



On the Monniaux Problem in Abstract Interpretation

NATHANAEL FIJALKOW, University of Bordeaux, CNRS, LaBRI, Bordeaux, France

ENGEL LEFAUCHEUX, Max-Planck-Institute for Software Systems, Saarbrücken, Germany

PIERRE OHLMANN, Institut de Recherche en Informatique Fondamentale, IRIF, Paris, France

JOËL OUAKNINE, Max-Planck-Institute for Software Systems, Saarbrücken, Germany and Computer Science, Oxford University, Oxford, United Kingdom of Great Britain and Northern Ireland

AMAURY POULY, Institut de Recherche en Informatique Fondamentale, IRIF, Paris, France

JAMES WORRELL, Department of Computer Science, University of Oxford, Oxford, United Kingdom of Great Britain and Northern Ireland

The Monniaux Problem in abstract interpretation asks, roughly speaking, whether the following question is decidable: Given a program P , a safety (e.g., non-reachability) specification φ , and an abstract domain of invariants \mathcal{D} , does there exist an inductive invariant \mathcal{I} in \mathcal{D} guaranteeing that program P meets its specification φ ? The Monniaux Problem is of course parameterised by the classes of programs and invariant domains that one considers. In this article, we show that the Monniaux Problem is undecidable for unguarded affine programs and semilinear invariants (unions of polyhedra). Moreover, we show that decidability is recovered in the important special case of simple linear loops.

CCS Concepts: • **Theory of computation** → **Invariants; Program verification**;

Additional Key Words and Phrases: Abstract interpretation, invariants, dynamical systems, Monniaux problem

ACM Reference Format:

Nathanael Fijalkow, Engel Lefauchaux, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2025. On the Monniaux Problem in Abstract Interpretation. *J. ACM* 72, 2, Article 11 (March 2025), 51 pages. <https://doi.org/10.1145/3704632>

1 Introduction

Invariants are one of the most fundamental and useful notions in the quantitative sciences, appearing in a wide range of contexts, from gauge theory, dynamical systems, and control theory

Joël Ouaknine was supported by ERC grant AVS-ISS (648701) and DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>). Joël Ouaknine is also affiliated with Keble College, Oxford as emmy.network Fellow. James Worrell was supported by EPSRC Fellowships EP/N008197/1 and EP/X033813/1.

An extended abstract appeared in the 26th International Symposium on Static Analysis (SAS 2019) [12].

Authors' Contact Information: Nathanael Fijalkow, University of Bordeaux, CNRS, LaBRI, Bordeaux, France; e-mail: nathanael.fijalkow@gmail.com; Engel Lefauchaux, Max-Planck-Institute for Software Systems, Saarbrücken, Saarland, Germany; e-mail: elefauch@mpi-sws.org; Pierre Ohlmann, Institut de Recherche en Informatique Fondamentale, IRIF, Paris, Île-de-France, France; e-mail: ohlmann@irif.fr; Joël Ouaknine, Max-Planck-Institute for Software Systems, Saarbrücken, Saarland, Germany and Computer Science, Oxford University, Oxford, Oxfordshire, United Kingdom of Great Britain and Northern Ireland; e-mail: joel@mpi-sws.org; Amaury Pouly, Institut de Recherche en Informatique Fondamentale, IRIF, Paris, Île-de-France, France; e-mail: amaury.pouly@irif.fr; James Worrell, Department of Computer Science, University of Oxford, Oxford, United Kingdom of Great Britain and Northern Ireland; e-mail: james.ben.worrell@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 0004-5411/2025/03-ART11

<https://doi.org/10.1145/3704632>

in physics, mathematics, and engineering to program verification, static analysis, abstract interpretation, and programming language semantics (among others) in computer science. In spite of decades of scientific work and progress, automated invariant synthesis remains a topic of active research, especially in the fields of program analysis and abstract interpretation, and plays a central role in methods and tools seeking to establish correctness properties of computer programs; see, e.g., Reference [23], and particularly Section 8 therein.

The focus of the present article is the **Monniaux Problem** on the decidability of the existence of separating invariants, which was formulated by David Monniaux in Reference [28] and also raised by him in a series of personal communications with various members of the theoretical computer science community over the past five years or so. There are in fact a multitude of versions of the Monniaux Problem—indeed, it would be more appropriate to speak of a *class* of problems rather than a single question—but at a high level the formulation below is one of the most general:

Consider a program P operating over some numerical domain (such as the integers or rationals), and assume that P has an underlying finite control-flow graph over the set of nodes $Q = \{q_1, \dots, q_r\}$. Let us assume that P makes use of d numerical variables, and each transition $q \xrightarrow{t} q'$ comprises a function $f_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ as well as a guard $g_t \subseteq \mathbb{R}^d$. Let $x, y \in \mathbb{Q}^d$ be two points in the ambient space. By way of intuition and motivation, we are interested in the reachability problem as to whether, starting in location q_1 with variables having valuation x , it is possible to reach location q_r with variables having valuation y , by following the available transitions and under the obvious interpretation of the various functions and guards. Unfortunately, in most settings this problem is well-known to be undecidable.

A collection $\{\mathcal{I}_q \mid q \in Q\}$ is called an (inductive¹) *invariant* if for each transition $q \xrightarrow{t} q'$, we have that $f_t(\mathcal{I}_q \cap g_t) \subseteq \mathcal{I}_{q'}$. If it additionally satisfies that $x \in \mathcal{I}_{q_1}$ and $y \notin \mathcal{I}_{q_r}$, then it is a *separating invariant* for program P . Clearly, the existence of a separating invariant constitutes a proof of non-reachability for P with the given x and y .

Let $\mathcal{D} \subseteq 2^{\mathbb{R}^d}$ be an “abstract domain” for P , i.e., a collection of subsets of \mathbb{R}^d . For example, \mathcal{D} could be the collection of all convex polyhedra in \mathbb{R}^d , or the collection of all closed semi-algebraic sets in \mathbb{R}^d , and so on.

The Monniaux Problem can now be formulated as a decision question: Is it possible to adorn each control location q with an element $\mathcal{I}_q \in \mathcal{D}$ such that the collection $\{\mathcal{I}_q \mid q \in Q\}$ forms a separating invariant?

Associated with this decision problem, in positive instances, one is also potentially interested in the synthesis problem, i.e., the matter of algorithmically producing a suitable separating invariant $\{\mathcal{I}_q : q \in Q\}$.

The Monniaux Problem is therefore parameterised by a number of items, key of which are (i) the abstract domain \mathcal{D} under consideration and (ii) the kind of functions and guards allowed in transitions.

Our main interest in this article lies in the *decidability* of the existence of separating invariants for various instances of the Monniaux Problem. We give below a cursory cross-sectional survey of existing work and results in this direction.

Arguably the earliest positive result in this area is due to Karr, who showed that strongest affine invariants (conjunctions of affine equalities) for affine programs (no guards, and all transition functions are given by affine expressions) could be computed algorithmically [22]. Note that the ability to synthesise strongest (i.e., smallest with respect to set inclusion) invariants immediately entails

¹In the remainder of this article, the term “invariant” shall always refer to the inductive kind.

the decidability of the Monniaux Problem instance, since the existence of *some* separating invariant is clearly equivalent to whether the *strongest* invariant is separating. Müller-Olm and Seidl later extended this work on affine programs to include the computation of strongest polynomial invariants of fixed degree [30], and a randomised algorithm for discovering affine relations was proposed by Gulwani and Necula [31]. In Reference [15], the least inductive invariant is computed by policy iteration for some families of abstract domains. More recently, Hrushovski et al. showed how to compute a basis for *all* polynomial relations at every location of a given affine program [19].

The approaches described above all compute invariants consisting exclusively of conjunctions of *equality* relations. By contrast, an early and highly influential paper by Cousot and Halbwachs considers the domain of convex closed polyhedra [9] for programs having polynomial transition functions and guards. While no decidability results appear in that paper, much further work was devoted to the development of restricted polyhedral domains for which theoretical guarantees could be obtained, leading (among others) to the *octagon domain* of Miné [27], the *octahedron domain* of Clarisó and Cortadella [6], and the *template polyhedra* of Sankaranarayanan et al. [33]. In fact, as observed by Monniaux [28], if one considers a domain of convex polynomial templates having a *uniformly bounded* number of faces (therefore subsuming in particular the domains just described), then for any class of programs with polynomial transition relations and guards, the existence of separating invariants becomes decidable, as the problem can equivalently be phrased in the first-order theory of the reals.

One of the central motivating questions for the Monniaux Problem is whether one can always compute separating invariants for the full domain of polyhedra. Unfortunately, on this matter very little is known at present. In recent work, Monniaux showed undecidability for the domain of convex polyhedra and the class of programs having affine transition functions and polynomial guards [28]. One of the main results of the present article is to show undecidability for the domain of *semilinear sets*² and the class of affine programs (without any guards)—in fact, affine programs with only a single control location and two transitions:

THEOREM 1.1. *Let $A, B \in \mathbb{Q}^{d \times d}$ be two rational square matrices of dimension d , and let x, y be two points in \mathbb{Q}^d . Then, the existence of a semilinear set $I \subseteq \mathbb{R}^d$ having the following properties:*

- (1) $x \in I$;
- (2) $AI \subseteq I$ and $BI \subseteq I$; and
- (3) $y \notin I$

is an undecidable problem.

It is worth pointing out that the theorem remains valid even for a fixed d (our proof shows undecidability for $d = 96$, but this value could be improved). If, moreover, one requires I to be topologically closed, one can lower d to having fixed value 24. Finally, an examination of the proof reveals that the theorem also holds for the domain of semi-algebraic sets, and in fact for any domain of o-minimal sets in the sense of Reference [1]. The proof also carries through whether one considers the domain of semilinear sets having rational, algebraic, or real coordinates.

Although the above is a negative (undecidability) result, it should be viewed in a positive light; as Monniaux writes in Reference [28], “We started this work hoping to vindicate forty years of research on heuristics by showing that the existence of polyhedral inductive separating invariants in a system with transitions in linear arithmetic (integer or rational) is undecidable.” Theorem 1.1 shows that, at least as regards non-convex invariants, the development and use of heuristics is indeed

²A semilinear set consists of a finite union of polyhedra or equivalently is defined as the solution set of a Boolean combination of linear inequalities.

vindicated and will continue to remain essential. Related questions of *completeness* of given abstraction scheme have also been examined by Giacobazzi et al. in References [17, 18]. We refer to Reference [29] for a recent and personal point of view on the Monniaux problem, by Monniaux himself.

It is important to note that our undecidability result requires at least *two* transitions. In fact, much research work has been expended on the class of simple *affine* loops, i.e., one-location programs equipped with a single self-transition. In terms of invariants, Fijalkow et al. establish in References [13, 14] the decidability of the existence of *semi-algebraic* separating invariants and specifically state the question of the existence of separating *semilinear* invariants as an open problem. Almagor et al. extend this line of work in Reference [1] to more complex targets (in lieu of the point y) and richer classes of invariants. The second main result of the present article is to settle the open question of References [13, 14] in the affirmative:

THEOREM 1.2. *Let $A \in \mathbb{Q}^{d \times d}$ be a rational square matrix of dimension d , and let $x, y \in \mathbb{Q}^d$ be two points in \mathbb{Q}^d . It is decidable whether there exists a closed semilinear set $\mathcal{I} \subseteq \mathbb{R}^d$ having algebraic coordinates such that:*

- (1) $x \in \mathcal{I}$;
- (2) $A\mathcal{I} \subseteq \mathcal{I}$; and
- (3) $y \notin \mathcal{I}$.

The proof shows that, in fixed dimension d , the decision procedure runs in polynomial time. It is worth noting that one also has decidability if A , x , and y are taken to have real-algebraic (rather than rational) entries.

Let us conclude this section by briefly commenting on the important issue of *convexity*. At its inception, abstract interpretation had a marked preference for domains of *convex* invariants, of which the interval domain, the octagon domain, and of course the domain of convex polyhedra are prime examples. Convexity confers several distinct advantages, including simplicity of representation, algorithmic tractability and scalability, ease of implementation, and better termination heuristics (such as the use of widening). The central drawback of convexity, however, is its poor expressive power. This has been noted time and again: “*convex polyhedra [...] are insufficient for expressing certain invariants, and what is often needed is a disjunction of convex polyhedra.*” [2]; “*the ability to express non-convex properties is sometimes required in order to achieve a precise analysis of some numerical properties*” [16]. Abstract interpretation can accommodate non-convexity either by introducing disjunctions (see, e.g., Reference [2] and references therein) or via the development of special-purpose domains of non-convex invariants such as *donut domains* [16]. The technology, data structures, algorithms, and heuristics supporting the use of disjunctions in the leading abstract-interpretation tool *ASTRÉE* are presented in great detail in Reference [8]. In the world of software verification, where predicate abstraction is the dominant paradigm, disjunctions—and hence non-convexity—are nowadays native features of the landscape.

