebraic groups. In Alan Baker, editor, *New Advances in Transcendence Theory*, pages 248–262. Cambridge University Press, 1988.
- [Mig82] Maurice Mignotte. Some useful bounds. In *Computer Algebra*, volume 4 of *Computing Supplementum 4*. Springer Vienna, 1982.
- [MS04a] Markus Müller-Olm and Helmut Seidl. A note on Karr’s algorithm. In *ICALP*, volume 3142 of *Lecture Notes in Computer Science*, pages 1016–1028. Springer, 2004.
- [MS04b] Markus Müller-Olm and Helmut Seidl. Precise interprocedural analysis through linear algebra. In *POPL*, pages 330–341. ACM, 2004.
- [OW14] Joël Ouaknine and James Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In *ICALP*, pages 330–341, 2014.
- [Tao08] Terence Tao. *Structure and Randomness*. AMS, 2008.