It is important to note that the two main results presented in this article, Theorems 1.1 and 1.2, have only been proven for families of invariants that are not necessarily convex. The Monniaux Problem restricted to families of *convex* invariants remains open and challenging.

2 Preliminaries

We start with necessary definitions and notations.

2.1 Real and Complex Numbers

We will mostly work in the field $\mathbb{A} \subseteq \mathbb{C}$ of algebraic numbers, that is, roots of polynomials with coefficients in \mathbb{Z} . It is possible to represent and manipulate algebraic numbers effectively,

by storing their minimal polynomial and a sufficiently precise numerical approximation. An excellent reference in computational algebraic number theory is Reference [7]. All standard algebraic operations such as sums, products, root-finding of polynomials, or computing Jordan normal forms of matrices with algebraic entries can be performed effectively.

The set of complex numbers is \mathbb{C} , and for a complex number z its modulus is $|z|$, its real part is $\text{Re}(z)$, its imaginary part is $\text{Im}(z)$ and its conjugate is z^* . Let \mathbb{C}^* denote the set of non-zero complex numbers. We write S^1 for the complex unit circle, i.e., the set of complex numbers of modulus 1. We let \mathbb{U} denote the set of roots of unity, i.e., complex numbers $z \in S^1$ such that $z^n = 1$ for some $n \in \mathbb{N}$. We write $\text{Diag}(\lambda_1, \dots, \lambda_d)$ for

$$\begin{bmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \ddots & \\ & & & \lambda_d \end{bmatrix}.$$

When working in \mathbb{C}^d , the norm of a vector z is $\|z\|$, defined as the maximum of the modulus of each complex component z_i for i in $\{1, \dots, d\}$, where z_i is the i th component of vector z . For $\varepsilon > 0$ and z in \mathbb{C}^d , we write $B(z, \varepsilon)$ for the open ball centered in z of radius ε . The topological closure of a set $\mathcal{I} \subseteq \mathbb{C}^d$ is $\overline{\mathcal{I}}$, its interior \mathcal{I}° , and its boundary $\partial\mathcal{I}$, defined as $\overline{\mathcal{I}} \cap \overline{\mathbb{C}^d} \setminus \mathcal{I}$.

For semilinear sets, defined below, we will need to view complex sets as real sets in twice the dimension. We introduce some notations to clarify this part. For any $d \in \mathbb{N}$, we introduce the mapping

$$(\cdot)_{\mathbb{R}} : \mathbb{C}^d \rightarrow \mathbb{R}^{2d} \\ (x_1, \dots, x_d) \mapsto (\text{Re}(x_1), \text{Im}(x_1), \dots, \text{Re}(x_d), \text{Im}(x_d)).$$

We naturally extend this mapping to matrices so if $A \in \mathbb{C}^{d \times d}$, then $A_{\mathbb{R}} \in \mathbb{R}^{2d \times 2d}$ is such that for all $x \in \mathbb{C}^d$, $A_{\mathbb{R}}x_{\mathbb{R}} = (Ax)_{\mathbb{R}}$. Furthermore, the following relation exists between the determinant of A and $A_{\mathbb{R}}$:

LEMMA 2.1. For any $A \in \mathbb{C}^{d \times d}$, $\det(A_{\mathbb{R}}) = \det(A)\det(A)^*$.

PROOF. Recall that the determinant is invariant when corresponding rows and columns undergo permutation. It is not hard to check that $A_{\mathbb{R}}$ is the block matrix $(R(A_{ij}))_{i,j}$ where

$$R(z) = \begin{bmatrix} \text{Re}(z) & -\text{Im}(z) \\ \text{Im}(z) & \text{Re}(z) \end{bmatrix}$$

for all $z \in \mathbb{C}$. Therefore, by permuting rows, we can write

$$A_{\mathbb{R}} = P^{-1} \begin{bmatrix} \text{Re}(A) & -\text{Im}(A) \\ \text{Im}(A) & \text{Re}(A) \end{bmatrix} P,$$

where P has determinant 1 and $\text{Re}(A)$ (respectively, $\text{Im}(A)$) is the matrix whose entries are the real (respectively, imaginary) parts of the entries of A . Now, note that for any matrices $X, Y \in \mathbb{R}^{d \times d}$,

$$\begin{bmatrix} X & -Y \\ Y & X \end{bmatrix} = \begin{bmatrix} -iI_d & iI_d \\ I_d & I_d \end{bmatrix} \begin{bmatrix} X - iY & 0 \\ 0 & X + iY \end{bmatrix} \begin{bmatrix} \frac{i}{2}I_d & \frac{1}{2}I_d \\ -\frac{i}{2}I_d & \frac{1}{2}I_d \end{bmatrix},$$

where I_d is the identity matrix of dimension d .

Therefore, $\det(A_{\mathbb{R}}) = \det(\text{Re}(A) - i\text{Im}(A)) \det(\text{Re}(A) + i\text{Im}(A)) = \det(A^*) \det(A) = \det(A)^* \det(A)$.

□

2.2 Linear Dynamical Systems

A dynamical system is given by a set of functions $f_t : \mathbb{R}^d \rightarrow \mathbb{R}^d$ for $t \in [1, k]$. Let x be an initial vector, the set of *reachable points* from x is the smallest subset R of \mathbb{R}^d containing x and closed under the functions f_t : if $z \in R$, then $f_t(z) \in R$. If there is a single function ($k = 1$), then the set of reachable points from x is called the *orbit* of x under f . We say that y is reachable from x if y belongs to the set of reachable points from x . The reachability problem is the following decision problem: Given a (dynamical) system $S = (x, \{f_t : t \in [1, k]\}, y)$, determine whether y is reachable from x .

We are in this article interested in linear dynamical systems, where the functions f_t are linear: f_t is induced by a square matrix $A_t \in \mathbb{A}^{d \times d}$ with $f_t(z) = A_t z$. We simply write $S = (x, \{A_t : t \in [1, k]\}, y)$ for a linear dynamical system. The special case of a single matrix ($k = 1$) is called “simple linear loops.”

2.3 Invariants

Natural certificates that y is not reachable from x are separating invariants: An *invariant* is a set $I \subseteq \mathbb{C}^d$ such that $f_t(I) \subseteq I$ for all $t \in [1, k]$. It is *separating* for (x, y) if additionally $x \in I$ and $y \notin I$.

Note that for a linear function $f_t(z) = A_t z$, the property $f_t(I) \subseteq I$ is equivalent to $A_t I \subseteq I$, and in that case, we say that I is invariant under A .

The following are equivalent:

- there exists a separating invariant.
- y is not reachable from x .

It is clear that the existence of a separating invariant implies that y is not reachable from x . A stronger statement is: The set R of reachable points from x is a separating invariant for (x, y) if and only if y does not belong to R . However, the set R may be very complicated, making it not so useful as a separating invariant. We therefore consider restrictions on the class of invariants.

2.4 Semilinear Sets

A set $I \subseteq \mathbb{R}^d$ is semilinear if it is the set of (real) solutions of some finite Boolean combination of linear inequalities with algebraic coefficients. We give an equivalent definition now using half-spaces and polyhedra. A half-space \mathcal{H} is a subset of \mathbb{R}^d of the form

$$\mathcal{H} = \{z \in \mathbb{R}^d : z \cdot u > a\},$$

for some u in \mathbb{A}^d , a in $\mathbb{A} \cap \mathbb{R}$ and $> \in \{\geq, >\}$. A polyhedron is a finite intersection of half-spaces, and a semilinear set a finite union of polyhedra.

If I is a semilinear set, then I^o , \bar{I} and ∂I are also semilinear sets. A classical (and non-trivial) result about semilinear sets is their closure under projections as stated below. We will also need some effective bounds on sections of semilinear sets.

LEMMA 2.2 (PROJECTIONS OF SEMILINEAR SETS). *Let I be a semilinear set in $\mathbb{R}^{d+d'}$. Then, the projection of I on the first d coordinates defined by $\{z \in \mathbb{R}^d : \exists t \in \mathbb{R}^{d'}, (z, t) \in I\}$ is a semilinear set.*

LEMMA 2.3 (SECTIONS OF SEMILINEAR SETS). *Let I be a semilinear set in $\mathbb{R}^{d+d'}$ and t in $\mathbb{R}^{d'}$. Then, the section of I along t defined by $\{z \in \mathbb{R}^d : (z, t) \in I\}$ is a semilinear set. Furthermore, there exists a bound B in \mathbb{R} such that for all t in $\mathbb{R}^{d'}$ of norm at most 1, if the section is non-empty, then it contains some z in \mathbb{R}^d of norm at most B .*

PROOFS FOR LEMMAS 2.2 AND 2.3. Lemma 2.2 is a reformulation of Fourier-Motzkin elimination, from which the first part of Lemma 2.3 also follows. We now prove existence of the bound B . Let us first assume that I is closed and write $I = \mathcal{P}_1 \cup \dots \cup \mathcal{P}_n$ where the \mathcal{P}_i 's are closed polyhedra.

For each i , $T_i = \{t \in \mathbb{R}^d \mid \|t\| \leq 1 \text{ and } \exists z, (z, t) \in \mathcal{P}_i\}$ is a compact polyhedron, and the map $f_i : T_i \rightarrow \mathbb{R}$ assigning $\min\{\|z\|, (z, t) \in \mathcal{P}_i\}$ to t in T_i is continuous. Now, since T_i is compact, f_i admits a maximum B_i over T_i . Then, we simply let B be the maximal B_i .

If \mathcal{I} is not closed, then apply the above to $\mathcal{I}' = \overline{\mathcal{I}}$, which yields a bound B' . Then, for each t , the section of \mathcal{I} along t contains the interior of the section of \mathcal{I}' along t ; therefore, the bound $B = B' + 1$ applies to \mathcal{I} . \square

For the reader's intuitions, note that the last part of this lemma does not hold for more expressive domains. For instance, consider the hyperbola defined by $\mathcal{I} = \{(x, y) \in \mathbb{R}^2 : xy = 1\}$. Choosing a small x forces to choose a large y , hence there exists no bound B as stated in the lemma for \mathcal{I} .

It will be convenient for the proofs to consider semilinear sets in \mathbb{C}^d by identifying \mathbb{C}^d with \mathbb{R}^{2d} . Formally, $\mathcal{I} \subseteq \mathbb{C}^d$ is a *complex semilinear set* if $\mathcal{I}_{\mathbb{R}}$ is a (real) semilinear set. Note that this definition is consistent: $\mathcal{I} \subseteq \mathbb{R}^d$ is semilinear if and only if $\mathcal{I}_{\mathbb{R}}$ is semilinear. We will refer to complex semilinear sets as simply semilinear sets when it is clear from the context.

2.5 The Semilinear Invariant Problem

The problem we study in this article is the semilinear invariant problem, which asks whether given a linear dynamical system there exists a semilinear separating invariant. The next section gives high-level overviews of the proofs for our two main results, namely, Theorems 1.1 and 1.2.

3 Main Results and Proof Overviews

3.1 Undecidability for Two Matrices

We sketch the proofs of two undecidability results; as an intermediate step and towards the (complicated) proof of Theorem 1.1, we provide a simpler undecidability result regarding closed semilinear invariants. We will only sketch proofs in this section and defer the full proofs to Section 4. As it will appear below, it is more convenient here to write matrix-multiplication from the left, and vectors in row convention. We adopt this convention locally to this proof overview (Section 3.1), as well as in the full proof (Section 4).

We will construct reductions from the ω -Post Correspondence Problem (in short: ω -PCP), an extension of the well-known Post Correspondence Problem to infinite words. For a word w , we let $|w|$ denote its length, and for $i \in [1, |w|]$, we write w_i for the letter of w in position i , so $w = w_1 w_2 \dots$. We write $w_{1..n}$ for the prefix of w of length n .

An instance of the ω -PCP is given by a set of pairs of non-empty words $(u^i, v^i)_{i \in [1, p]}$ over some alphabet Σ . The objective is to determine whether there exists an infinite word $w = w_1 w_2 \dots$ over the alphabet $[1, p]$ such that the following equality over infinite words holds: $u^{w_1} u^{w_2} \dots = v^{w_1} v^{w_2} \dots$, and in that case, we say that w is a solution of $(u^i, v^i)_{i \in [1, p]}$. A pair (u^i, v^i) is called a tile (see Figure 1 for a graphical representation).

This problem is known to be undecidable [10] even for a fixed p and an alphabet of size 2. For the remainder of this section, we let p denote the smallest number such that the ω -PCP is undecidable with a fixed number of tiles p . The latest improvement on this result shows that $p \leq 8$ [10].

A First Undecidability Result for Closed Semilinear Invariants

The first undecidability result we prove is for (topologically) closed invariants: It does not yet imply Theorem 1.1; the reduction will be further refined later on.

THEOREM 3.1. *The semilinear invariant problem is undecidable for closed invariants with p matrices of dimension 3.*

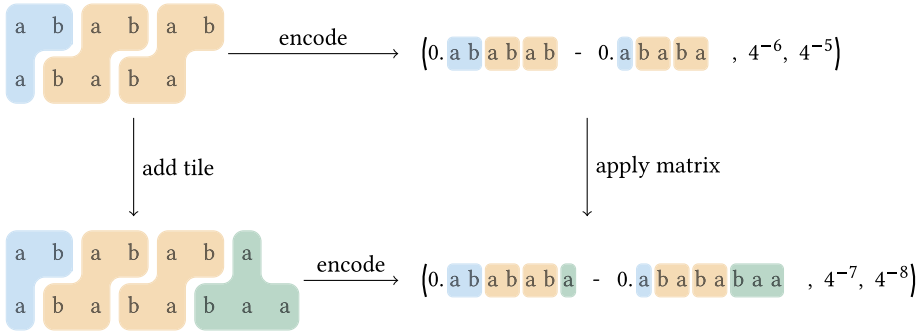


Fig. 1. Encoding using matrices: The partial solution consisting of three tiles on the left is encoded as three real numbers on the right (here, a and b are digits). For each tile, we construct a matrix such that concatenating the tile on the left is equivalent to multiplying this vector by the matrix corresponding to the tile.

Let us consider an ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$. Without loss of generality the alphabet is $\Sigma = \{0, 2\}$: This way, a word $u = u_1 \dots u_n \in \Sigma^*$ is encoded as the digits of some real number in $[0, 1]$ in base 4 (with least significant digits to the right):

$$[u] = \sum_{i=1}^n u_i 4^{-i}.$$

The choice of base 4 and digits in $\{0, 2\}$ instead of the more canonical base 2 is for having a sparse encoding, which will be useful for defining invariants. A finite word $w \in [1, p]^*$ induces two finite words $u^w, v^w \in \Sigma^*$:

$$u^w = u^{w_1} u^{w_2} \dots u^{w_n} \quad ; \quad v^w = v^{w_1} v^{w_2} \dots v^{w_n}.$$

We say that w is a partial solution if either u^w is a prefix of v^w or v^w a prefix of u^w .

We encode w by the vector $([u^w] - [v^w], 4^{-|u^w|}, 4^{-|v^w|})$ of dimension 3. Figure 1 illustrates the encoding of ω -PCP. The remarkable property of this encoding is that adding the tile (u^i, v^i) to w , meaning considering wi , corresponds to multiplying the vector by a fixed matrix A_i . Formally, for a tile (u^i, v^i) , we construct a 3×3 matrix A_i such that

$$\left([u^w] - [v^w], 4^{-|u^w|}, 4^{-|v^w|}\right) \cdot A_i = \left([u^{wi}] - [v^{wi}], 4^{-|u^{wi}|}, 4^{-|v^{wi}|}\right).$$

For a word $w \in [1, p]^*$, we define A_w : It is obtained by multiplying the matrices A_i following w . For instance,³ $A_{13,422} = A_1 A_3 A_4 A_2 A_2$. Note that the set of reachable points from x is $\{x \cdot A_w : w \in [1, p]^*\}$.

Let $x = (0, 1, 1)$, the equality above implies that

$$x \cdot A_w = \left([u^w] - [v^w], 4^{-|u^w|}, 4^{-|v^w|}\right).$$

Let $y = (0, 0, 0)$. Let us consider the system $S = (\{A_i\}_{i \in [1, p]}, x, y)$. We now argue the ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ does not have a solution if and only if there exists a closed separating semilinear invariant for S .

We first show that the existence of a solution implies the non-existence of closed separating invariants. Considering the prefixes of a solution of the ω -PCP yields a sequence of vectors that converges to the zero vector. In other words, in that case, y is in the topological closure of the reachable set from x . This implies that there cannot exist a closed separating invariant (semilinear or not).

³Here, it is convenient to use row-vectors and multiply from the left, otherwise, the order would be reversed, e.g., $A_{13,422} = A_2 A_2 A_4 A_3 A_1$. This is the reason why we adopt this convention for undecidability proofs.

Conversely, if there is no solution to the ω -PCP instance, then an application of König’s lemma implies that there exists a bound $N \in \mathbb{N}$ such that there are no partial solutions of length greater than N . It follows that for any $w \in [1, p]^+$ the first coordinate of $x \cdot A_w$, which is $[u^w] - [v^w]$, is lower bounded in absolute value by 4^{-N} . From this observation, we can construct a closed separating semilinear invariant (we refer to the full proof for details).

The Main Undecidability Result

The above reduction strongly relies on the fact that if the ω -PCP instance has a solution, then the target belongs to the closure of the set of reachable points, since this property implies that there cannot exist a *closed* separating invariant. To obtain the undecidability for the class of all semilinear invariants, we refine the reduction above.

THEOREM 3.2. *The semilinear invariant problem is undecidable with $p + 4$ matrices in dimension 8.*

Reducing to Two Matrices

In the reductions above, we used p matrices in dimension 3×3 and $p + 4$ matrices in dimension 8×8 . A standard transformation reduces the number of matrices by combining all matrices into one large matrix A and adding a shift matrix A_{shift} , yielding the following result strengthening of Theorem 1.1:

COROLLARY 3.3. *The closed semilinear invariant problem is undecidable with 2 matrices of dimension $3p$, and the semilinear invariant problem is undecidable with 2 matrices of dimension $8(p + 4)$.*

3.2 Decidability for Simple Linear Loops

Our positive result concerns the case of a single matrix, classically called “simple linear loops.” In this case, a system is (x, A, y) and called an Orbit instance, and the orbit of A under x is $\{A^n x : n \in \mathbb{N}\}$. The objective is to determine whether there exists a semilinear invariant, meaning a semilinear set \mathcal{I} such that $x \in \mathcal{I}$, $A\mathcal{I} \subseteq \mathcal{I}$, and $y \notin \mathcal{I}$. We say that an Orbit instance (x, A, y) is a reach-instance if $y = A^n x$ for some n ; since it is decidable in polynomial time [20, 21] whether (x, A, y) is a reach-instance, we may always assume that the answer is negative. Our decidability results are only concerned with *closed invariants*, which is crucial in several proofs. We might sometimes omit the adjective “closed,” but it is understood that whenever we consider an invariant it is closed.

THEOREM 1.2. *There is an algorithm that decides whether an Orbit instance with real algebraic coefficients admits a closed semilinear invariant. Furthermore, for instances with rational inputs, it runs in polynomial time assuming the dimension d is fixed.*

Before giving an overview of the proof of Theorem 1.2, we highlight some of the difficulties that occur by discussing a few examples.

Example 3.4. Consider the Orbit instance $\ell = (x, A, y)$ in dimension 2 where

$$A = \frac{1}{2} \begin{bmatrix} 1 & -2 \\ 2 & 1 \end{bmatrix},$$

$x = (0, 1)$ and $y = (-\frac{3}{2}, 0)$. The orbit is depicted on Figure 2 as the sequence of red dots.

The matrix A is a counterclockwise rotation around the origin with an expanding scaling factor. A suitable semilinear invariant can be constructed by taking the complement of the convex hull of a large enough number of points of the orbit and adding the missing points. In this example, we can take

$$\mathcal{I} = \{x, Ax\} \cup \overline{\text{Conv}(\{A^n x : n \leq 7\})^c},$$




Fig. 2. An invariant for Example 3.4.

which corresponds to the shaded region in Figure 2.

Example 3.5. Let us remove the expanding factor from the previous example by considering instead the following matrix:

$$A_1 = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 & -2 \\ 2 & 1 \end{bmatrix}.$$

Now, A_1 being a rotation of an irrational angle, the orbit of x is dense in the circle of radius 1. It is not too difficult to prove that there are no closed semilinear invariants (except for the whole space \mathbb{R}^2) for this instance, for any value of y . This gives a first instance of non-existence of a semilinear invariant. Many such examples exist, and we will next describe a more subtle one. Note that natural (non-semilinear) invariants do exist, such as the unit circle, which is a semi-algebraic set but not a semilinear one.

Example 3.6. Consider $\ell = (A_2, x, y)$ in dimension 4 with

$$A_2 = \begin{bmatrix} A_1 & I_2 \\ 0 & A_1 \end{bmatrix},$$

where A_1 is the matrix from Example 3.5, $x = (0, 0, 1, 0)$ and y is arbitrary. When repeatedly applying A_1 to x , the last two coordinates describe a circle of radius 1 as in the previous example. However, the first two coordinates diverge: At each step, they are rotated and the last two coordinates are added. Again, it is the case that there are no semilinear invariants (except again for the whole space \mathbb{R}^4), but it is much harder to prove than for Example 3.5.

However, even in instances similar to the one above, it may be the case that some coarse information in the input can still be captured by a semilinear invariant; for instance, if two synchronised blocks have some identical components. Let us illustrate this on an example.

Example 3.7. Consider the Orbit instance $\ell = (x, A, y)$ in dimension 8 where

$$A = \begin{bmatrix} A_2 & \\ & A_2 \end{bmatrix},$$

$x = (0, 0, 1, 0, 0, 0, 1, 0)$ and $y = (0, 0, 0, 1, 0, 0, 0, 2)$. We know from the previous example that when projecting on each block separately, there are only trivial semilinear invariants. However, there is indeed a semilinear invariant that exploits the synchronous behavior between the two blocks:

$$\mathcal{I} = \{z \in \mathbb{R}^4 : z_1 = z_3 \text{ and } z_2 = z_4\}.$$

This invariant has the property of being “strongly minimal”: It is contained in any semilinear set \mathcal{J} satisfying $A\mathcal{J} \subseteq \mathcal{J}$ and $A^n x \in \mathcal{J}$ for some n (this will be defined more formally below).




Fig. 3. The minimal semilinear invariant for Example 3.8.

Let us discuss two more examples having such strongly minimal semilinear invariants.

Example 3.8. Consider (A, x) in dimension 3 with

$$A = \begin{bmatrix} A_1 & 0 \\ 0 & -1 \end{bmatrix},$$

where A_1 is the matrix of Example 3.5 and $x = (1, 0, 1)$. As we iterate the matrix A , the two first coordinates describe a circle, and the third coordinate alternates between 1 and -1 : The orbit is dense in the union of two parallel circles (see Figure 3). In this example, the strongly minimal semilinear invariant is the union of the two planes containing these circles.

Example 3.9. Consider (A, x) in dimension 8 with

$$A = \begin{bmatrix} A_2 & 0 \\ 0 & -A_2 \end{bmatrix},$$

where A_2 is the matrix from Example 3.6. This can be seen as two instances of Example 3.6 running in parallel. Let $x = (0, 0, 1, 0, 0, 0, -7, 0)$, and note that both blocks of x are initially related by a multiplicative factor, namely, $-7(x_1, x_2, x_3, x_4) = (x_5, x_6, x_7, x_8)$. Moreover, as the first block is multiplied by the matrix A_2 while the second one is multiplied by $-A_2$, the multiplicative factor relating the two blocks alternates between 7 and -7 . Hence,

$$\mathcal{I} = \{u \in \mathbb{R}^8 : (u_1, u_2, u_3, u_4) = \pm 7(u_5, u_6, u_7, u_8)\}$$

is a semilinear invariant, which one can prove to be strongly minimal. Note that \mathcal{I} has dimension 4. If, however, we let $x = (0, 0, 1, 0, 0, 1, -7, 0)$, then the strongly minimal semilinear invariant becomes

$$\mathcal{I} = \{u \in \mathbb{R}^8 : (u_3, u_4) = \pm 7(u_7, u_8)\},$$

which has dimension 6. This shows that the strongly minimal semilinear invariant depends on x . Intuitively, in the second case no semilinear relation holds between (u_1, u_2) and (u_5, u_6) .

Proof overview. To prove Theorem 1.2, we proceed in three steps.

- (1) Identify positive cases, such as the one in Example 3.4, in which semilinear invariants always exist. These instances are called “simple instances.”
- (2) Reduce a non-simple instance to a so called “core instance.”
- (3) Prove that core instances admit only trivial semilinear invariants.

We now provide more details for each step separately.

Positive cases. The positive cases we identify fall into three categories:

- (i) There is a Jordan block J whose eigenvalue has modulus > 1 and such that $x_J \neq 0$ (this is just like Example 3.4).
- (ii) There is a Jordan block J whose eigenvalue has modulus < 1 and such that $y_J \neq 0$.
- (iii) There is a non-diagonal Jordan block J whose eigenvalue is a root of unity and such that $x_{J, >1} \neq 0$, which means that x has a non-zero coordinate on block J , which is not the first one.

We say that an instance is simple if it satisfies one of the three cases above and that it is non-simple otherwise. In each of these cases, we rely on the divergent behavior of the orbit to construct a semilinear invariant.

THEOREM 3.10. *Simple instances admit semilinear invariants.*

While cases (i) and (iii) are fairly straightforward, (ii) is more involved. Details are presented in Section 5.2.

Core instances. We now explain the third step, which amounts to identifying a class of core instances for which no non-trivial semilinear invariant exist. We say that a pair (x, A) defines a core pair⁴ if

- A is in Jordan normal form.
- All eigenvalues of A have modulus 1.
- No eigenvalue of A is a root of unity.
- Two different Jordan blocks of A are associated with different eigenvalues $\lambda \neq \lambda'$, and such that neither their product nor their quotient is a root of unity.
- For all Jordan blocks J , we have $x_{J, d(J)} \neq 0$, meaning the last coordinate of x on each block is non-zero.

Intuitively, in a core pair, no synchronisation phenomena may occur, so there exist only trivial invariants.

THEOREM 3.11. *Let (x, A) be a core pair of dimension d . The only closed semilinear set stable under A and containing x is \mathbb{C}^d .*

The proof of Theorem 3.11 is long and technical, it is the object of Section 6.

Reductions. The second step hence aims at reducing a non-simple orbit instance (x, A, y) to a core pair (x', A') . Towards this goal, we develop different reductions allowing for instance to remove Jordan blocks where x is zero, or to identify blocks J and J' when the associated eigenvalues λ and λ' satisfy some equations of the form $(\lambda\lambda')^n = 1$ or $\lambda^n = \lambda'^n$. In essence, this allows to capture linear relations that may hold for x among Jordan blocks with a synchronised behavior, as in Examples 3.6, or 3.9. Figure 4 displays our full pipeline of reductions; formal definitions for all classes of instances we consider appear in Section 5.3.

This allows us to conclude with the implications on the right-hand side of the figure. In the statement below, conjugated instances refer to those instances in Jordan normal form which originate from real matrices; for formal definition, we refer to Section 5.3.

THEOREM 3.12. *Assuming Theorem 3.11, there is a polynomial time algorithm deciding whether a non-simple conjugated Orbit instance admits a semilinear invariant.*

Details about reductions and establishing Theorem 3.12 are given in Section 5.3.

⁴Here, y is irrelevant, so we speak of pairs (x, A) rather than Orbit instances (x, A, y) .

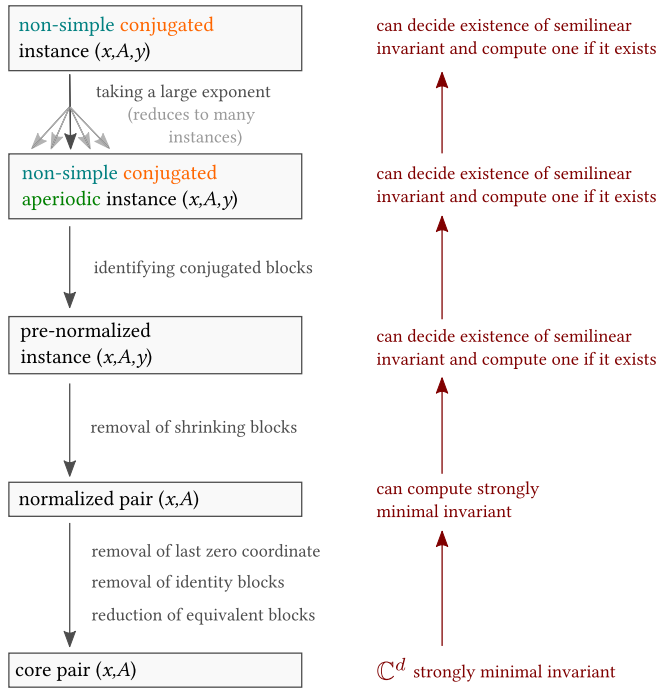


Fig. 4. The pipeline of reductions, from non-simple real orbit instances to core pairs.

4 Undecidability Proofs

The structure of this section follows the outline given in Section 3; we rely on the explanations given there but state and prove all technical details here.

4.1 Proof of Theorem 3.1

We start with proving Theorem 3.1: The semilinear invariant problem is undecidable for closed invariants with p matrices of dimension 3. We refer to Section 3.1 for the definition of ω -PCP, a sketch of the proofs, and the statements. Recall that p is the smallest number such that the ω -PCP is undecidable with a fixed number of tiles p for an alphabet of size 2; we know that $p \leq 8$ [10].

Let us consider an ω -PCP instance $(u^i, v^i)_{i \in [1,p]}$ over some alphabet Σ of size 2. A finite word $w \in [1,p]^*$ induces two finite words $u^w, v^w \in \Sigma^*$:

$$u^w = u^{w_1} u^{w_2} \dots u^{w_n}; v^w = v^{w_1} v^{w_2} \dots v^{w_n}.$$

We say that w is a partial solution if either u^w is a prefix of v^w or v^w a prefix of u^w . We state (and prove for the sake of completeness) a classical lemma on ω -PCP.

LEMMA 4.1. *Let $(u^i, v^i)_{i \in [1,p]}$ be an ω -PCP instance and $w \in [1,p]^\omega$.*

- *The infinite word $w \in [1,p]^\omega$ is a solution if and only if all prefixes of w are partial solutions.*
- *If there are no solutions, then there exists a bound N such that all partial solutions have length at most N .*

PROOF. The first item is clear, so we focus on the second. We consider the infinite tree with branching $[1,p]$: The set of nodes is $[1,p]^*$. We remove from the tree a node w if w is not a partial solution (note that we remove all of the descendants of w , since they are also not partial solutions).

Since the ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ does not have a solution, there are no infinite paths in this tree. The tree is finitely branching, so König's lemma implies that it is finite. Let N be the depth of this finite tree, then there are no partial solutions of length greater than N . \square

Let us write 0 and 2 for the two letters in Σ , meaning $\Sigma = \{0, 2\}$: This way, a word $u = u_1 \dots u_n \in \Sigma^*$ is encoded as the digits of some real number in $[0, 1]$ in base 4 (with least significant digit to the right):

$$[u] = \sum_{i=1}^n u_i 4^{-i}.$$

The choice of base 4 and digits in $\{0, 2\}$ instead of the more canonical base 2 is for having a "sparse" encoding as explained later. We encode $w \in [1, p]^*$ by the vector $([u^w] - [v^w], 4^{-|u^w|}, 4^{-|v^w|})$ of dimension 3. The remarkable property of this encoding is that adding the tile (u^i, v^i) to w , meaning considering wi , corresponds to multiplying the vector by the following matrix A_i :

$$A_i = \begin{bmatrix} 1 & 0 & 0 \\ [u^i] & 4^{-|u^i|} & 0 \\ -[v^i] & 0 & 4^{-|v^i|} \end{bmatrix}.$$

For $w \in [1, p]^\omega$, we write $w_1 \dots w_n$ for the prefix of length n of w . For $w \in [1, p]^*$, we define A_w as follows: A_w is obtained by multiplying the matrices A_i following w , for instance $A_{13,422} = A_1 A_3 A_4 A_2 A_2$.

Let $x = (0, 1, 1)$. We state in the following lemma the key properties of the encoding:

LEMMA 4.2. *Let $(u^i, v^i)_{i \in [1, p]}$ be an ω -PCP instance.*

- (1) *Let $w \in [1, p]^*$, we have $x \cdot A_w = ([u^w] - [v^w], 4^{-|u^w|}, 4^{-|v^w|})$.*
- (2) *Let us write $x \cdot A_w = (s, c, d)$. Then:*
 - *If w is a partial solution, then $|s| \leq \frac{2}{3}(c + d)$.*
 - *If w is not a partial solution, then $|s| > \frac{2}{3}(c + d)$.*
- (3) *Let $z = (s, c, d)$ and $\alpha \geq 0$. Let us write $z \cdot A_i = (s', c', d')$ for some $i \in [1, p]$. If $|s| \geq \frac{2}{3}(c + d) + \alpha$, then $|s'| \geq \frac{2}{3}(c' + d') + \alpha$.*

PROOF. We prove the three items.

- (1) The calculation for $x \cdot A_w$ is done by induction on w , noting that:

$$([u^w] - [v^w], 4^{-|u^w|}, 4^{-|v^w|}) \cdot A_i = ([u^w u_i] - [v^w v_i], 4^{-|u^w u_i|}, 4^{-|v^w v_i|}),$$

since $[u^w u_i] = [u^w] + 4^{-|u^w|}[u_i]$ and $[v^w v_i] = [v^w] + 4^{-|v^w|}[v_i]$.

- (2) Let $x \cdot A_w = (s, c, d)$.

– Assume that w is a partial solution: Either u^w is a prefix of v^w or the other way around. Assume the former holds: We have $v^w = u^w u'$ for some u' . This implies that $[v^w] = [u^w] + 4^{-|u^w|}[u']$. Since $[u'] \leq 1$, we obtain

$$|s| = |[v^w] - [u^w]| \leq 4^{-|u^w|} = c.$$

In the other case, the same reasoning yields $|s| \leq d$. Thus, $|s| \leq \frac{1}{2}(c + d) \leq \frac{2}{3}(c + d)$.

– Assume that w is not a partial solution, and let us write n for the smallest position such that $u_n^w \neq v_n^w$. Then,

$$[u^w] - [v^w] = (u_n^w - v_n^w) \frac{1}{4^n} + \sum_{j \geq n+1} (u_j^w - v_j^w) \frac{1}{4^j}.$$

The choice of base 4 and digits in $\{0, 2\}$ is all contained in the following calculations. Since $u_n^w \neq v_n^w$ and they are digits in $\{0, 2\}$, we have $|u_n^w - v_n^w| = 2$. For $j \geq n + 1$, we have $|u_j^w - v_j^w| \leq 2$ so

$$\left| \sum_{j \geq n+1} (u_j^{w_i} - v_j^{w_i}) \frac{1}{4^j} \right| < \frac{2}{4^{n+1}} \cdot \sum_{j \geq 0} \frac{1}{4^j} = \frac{2}{3} \cdot \frac{1}{4^n}.$$

It follows that

$$\begin{aligned} |s| = |[u^w] - [v^w]| &> 2 \cdot \frac{1}{4^n} - \frac{2}{3} \cdot \frac{1}{4^n} \\ &= \frac{4}{3} \cdot \frac{1}{4^n} \\ &\geq \frac{2}{3} (c + d). \end{aligned}$$

In the last inequality, we use $n \leq |u^w|$ and $n \leq |v^w|$.

(3) Let $z = (s, c, d)$ and $z \cdot A_i = (s', c', d')$ for some $i \in [1, p]$. Assume that $|s| \geq \frac{2}{3} (c + d) + \alpha$.

$$\begin{aligned} |s'| = |s + c[u^i] - d[v^i]| &\geq |s| - c[u^i] - d[v^i] \\ &\geq \frac{2}{3} (c + d) + \alpha - c[u^i] - d[v^i] \\ &= \underbrace{\left(\frac{2}{3} - [u^i] \right) c}_{\geq \frac{2}{3} \cdot 4^{-|u^i|}} + \underbrace{\left(\frac{2}{3} - [v^i] \right) d}_{\geq \frac{2}{3} \cdot 4^{-|v^i|}} + \alpha \\ &\geq \frac{2}{3} (c' + d') + \alpha. \end{aligned}$$

We have used the inequality $\frac{2}{3} - [u] \geq \frac{2}{3} \cdot 4^{-|u|}$, valid for $|u| \geq 1$. Thus, $|s'| \geq \frac{2}{3} (c' + d') + \alpha$. \square

Let $y = (0, 0, 0)$. We construct the linear dynamical system $S = (\{A_i\}_{i \in [1, p]}, x, y)$.

LEMMA 4.3. *The ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ does not have a solution if and only if there exists a closed separating semilinear invariant for S .*

PROOF. We distinguish two cases.

– Either the ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ has a solution $w \in [1, p]^\omega$. Thanks to Lemma 4.2, we have that for all $n \in \mathbb{N}$,

$$x \cdot A_{w_1 \dots w_n} = \left([u^{w_1 \dots w_n}] - [v^{w_1 \dots w_n}], 4^{-|u^{w_1 \dots w_n}|}, 4^{-|v^{w_1 \dots w_n}|} \right).$$

Since w is a solution, $w_1 \dots w_n$ is a partial solution, so, again, thanks to Lemma 4.2:

$$|[u^{w_1 \dots w_n}] - [v^{w_1 \dots w_n}]| \leq \frac{2}{3} \left(4^{-|u^{w_1 \dots w_n}|} + 4^{-|v^{w_1 \dots w_n}|} \right),$$

implying that $\lim_n x \cdot A_{w_1 \dots w_n} = (0, 0, 0) = y$. In other words, $y \in \overline{\{x \cdot A_w : w \in [1, p]^*\}}$, the topological closure of the set of reachable points from x .

Note that an invariant set \mathcal{I} for S containing x also contains the set of reachable points from x . If additionally \mathcal{I} is closed, then it contains its closure, hence it contains y . Thus, there are no closed semilinear invariants for S (note that we did not use semilinearity here).

- Or the ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ does not have a solution. Thanks to Lemma 4.1, there exists a bound N such that all partial solutions have length less than N . Let

$$\alpha = \min \left\{ |s| - \frac{2}{3}(c + d) : x \cdot A_w = (s, c, d) \text{ and } |w| = N \right\},$$

thanks to Lemma 4.2, we have $\alpha > 0$.

Let us define the sets

$$\begin{aligned} \mathcal{I}_1 &= \{x \cdot A_w : |w| < N\}, \\ \mathcal{I}_2 &= \{(s, c, d) : |s| \geq \frac{2}{3}(c + d) + \alpha\}, \\ \mathcal{I} &= \mathcal{I}_1 \cup \mathcal{I}_2. \end{aligned}$$

We argue that \mathcal{I} is a separating closed semilinear invariant. It is easy to see that \mathcal{I} is closed, semilinear, contains x , and does not contain y . We show that \mathcal{I} is indeed an invariant: Let $z \in \mathcal{I}$, we show that $z \cdot A_i \in \mathcal{I}$. We distinguish two cases.

- Either $z \in \mathcal{I}_1$, meaning $z = x \cdot A_w$ for $|w| < N$. Then, $z \cdot A_i = x \cdot A_{wi}$. If $|w| < N - 1$, then $|wi| < N$, so $z \cdot A_i = x \cdot A_{wi} \in \mathcal{I}_1$. Otherwise, $|wi| = N$, let us write $z \cdot A_i = (s, c, d)$. Since there are no partial solutions of length N , thanks to Lemma 4.2 and the definition of α , we have $|s| \geq \frac{2}{3}(c + d) + \alpha$. This shows that $z \cdot A_i \in \mathcal{I}_2$.
- Or $z \in \mathcal{I}_2$. Thanks to Lemma 4.2, we have $z \cdot A_i \in \mathcal{I}_2$. □

4.2 Proof of Theorem 3.2

For technical convenience, it will be useful to use affine transitions instead of linear ones; an affine transition is of the form $z \leftarrow z \cdot A + a$ for a matrix A and a vector a . A classical transformation reduces affine transitions to linear ones by adding a single dimension, as stated in the following lemma:

LEMMA 4.4. *Let S be a dynamical system with affine transitions in dimension d , we can construct a linear dynamical system S' in dimension $d+1$ such that there exists a (semilinear) separating invariant for S if and only if there exists a (semilinear) separating invariant for S' .*

We work in dimension 7 and divide a vector $z = (s, c, d, n, u, v, m)$ in two blocks: (s, c, d, n) and (u, v, m) . Let us define operations on each block:

- Resetting (s, c, d, n) is to perform the following operations, abbreviated $\text{Reset}(s, c, d, n)$:

$$s \leftarrow 0 ; c \leftarrow 1 ; d \leftarrow 1 ; n \leftarrow 0.$$

We say that (s, c, d, n) is “reset” if $(s, c, d, n) = (0, 1, 1, 0)$.

- Simulating i on (s, c, d, n) is to perform the following operations, abbreviated $\text{Simulation}_i(s, c, d, n)$, where $m = \max(|u_i|, |v_i|)$:

$$s \leftarrow 4^m (s + [u_i]c - [v_i]d) ; c \leftarrow 4^{m-|u_i|}c ; d \leftarrow 4^{m-|v_i|}d ; n \leftarrow n + 2.$$

- Resetting (u, v, m) is to perform the following operations, abbreviated $\text{Reset}(u, v, m)$:

$$u \leftarrow 0 ; v \leftarrow 0 ; m \leftarrow 0.$$

We say that (u, v, m) is “reset” if $(u, v, m) = (0, 0, 0)$.

We can now define the transitions.

- For each $i \in [1, p]$, the transition t_i does the following: $\text{Simulation}_i(s, c, d, n)$; $\text{Reset}(u, v, m)$.
- The transition t_{transfer} does the following: $u \leftarrow 3s - 2c - 2d$; $v \leftarrow -3s - 2c - 2d$; $m \leftarrow n$; $\text{Reset}(s, c, d, n)$.
- The transition $t_{\text{increase}(u)}$ does the following: $\text{Reset}(s, c, d, n)$; $u \leftarrow u + 1$.
- The transition $t_{\text{increase}(v)}$ does the following: $\text{Reset}(s, c, d, n)$; $v \leftarrow v + 1$.
- The transition $t_{\text{decrease}(m)}$ does the following: $\text{Reset}(s, c, d, n)$; $m \leftarrow m - 2$.

For a word $w \in [1, p]^*$, we write t_w for the composition of the transitions t_i following w : for instance, $t_{1423} = t_1 t_4 t_2 t_3$.

Let $\widehat{x} = (0, 1, 1, 0, 0, 0, 0)$ and $\widehat{y} = (0, 1, 1, 0, 0, 0, 1)$. We consider the system

$$\widehat{S} = (\{t_i : i \in [1, p]\} \cup \{t_{\text{transfer}}, t_{\text{increase}(u)}, t_{\text{increase}(v)}, t_{\text{decrease}(m)}\}, \widehat{x}, \widehat{y}).$$

LEMMA 4.5. *The ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ does not have a solution if and only if there exists a separating semilinear invariant for \widehat{S} .*

Since \widehat{S} uses affine transitions in dimension 7, we obtain an equivalent system using linear transitions in dimension 8 using Lemma 4.4.

PROOF. We distinguish two cases.

- Either the ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ has a solution $w \in [1, p]^\omega$. Let us consider a semilinear invariant \mathcal{I} for \widehat{S} containing \widehat{x} , and show that it necessarily contains \widehat{y} .

Let us consider the set $\mathcal{I}' = \{m \in \mathbb{R} : (0, 1, 1, 0, 0, 0, m) \in \mathcal{I}\}$. It is semilinear by closure under sections (Lemma 2.3). We argue that it contains all even natural numbers.

Let $n \in \mathbb{N}$. Starting from \widehat{x} and applying the transitions $t_{w_1}, t_{w_2}, \dots, t_{w_n}$, we reach $(s_n, c_n, d_n, 2n, 0, 0, 0)$ with $s_n, c_n, d_n \in \mathbb{Z}$ satisfying $|s_n| \leq \frac{2}{3}(c_n + d_n)$. Then, applying the transition t_{transfer} , we obtain $(0, 1, 1, 0, u_n, v_n, 2n)$ with $u_n, v_n \in \mathbb{Z}$ satisfying $u_n \leq 0$ and $v_n \leq 0$. From there, applying the transitions $t_{\text{increase}(u)}$ exactly $-u_n$ times and $t_{\text{increase}(v)}$ exactly $-v_n$ times yields $(0, 1, 1, 0, 0, 0, 2n)$. Since \mathcal{I} is invariant and contains \widehat{x} this implies that $(0, 1, 1, 0, 0, 0, 2n) \in \mathcal{I}$, so $2n \in \mathcal{I}'$.

Since any infinite semilinear set in dimension 1 over the reals must contain an odd natural number, it follows that \mathcal{I} contains $(0, 1, 1, 0, 0, 0, 2m+1)$ for some $m \in \mathbb{N}$. Indeed, a semilinear set in dimension 1 is a finite union of intervals, so if it contains all even natural numbers, then it must contain at least one odd number. Applying the transition $t_{\text{decrease}(m)}$ exactly m times, we conclude that \mathcal{I} contains $\widehat{y} = (0, 1, 1, 0, 0, 0, 1)$. Thus, there are no separating semilinear invariants for \widehat{S} .

- Or the ω -PCP instance $(u^i, v^i)_{i \in [1, p]}$ does not have a solution. Thanks to Lemma 4.1, there exists a bound N such that all partial solutions have length less than N . Let us define

$$\begin{aligned} \mathcal{I}_1 &= \{t_w(\widehat{x}) : w \in [1, p]^* \text{ with } |w| < N, \text{ and } (u, v, m) \text{ is reset}\}, \\ \mathcal{I}_2 &= \{z : |s| > \frac{2}{3}(c + d) \text{ and } (u, v, m) \text{ is reset}\}, \\ \mathcal{I}_3 &= \{z : (m \leq 0 \text{ or } m \in 2 \cdot [0, N] \text{ or } u > 0 \text{ or } v > 0), \text{ and } (s, c, d, n) \text{ is reset}\}, \\ \mathcal{I} &= \mathcal{I}_1 \cup \mathcal{I}_2 \cup \mathcal{I}_3. \end{aligned}$$

We argue that \mathcal{I} is a separating semilinear invariant for \widehat{S} . First \mathcal{I} is semilinear, contains \widehat{x} (because \mathcal{I}_1 does) and not \widehat{y} .

We show that \mathcal{I} is invariant. Let $z = (s, c, d, n, u, v, m) \in \mathcal{I}$, in the following case distinction, we write $t(z) = (s', c', d', n', u', v', m')$. We distinguish three cases, and for each consider all types of transitions:

- If $z \in \mathcal{I}_1$, then $z = t_w(\widehat{x})$ for some $w \in [1, p]^*$ with $|w| < N$.
 - * For $i \in [1, p]$, we have $t_i(z) \in \mathcal{I}_1$ or $t_i(z) \in \mathcal{I}_2$: if $|w| < N - 1$, then $t_i(z) = t_{wi}(\widehat{x}) \in \mathcal{I}_1$, since $|wi| < N$, otherwise, $|wi| = N$ and, since there are no partial solutions of length N , wi is not a partial solution, so, thanks to Lemma 4.2, we have $|s'| > \frac{2}{3}(c' + d')$, implying that $t_i(z) \in \mathcal{I}_2$.
 - * We have $t_{\text{transfer}}(z) \in \mathcal{I}_3$: since $n \in 2 \cdot [0, N]$, we have $m' = n \in 2 \cdot [0, N]$.
 - * We have $t_{\text{increase}(u)}(z) \in \mathcal{I}_3$: we have $m = 0$ so $m' = 0$.
 - * We have $t_{\text{increase}(v)}(z) \in \mathcal{I}_3$: we have $m = 0$ so $m' = 0$.

- * We have $t_{\text{decrease}(m)}(z) \in \mathcal{I}_3$: we have $m = 0$ so $m' = -2$.
 - If $z \in \mathcal{I}_2$, then $|s| > \frac{2}{3}(c + d)$.
 - * For $i \in [1, p]$, we have $t_i(z) \in \mathcal{I}_2$: thanks to Lemma 4.2, because (u, v, m) is reset.
 - * We have $t_{\text{transfer}}(z) \in \mathcal{I}_3$: either $u' > 0$ or $v' > 0$.
 - * We have $t_{\text{increase}(u)}(z) \in \mathcal{I}_3$: we have $m = 0$ so $m' = 0$.
 - * We have $t_{\text{increase}(v)}(z) \in \mathcal{I}_3$: we have $m = 0$ so $m' = 0$.
 - * We have $t_{\text{decrease}(m)}(z) \in \mathcal{I}_3$: we have $m = 0$ so $m' = -2$.
 - If $z \in \mathcal{I}_3$, then $m \leq 0$ or $(m \in 2 \cdot [0, N])$ or $u > 0$ or $v > 0$.
 - * For $i \in [1, p]$, we have $t_i(z) \in \mathcal{I}_1$: indeed, $t_i(z) = t_i(\widehat{x})$.
 - * We have $t_{\text{transfer}}(z) \in \mathcal{I}_3$: indeed, $n = 0$ so $m' = 0$.
 - * We have $t_{\text{increase}(u)}(z) \in \mathcal{I}_3$: indeed, $m \leq 0$ or $m \in 2 \cdot [0, N]$ or $u > 0$ or $v > 0$, so $m' \leq 0$ or $m' \in 2 \cdot [0, N]$ or $u' > 0$ or $v' > 0$.
 - * We have $t_{\text{increase}(v)}(z) \in \mathcal{I}_3$: indeed, $m \leq 0$ or $m \in 2 \cdot [0, N]$ or $u > 0$ or $v > 0$, so $m' \leq 0$ or $m' \in 2 \cdot [0, N]$ or $u' > 0$ or $v' > 0$.
 - * We have $t_{\text{decrease}(m)}(z) \in \mathcal{I}_3$: indeed, $m \leq 0$ or $m \in 2 \cdot [0, N]$, so $m' \leq 0$ or $m' \in 2 \cdot [0, N - 1]$.
- It follows that \mathcal{I} is a semilinear invariant for \widehat{S} . \square

4.3 Proof of Corollary 3.3

Consider a linear dynamical system $S_d = (\{A_i\}_{i \in [1, p]}, x, y)$ in dimension d , we construct a second linear dynamical system $S_{pd} = (\{A, A_{\text{shift}}\}, x', y')$ in dimension pd using only two matrices such that there exists a (closed, semilinear) separating invariant for S_d if and only if there exists a (closed, semilinear) separating invariant for S_{pd} .

We let I_d denote the identity matrix of size $d \times d$, $0_{d, d'}$ the zero matrix of size $d \times d'$, and 0_d the zero vector of size d . In particular, for $d' = 1$, we write 0_d for the zero vector of size d . We now define A and A_{shift} :

$$A = \begin{bmatrix} A_1 & & \cdots & 0 \\ & A_2 & & \vdots \\ \vdots & & \ddots & \\ 0 & \cdots & & A_p \end{bmatrix}, \quad A_{\text{shift}} = \begin{bmatrix} 0_{d, d} & I_d & & 0_{d, d} \\ & & \ddots & \\ & & & I_d \\ I_d & & & 0_{d, d} \end{bmatrix}.$$

For $z \in \mathbb{R}^d$ and $i \in [1, p]$, the i th shift $z^{\downarrow i} \in \mathbb{R}^{pd}$ of z is

$$z^{\downarrow i} = \begin{bmatrix} 0_{d(i-1)} \\ z \\ 0_{d(p-i)} \end{bmatrix}.$$

Note that $z^{\downarrow i} \cdot A_{\text{shift}} = z^{\downarrow (i \bmod p) + 1}$, justifying the name “shift.”

We let $x' = x^{\downarrow 1}$ and $y' = y^{\downarrow 1}$.

LEMMA 4.6. *Let S_d be a linear dynamical system using p matrices, the linear dynamical system S_{pd} constructed above satisfies the following: There exists a (closed, semilinear) separating invariant for S_d if and only if there exists a (closed, semilinear) separating invariant for S_{pd} .*

PROOF. Let \mathcal{I} be a separating invariant for S_d . Let

$$\mathcal{J} = \bigcup_{i=1}^p \left\{ z^{\downarrow i} \in \mathbb{R}^{pd} : z \in \mathcal{I} \right\}.$$

We argue that \mathcal{J} is a separating invariant for S_{pd} . Clearly, $x' \in \mathcal{J}$ and $y' \notin \mathcal{J}$. Let $z^{\downarrow i} \in \mathcal{J}$ for $i \in [1, p]$, then $z^{\downarrow i} \cdot A_{\text{shift}} = z^{\downarrow (i \bmod p) + 1}$, which is in \mathcal{J} , and $z^{\downarrow i} \cdot A = (z \cdot A_i)^{\downarrow i} \in \mathcal{J}$ is also in

\mathcal{J} . Thus, \mathcal{J} is a separating invariant for S_{pd} , and it is closed and semilinear if \mathcal{I} is closed and semilinear.

Conversely, let \mathcal{J} be a separating invariant for S_{pd} . Let

$$\mathcal{I} = \left\{ z \in \mathbb{R}^d : z^{\downarrow 1} \in \mathcal{J} \right\}.$$

We argue that \mathcal{I} is a separating invariant for S_d . Clearly, $x \in \mathcal{I}$ and $y \notin \mathcal{I}$. Let $z \in \mathcal{I}$ and $i \in [1, p]$, we show that $z \cdot A_i \in \mathcal{J}$, i.e., $(z \cdot A_i)^{\downarrow 1} \in \mathcal{I}$. We have

$$(z \cdot A_i)^{\downarrow 1} = z^{\downarrow 1} \cdot A_{\text{shift}}^{i-1} \cdot A \cdot A_{\text{shift}}^{d-i+1} \in \mathcal{J},$$

since $z^{\downarrow 1}$ and \mathcal{I} is invariant under A and A_{shift} . Thus, \mathcal{I} is a separating invariant for S_d , and it is closed and semilinear if \mathcal{J} is closed and semilinear. \square

Corollary 3.3 directly follows:

- Theorem 3.1 yields undecidability for p matrices of dimension 3, and using Lemma 4.6 this implies undecidability for 2 matrices of dimension $3p$.
- Theorem 3.2 yields undecidability for closed invariants for $p + 4$ matrices of dimension 8, and using Lemma 4.6 this implies undecidability for 2 matrices of dimension $8(p + 4)$.

5 Decidability Proof: Reducing to Core Instances

This section details the two first steps in our proof of Theorem 1.2. Section 5.1 introduces some terminology about reductions and Jordan normal form. Then, in Section 5.2, we eliminate simple instances, thereby proving Theorem 3.10. Last, in Section 5.3, we proceed to reduce from real non-simple instances to core pairs, establishing Theorem 3.12.

5.1 Reductions and Jordan Normal Form

Reductions between Orbit instances. Recall that an Orbit instance is (x, A, y) , where x is the initial vector, A is a matrix, and y the target vector.

A reduction from a class of Orbit instances C to another class of Orbit instances C' consists of the following:

- A function R mapping an Orbit instance $\ell \in C$ to an Orbit instance $R(\ell) \in C'$.
- For each Orbit instance $\ell \in C$, a function ϕ mapping any semilinear invariant \mathcal{I} of ℓ into a semilinear invariant $\phi(\mathcal{I})$ of $R(\ell)$.
- A function ψ mapping any semilinear invariant \mathcal{I}' of $R(\ell)$ into a semilinear invariant $\psi(\mathcal{I}')$ of ℓ .

We say that the reduction is polynomial time if all involved functions are computable in polynomial time. Clearly, if C reduces to C' , then for all $\ell \in C$, we have that ℓ and $R(\ell)$ are equivalent: One admits a semilinear invariant if and only if the other one does.

We will also consider reductions where we construct many Orbit instances instead of a single one; the definitions above are easily adapted to this scenario.

From real to complex Orbit instances. It is crucial in our proof to reduce a matrix to its Jordan normal (recalled below). This requires working with complex semilinear invariants, which is not an issue, thanks to the following Lemma:

LEMMA 5.1. *There exists a polynomial time reduction from real Orbit instances to complex Orbit instances using complex semilinear invariants.*

PROOF. Let $\ell = (x, A, y)$ be a real Orbit instance. We show that ℓ admits a real semilinear invariant if and only if it admits a complex one. Let \mathcal{I} be a real semilinear invariant for ℓ and let

$\mathcal{I}' = \{z \in \mathbb{C}^d : \text{Re}(z) \in \mathcal{I} \text{ and } \text{Im}(z) = 0\}$. Then, \mathcal{I}' is a complex semilinear set, $x \in \mathcal{I}'$, $y \notin \mathcal{I}'$ and $A\mathcal{I}' \subseteq \mathcal{I}'$, so it is complex invariant for ℓ . Conversely, let \mathcal{I} be a complex semilinear invariant for ℓ and let \mathcal{I}' be the section of \mathcal{I} along the real numbers:

$$\mathcal{I}' = \{v \in \mathbb{R}^d : \exists z \in \mathcal{I}, v = \text{Re}(z) \text{ and } \text{Im}(z) = 0\}.$$

Then, \mathcal{I}' is a (real) semilinear set by Lemma 2.3, $x \in \mathcal{I}'$, since x is real, $y \notin \mathcal{I}'$ for the same reason and $A\mathcal{I}' \subseteq \mathcal{I}'$, since A has real coefficients. Therefore, \mathcal{I}' is a (real) semilinear invariant for ℓ . \square

Jordan normal form. 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 (JNF)**, meaning that J is a diagonal block matrix where the blocks (called Jordan blocks) are of the form:

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

The complex number λ is an eigenvalue of A . We will sometimes use notation $\mathcal{J}_d(\lambda)$ for the Jordan block of size d with eigenvalue λ . A Jordan block of dimension one is called diagonal, and A is diagonalisable if and only if all Jordan blocks are diagonal. Note that the transformation into the Jordan normal form can be performed in polynomial time [3, 4].

The following convenient lemma states that changes of bases define structured reductions.

LEMMA 5.2. *Consider a class of Orbit instances C along with a fixed invertible matrix Q_ℓ for each $\ell = (x, A, y) \in C$, computable in polynomial time. There is a polynomial time reduction from C to $C' = \{(Q_\ell^{-1}x, Q_\ell^{-1}AQ_\ell, Q_\ell^{-1}y) \mid \ell = (x, A, y) \in C\}$.*

PROOF. Fix $\ell = (x, A, y) \in C$ and let Q denote Q_ℓ and $\ell' = (Q^{-1}x, Q^{-1}AQ, Q^{-1}y) = (x', A', y')$.

Let \mathcal{I} be an semilinear invariant for ℓ : $x \in \mathcal{I}$, $A\mathcal{I} \subseteq \mathcal{I}$ and $y \notin \mathcal{I}$. Let $\mathcal{I}' = Q^{-1}\mathcal{I}$. Then, $x' = Q^{-1}x \in Q^{-1}\mathcal{I} = \mathcal{I}'$, likewise, $y' = Q^{-1}y \in \mathcal{I}'$ and $A'\mathcal{I}' = Q^{-1}AQQ^{-1}\mathcal{I} = Q^{-1}A\mathcal{I} \subseteq Q^{-1}\mathcal{I} = \mathcal{I}'$. Therefore, \mathcal{I}' is a semilinear invariant for ℓ' . Conversely, given a semilinear invariant \mathcal{I}' for ℓ' , the proof that $\mathcal{I} = Q\mathcal{I}'$ is a semilinear invariant for ℓ follows exactly the same lines. \square

Notations regarding coordinates and Jordan blocks. When A is in JNF, we index the d coordinates in the matrix A by pairs (J, k) , where J ranges over the Jordan blocks of A and $k \in \{1, \dots, d(J)\}$, with $d(J)$ being 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 $(J_1, 1)$ and $(J_2, 1), (J_2, 2)$.

For $z \in \mathbb{C}^d$ and a subset S of dimensions, we let z_S be the projection of z on the dimensions in S , and extend this notation to matrices. For instance, $z_J \in \mathbb{C}^J$ is the vector corresponding to the dimensions of the Jordan block J , and $z_{J, >k}$ is its projection on the coordinates of the Jordan block J whose indices are greater than k . We write S^c for the dimensions that are not in S . We also write $\pi_S : \mathbb{C}^d \rightarrow \mathbb{C}^S$, where S is a set of coordinates, for the projection $z \mapsto z_S$.

Conjugated instances. We say that a matrix A is conjugated if it is in JNF and there is an involution $J \mapsto J^*$ between its Jordan blocks such that for all blocks, $A_{J^*} = A_J^*$. We say that an orbit instance (x, A, y) is conjugated if A is conjugated and, moreover, for all blocks J , we have $x_{J^*} = x_J^*$ and $y_{J^*} = y_J^*$.

Let $\gamma_i = \frac{\alpha_i}{\rho_i}$ and $\Gamma = \sum_{i=1}^n \gamma_i$. Let $\alpha'_i = \frac{\gamma_i}{\Gamma}$, then $\alpha'_i \in [0, 1]$ and $\sum_{i=1}^n \alpha'_i = 1$ by definition. Therefore, $\sum_{i=1}^n \alpha'_i \rho_i x_i \in C'$ but $\sum_{i=1}^n \alpha'_i \rho_i x_i = \frac{1}{\Gamma} \sum_{i=1}^n \alpha_i x_i = \frac{\xi}{\Gamma}$. This shows that $B(0, \frac{\epsilon}{\Gamma}) = \frac{1}{\Gamma} B(0, \epsilon) \subseteq C'$. \square

LEMMA 5.5. *Let λ be a complex non-real number and x be a non-zero complex number. Then, there exists $n \in \mathbb{N}$ such that $\text{Conv}(\{\lambda^i x : i \in [0, n]\})$ contains an open ball centered at 0.*

PROOF. Let $\alpha = \frac{\lambda}{|\lambda|}$, which is also non-real. We claim that there exists n such that $C_n := \text{Conv}(\{\alpha^i : i \in [0, n]\})$ contains an open ball $B(0, \epsilon)$ for some $\epsilon > 0$. Indeed, let $\theta = \text{Arg}(\alpha)$ where $\text{Arg}(\cdot) \in (-\pi, \pi]$ denotes the principal argument. Then, $\theta \notin \{0, \pi\}$, since α is not real.

– If $\theta \notin \pi\mathbb{Q}$, then $\{\alpha^i : i \in \mathbb{N}\}$ is dense in the unit circle. Therefore, for sufficiently large n , C_n contains four points at distance at most $\frac{1}{2}$ from 1, i , -1 and $-i$. The resulting four points will then form a polygon that contains the open ball $B(0, \frac{1}{\sqrt{2}})$.

– Otherwise, the set $\{\alpha^i : i \in \mathbb{N}\}$ is finite and equal to the group of the n th roots of unity for some $n \geq 3$ ($n \neq 1, 2$ for otherwise $\theta \in \{0, \pi\}$). Therefore, C_n is a regular polygon with n faces, centered at the origin, so it contains an open ball of radius $1/2$, centered at the origin.

It follows by Lemma 5.4, for $\rho_i = |\lambda^i| > 0$, that $C'_n := \text{Conv}(\{\lambda^i : i \in [0, n]\})$ contains an open ball $B(0, \epsilon')$ for some $\epsilon' > 0$. But then, $\text{Conv}(\{\lambda^i x : i \in [0, n]\}) = C'_n x \supset B(0, \epsilon')x = B(0, \epsilon'|x|)$, which is open, since $x \neq 0$. \square

LEMMA 5.6. *Let λ be a complex non-real number of modulus greater than 1 and x be a non-zero complex number. Then, the sequence of polyhedra $(\text{Conv}(\{\lambda^i x : i \in [0, n]\}))_{n \in \mathbb{N}}$ is strictly increasing and its union is \mathbb{C} .*

PROOF. Let $C_n = \text{Conv}(\{\lambda^i x : i \in [0, n]\})$ for all $n \in \mathbb{N}$. To see that the sequence is strictly increasing, observe that for all n in \mathbb{N} , we have $C_n \subseteq \overline{B}(0, |\lambda|^n \cdot |x|)$. It follows that $\lambda^{n+1}x$ is not in C_n . To see that its union is \mathbb{C} , apply Lemma 5.5 to get n_0 such that C_{n_0} contains an open ball $B(0, \epsilon)$ for some $\epsilon > 0$. Then, note that for any $n \in \mathbb{N}$,

$$C_{n_0+n} \supseteq \text{Conv}(\{\lambda^{n_0+i} x : i \in [0, n]\}) = \lambda^k C_{n_0} \supset \lambda^k B(0, \epsilon) = B(0, |\lambda|^k \epsilon).$$

This concludes because the union of all such balls for $n \in \mathbb{N}$ is \mathbb{C} , since $|\lambda| > 1$. \square

THEOREM 5.7. *Let $\ell = (x, A, y)$ be a non-reach Orbit instance in JNF. If there exists a Jordan block J associated with an eigenvalue whose modulus is greater than 1 and such that $x_J \neq 0$, then there exists a semilinear invariant for ℓ .*

On an intuitive level first: Some coordinate of $(A^n x)_{n \in \mathbb{N}}$ diverges to infinity, so eventually gets larger in absolute value than the corresponding coordinate in y . This allows us to construct an invariant for ℓ by taking the first points and then all points having a large coordinate in the diverging dimension. For the invariant to be semilinear, we consider the complement of the convex envelope of an initial segment of points.

PROOF. We distinguish two cases. Let (J, s) denote the last coordinate of the Jordan block J such that $x_{J,s} \neq 0$; observe that $(A^n x)_{J,s} = \lambda^n x_{J,s}$.

– Suppose that λ is a real number.

For all $n \in \mathbb{N}$, we have $(A^n x)_{J,s} = \lambda^n x_{J,s}$, so it diverges to infinity in modulus. It follows that there exists n_0 in \mathbb{N} such that $|(A^{n_0} x)_{J,s}| \geq 2\sqrt{2} \cdot |y_{J,s}|$. Let

$$\mathcal{I} = \{x, Ax, \dots, A^{n_0-1}x\} \cup \{z \in \mathbb{C}^d : |\text{Re}(z_{J,s})| + |\text{Im}(z_{J,s})| \geq 2|y_{J,s}|\}.$$

We argue that \mathcal{I} is a semilinear invariant for ℓ . The non-trivial point is that \mathcal{I} is invariant under A . First, $A^{n_0}x$ is in \mathcal{I} because

$$|\operatorname{Re}((A^{n_0}x)_{J,s})| + |\operatorname{Im}((A^{n_0}x)_{J,s})| \geq \frac{1}{\sqrt{2}} \cdot |(A^{n_0}x)_{J,s}| \geq 2|y_{J,s}|.$$

Then, let $z \in \mathbb{C}^d$ such that $|\operatorname{Re}(z_{J,s})| + |\operatorname{Im}(z_{J,s})| \geq 2|y_{J,s}|$, we have that $(Az)_{J,s} = \lambda z_{J,s}$, so, since λ is real,

$$|\operatorname{Re}((Az)_{J,s})| + |\operatorname{Im}((Az)_{J,s})| = |\lambda| (|\operatorname{Re}(z_{J,s})| + |\operatorname{Im}(z_{J,s})|) \geq 2|y_{J,s}|,$$

thus Az is in \mathcal{I} .

– Suppose that λ is not a real number.

For any $n \in \mathbb{N}$, let $C_n = \operatorname{Conv}(\{\lambda^i x_{J,s} : i \in [1, n]\})$. By Lemma 5.6, the sequence $(C_n)_{n \in \mathbb{N}}$ of polyhedra in \mathbb{C} is strictly increasing and its union is \mathbb{C} . Let $n_0 \in \mathbb{N}$ be such that $y_{J,s}$ is in the interior of C_{n_0} . Finally, let

$$\mathcal{I} = \{x, Ax, \dots, A^{n_0}x\} \cup \overline{P}, \quad \text{where } P = \{z \in \mathbb{C}^d : z_{J,s} \notin C_{n_0}\}.$$

Note that \mathcal{I} is a closed semilinear set. We argue that \mathcal{I} is a semilinear invariant for ℓ . The non-trivial point is that \mathcal{I} is invariant under A .

We first need to prove that $A^{n_0+1}x$ is in \mathcal{I} . We have $(A^{n_0+1}x)_{J,s} = \lambda^{n_0+1}x_{J,s}$, which is not in C_{n_0} , because $C_{n_0+1} = \operatorname{Conv}(\{\lambda^{n_0+1}x_{J,s}\} \cup C_{n_0})$ and we have argued that the sequence $(C_n)_{n \in \mathbb{N}}$ is strictly increasing. Thus, $A^{n_0+1}x \in P \subseteq \overline{P} \subseteq \mathcal{I}$.

Second, we will show that $A\overline{P} \subseteq \overline{P}$; by continuity of matrix multiplication, it is sufficient to show that $AP \subseteq P$. Let $z \in P$, i.e., $z_{J,s} \notin C_{n_0}$ and assume towards contradiction that $Az \notin P$, i.e., $(Az)_{J,s} \in C_{n_0}$. But note that $(Az)_{J,s} = \lambda z_{J,s}$ so $z_{J,s} \subseteq \lambda^{-1}C_{n_0}$. However,

$$\lambda^{-1}C_{n_0} = \operatorname{Conv}(\{\lambda^i x_{J,s} : i \in [0, n_0]\}) = \operatorname{Conv}(\{x_{J,s}\} \cup C_{n_0-1}) \subseteq C_{n_0}$$

by convexity, since $x_{J,s} \in C_{n_0}$ and the sequence $(C_n)_n$ is increasing. Hence, $z_{J,s} \in C_{n_0}$, a contradiction. \square

5.2.2 Some Eigenvalue Has Modulus Less than 1. We now move on to the second case, which is the most involved of the three. We start with a simple lemma.

LEMMA 5.8. *Let λ be a complex non-real number of modulus less than 1 and x be a non-zero complex number. Then, the sequence $(\operatorname{Conv}(\{\lambda^i x : i \in [0, n]\}))_{n \in \mathbb{N}}$ of polyhedra in \mathbb{C} is ultimately constant, and its union contains an open neighbourhood of 0.*

PROOF. Let $C_n = \operatorname{Conv}(\{\lambda^i x : i \in [0, n]\})$ for all $n \in \mathbb{N}$. Apply Lemma 5.5 to get n_0 such that C_{n_0} contains an open ball $B(0, \varepsilon)$ for some $\varepsilon > 0$. Since $|\lambda| < 1$, there exists $n_1 \geq n_0$ such that $|\lambda|^{n_1} \cdot |x| < \varepsilon$. Note that $C_n \subseteq B(0, |x|)$ for all n , since $|\lambda| < 1$. Therefore,

$$\lambda^{n_1}C_n \subseteq \lambda^{n_1}B(0, |x|) = B(0, |\lambda|^{n_1} \cdot |x|) \subseteq B(0, \varepsilon) \subseteq C_{n_0}.$$

It follows that for any $n \geq n_1$,

$$C_{n_1} \subseteq C_n = \operatorname{Conv}(C_{n_1} \cup \lambda^{n_1}C_{n-n_1}) \subseteq \operatorname{Conv}(C_{n_1} \cup C_{n_0}) = C_{n_1}. \quad \square$$

The following lemma is the cornerstone for this section:

LEMMA 5.9. *Let $\varepsilon > 0$ and $\lambda \in \mathbb{C}$ with $|\lambda| < 1$. There exists a convex closed semilinear set $\mathcal{I} \subseteq B(0, \varepsilon) \subseteq \mathbb{C}^d$ that is invariant under the Jordan block $\mathcal{J}_d(\lambda)$ and contains $B(0, \varepsilon')$ for some $0 < \varepsilon' < \varepsilon$.*

PROOF. We let J denote $\mathcal{J}_d(\lambda)$. Note that $|z| \leq |\operatorname{Re}(z)| + |\operatorname{Im}(z)| \leq \sqrt{2}|z|$ for any $z \in \mathbb{C}$. We first treat the case where $\lambda \in \mathbb{R}$. Let

$$\mathcal{I} = \{z \in \mathbb{C}^d : \forall i \in [1, d], |\operatorname{Re}(z_i)| + |\operatorname{Im}(z_i)| \leq \varepsilon(1 - |\lambda|)^i\} \subseteq B(0, \varepsilon).$$

Then, $B(0, \varepsilon(1 - |\lambda|)^d / \sqrt{2}) \subseteq \mathcal{I}$. We show that $J\mathcal{I} \subseteq \mathcal{I}$. Let $z \in \mathcal{I}$. Then, $(Jz)_d = \lambda z_d$, so, since λ is real, $|\operatorname{Re}((Jz)_d)| + |\operatorname{Im}((Jz)_d)| \leq |\lambda|(|\operatorname{Re}(z_d)| + |\operatorname{Im}(z_d)|) \leq \varepsilon(1 - |\lambda|)^d$. Now, if $i < d$, then $(Jz)_i = \lambda z_i + z_{i+1}$, so

$$\begin{aligned} |\operatorname{Re}((Jz)_i)| + |\operatorname{Im}((Jz)_i)| &= |\lambda \operatorname{Re}(z_i) + \operatorname{Re}(z_{i+1})| + |\lambda \operatorname{Im}(z_i) + \operatorname{Im}(z_{i+1})| \\ &\leq |\lambda|(|\operatorname{Re}(z_i)| + |\operatorname{Im}(z_i)|) + (|\operatorname{Re}(z_{i+1})| + |\operatorname{Im}(z_{i+1})|) \\ &\leq |\lambda|\varepsilon(1 - |\lambda|)^i + \varepsilon(1 - |\lambda|)^{i+1} = \varepsilon(1 - |\lambda|)^i. \end{aligned}$$

Hence, \mathcal{I} is invariant under J , which concludes this first case.

We now assume that $\lambda \notin \mathbb{R}$ and prove the result by induction on d . We start with the base case $d = 1$. Fix $u \in \mathbb{C}$ of modulus ε ; for instance, $u = \varepsilon$. By Lemma 5.8, there exists n such that $\mathcal{I} := \operatorname{Conv}(\{\lambda^i u : i \in [0, n]\})$ contains an open ball centered at 0 and $\operatorname{Conv}(\{\lambda^i u : i \in [0, m]\}) = \mathcal{I}$ for all $m \geq n$. Since the extremal points of \mathcal{I} are of the form $\lambda^i u$, of modulus $|\lambda|^i \varepsilon < \varepsilon$, it holds that $\mathcal{I} \subseteq B(0, \varepsilon)$. Finally,

$$J\mathcal{I} = \operatorname{Conv}(\{\lambda^i u : i \in [1, n+1]\}) \subseteq \operatorname{Conv}(\{\lambda^i u : i \in [0, n+1]\}) = \mathcal{I}.$$

For $d > 1$, let $\varepsilon' > 0$ to be fixed later on. By induction, there exists a convex closed semilinear subset \mathcal{I}' of \mathbb{C}^{d-1} , invariant under $\mathcal{J}_{d-1}(\lambda)$, and such that

$$B(0, \varepsilon'') \subseteq \mathcal{I}' \subseteq B(0, \varepsilon') \subseteq \mathbb{C}^{d-1}, \quad (1)$$

for some $\varepsilon'' > 0$. Intuitively, we want to define \mathcal{I} of the form $\mathcal{I} = C \times \mathcal{I}'$ for some semilinear set $C \subseteq \mathbb{C}$. Note that the action of J on such a set satisfies

$$J(C \times \mathcal{I}') \subseteq (\lambda C + \pi_1(\mathcal{I}')) \times \mathcal{J}_{d-1}(\lambda)\mathcal{I}' \subseteq (\lambda C + \pi_1(\mathcal{I}')) \times \mathcal{I}'.$$

Therefore, we want to find C such that $C \subseteq \lambda C + \pi_1(\mathcal{I}')$. The idea to find C is to start from an arbitrary point u and then add what we need until the set is stable. We will then see that this process converges (after infinitely many steps), so, eventually (after n steps), all those sets are contained in a small ball. We then define C to be the convex hull of the first n sets: The first $n - 1$ sets will be stable by construction, and the last element will be contained in the small ball, itself contained in $\operatorname{Conv}(\{\lambda^i u : i \in [0, n]\})$, thanks to Lemma 5.8.

Formally, let u be a complex number of modulus $\varepsilon/2$, for instance, $u = \varepsilon/2 \in \mathbb{C}$. By Lemma 5.8, there exists n_0 such that $\operatorname{Conv}(\{\lambda^i u : i \in [0, n_0]\})$ contains an open ball $B(0, \delta)$ for some $\delta > 0$. Let $\varepsilon' = |1 - \lambda|\delta/4$ and \mathcal{I}' defined as in Equation (1). Note that $B(0, \delta) \subseteq \operatorname{Conv}(\{\lambda^i u : i \in [0, n_0]\}) \subseteq B(0, |u|) = B(0, \varepsilon/2)$ so $\delta \leq \varepsilon/2$, and $\varepsilon' \leq \delta/2 \leq \varepsilon$. We then let

$$C_0 = \{u\}, \quad \text{and } C_{n+1} = \lambda C_n + \pi_1(\mathcal{I}')$$

for all $n \in \mathbb{N}$. Since for convex sets S and reals a, b it holds that $aS + bS = (a + b)S$, it follows from convexity of $\pi_1(\mathcal{I}')$ that for all $n \in \mathbb{N}$,

$$C_n = \lambda^n C_0 + \frac{1 - \lambda^n}{1 - \lambda} \pi_1(\mathcal{I}').$$

Recall that $\mathcal{I}' \subseteq B(0, \varepsilon')$ so $\pi_1(\mathcal{I}') \subseteq B(0, \varepsilon')$ and therefore,

$$C_n \subseteq |\lambda|^n B(0, |u|) + \left| \frac{1 - \lambda^n}{1 - \lambda} \right| B(0, \varepsilon') \subseteq B\left(0, |\lambda|^n |u| + \left| \frac{1 - \lambda^n}{1 - \lambda} \right| \varepsilon'\right). \quad (2)$$

Since $|\lambda|^n|u| + \frac{1-\lambda^n}{1-\lambda}|\varepsilon'| \rightarrow \frac{\varepsilon'}{|1-\lambda|}$ as $n \rightarrow \infty$, there exists $n_1 \geq n_0$ such that $C_n \subseteq B(0, \frac{2\varepsilon'}{|1-\lambda|}) = B(0, \delta/2)$ for all $n \geq n_1$. We now define

$$\mathcal{I} = C \times \mathcal{I}', \quad \text{where } C = \text{Conv}(C_0 \cup \dots \cup C_{n_1}).$$

It is clear that \mathcal{I} is a convex closed semilinear set. We now claim that:

– $\mathcal{I} \subseteq B(0, \varepsilon)$: We have that $\mathcal{I}' \subseteq B(0, \varepsilon') \subseteq B(0, \varepsilon)$ by construction, and $C \subseteq B(0, \varepsilon)$ by Equation (2), since for all $n \in \mathbb{N}$,

$$|\lambda|^n|u| + \left| \frac{1-\lambda^n}{1-\lambda} \right| \varepsilon' \leq |u| + 2 \frac{\varepsilon'}{|1-\lambda|} \leq \frac{\varepsilon}{2} + \frac{\delta}{2} \leq \varepsilon,$$

since $\delta \leq \varepsilon/2$, as argued above. This concludes because $B(0, \varepsilon) \times B(0, \varepsilon) = B(0, \varepsilon)$ for the infinity norm.

– \mathcal{I} contains the open ball $B(0, \delta)$: We have that $B(0, \delta) \subseteq \text{Conv}(\{\lambda^i u : i \in [0, n_0]\})$ by construction. Furthermore, $\lambda^n u \in C_n$, because $0 \in \mathcal{I}'$, hence $0 \in \pi(\mathcal{I}')$, therefore, $C_{n+1} = \lambda C_n + \pi(\mathcal{I}') \supseteq \lambda C_n$ for all n . It follows that C contains $\{\lambda^j u : j \in [0, n_1]\}$ and this concludes because $n_1 \geq n_0$ and C is convex.

– \mathcal{I} is stable under J : Recall that

$$J\mathcal{I} = J(C \times \mathcal{I}') \subseteq (\lambda C + \pi_1(\mathcal{I}')) \times \mathcal{J}_{d-1}(\lambda)\mathcal{I}' \subseteq (\lambda C + \pi_1(\mathcal{I}')) \times \mathcal{I}',$$

since \mathcal{I}' is stable under $\mathcal{J}_{d-1}(\lambda)$. Therefore, it suffices to show that $\lambda C + \pi_1(\mathcal{I}') \subseteq C$. We first claim that $\lambda C_n + \pi_1(\mathcal{I}') \subseteq C$ for all $n \in [0, n_1]$. Indeed, for $n \in [0, n_1 - 1]$, we have $\lambda C_n + \pi_1(\mathcal{I}') = C_{n+1} \subseteq C$, since $n+1 \leq n_1$. And for C_{n_1} , we have

$$\lambda C_{n_1} + \pi_1(\mathcal{I}') \subseteq B(0, |\lambda|\delta/2) + B(0, \varepsilon') \subseteq B(0, \delta/2 + \varepsilon') \subseteq B(0, \delta) \subseteq \text{Conv}(\{\lambda^i u : i \in [0, n_0]\}) \subseteq C$$

by Equations (1), (2), and the definition of n_1 , the convexity of C , and the fact that $n_1 \geq n_0$. Now, let $x \in \lambda C + \pi_1(\mathcal{I}')$ and write $x = \lambda \sum_{i=0}^{n_1} \alpha_i x_i + y$ where $\sum_{i=0}^{n_1} \alpha_i = 1$, $x_i \in C_0 \cup \dots \cup C_{n_1}$ and $y \in \pi_1(\mathcal{I}')$. We can rewrite x as $x = \sum_{i=1}^{n_1} \alpha_i (\lambda x_i + y)$. Observe that for each i , $x_i \in C_j$ for some j so $\lambda x_i + y \in \lambda C_j + \pi_1(\mathcal{I}') \subseteq C$ by the above. Therefore, $x \in \text{Conv}(C) = C$. \square

We may now prove the following theorem:

THEOREM 5.10. *Let $\ell = (x, A, y)$ be a non-reach Orbit instance in \mathfrak{JNF} . If A has a Jordan block J associated with an eigenvalue whose modulus is less than 1 and such that $y_J \neq 0$, then there exists a semilinear invariant for ℓ .*

PROOF. Let $\varepsilon = \|y_J\|/2$. Thanks to Lemma 5.9, there exist $\varepsilon' > 0$ and a closed semilinear set $\mathcal{I} \subseteq \mathbb{C}^{d(J)}$ such that $J\mathcal{I} \subseteq \mathcal{I}$ and $B(0, \varepsilon') \subseteq \mathcal{I} \subseteq B(0, \varepsilon)$. Now, $(A^n x)_J \rightarrow 0$, so there exists n_0 such that $(A^{n_0} x)_J \in B(0, \varepsilon') \subseteq \mathcal{I}$. Hence,

$$\{x, Ax, \dots, A^{n_0-1}x\} \cup \{z \in \mathbb{C}^d : z_J \in \mathcal{I}\}$$

is a semilinear invariant for ℓ . \square

5.2.3 Some Non-diagonalisable Eigenvalue Is a Root of Unity. We now move on to the third positive case.

THEOREM 5.11. *Let $\ell = (x, A, y)$ be a non-reach Orbit instance in \mathfrak{JNF} . If there exists a non-diagonal Jordan block J associated with an eigenvalue which is a root of unity and such that $x_{J, >1} \neq 0$, then there exists a semilinear invariant for ℓ .*

PROOF. Let m be such that $\lambda^m = 1$, and let (J, s) be the maximal coordinate such that $x_{J,s}$ is non-zero. We rely on the divergence of the coordinate $(J, s - 1)$ to construct an invariant. For any $n \in \mathbb{N}$, we have $(A^n x)_{J,s-1} = \lambda^n x_{J,s-1} + n\lambda^{n-1} x_{J,s}$ and $(A^n x)_{J,s} = \lambda^n x_{J,s}$. Recall that z^* denotes the complex conjugate of $z \in \mathbb{C}$. Hence,

$$\operatorname{Re}(\lambda(A^n x)_{J,s-1}(A^n x)_{J,s}^*) = \operatorname{Re}(\lambda x_{J,s-1} x_{J,s}^*) + n|x_{J,s}|^2,$$

which goes to infinity when n grows. Note that this condition is quadratic, however, since $(A^n x)_{J,s} = \lambda^n x_{J,s}$ only takes a finite number of values, we will be able to construct a semilinear set from it. Let n_0 be such that

$$M := \operatorname{Re}(\lambda(A^{n_0} x)_{J,s-1}(A^{n_0} x)_{J,s}^*) > \operatorname{Re}(\lambda y_{J,s-1} y_{J,s}^*).$$

Finally, let

$$\mathcal{I} = \{x, Ax, \dots, A^{n_0-1}x\} \cup \bigcup_{i=0}^{m-1} \mathcal{I}_i, \quad \text{where } \mathcal{I}_i = \left\{z \in \mathbb{C}^d : z_{J,s} = \lambda^i x_{J,s} \text{ and } \operatorname{Re}(\lambda z_{J,s-1} z_{J,s}^*) \geq M\right\}.$$

It is clear that $x \in \mathcal{I}$ and $y \notin \mathcal{I}$. Each \mathcal{I}_i is semilinear, because the second condition is actually semilinear assuming $z_{J,s} = \lambda^i x_{J,s}$. There remains to see that $A^{n_0} x \in \mathcal{I}_{n_0 \bmod m}$. Indeed, $(A^{n_0} x)_{J,s} = \lambda^{n_0} x_{J,s} = \lambda^{n_0 \bmod m} x_{J,s}$ and we have defined M above so $\operatorname{Re}(\lambda z_{J,s-1} z_{J,s}^*) \geq M$ for $z = A^{n_0} x$. Now, if $z \in \mathcal{I}_i$, then we obtain that $(Az)_{J,s} = \lambda z_{J,s} = \lambda^{i+1} x_{J,s}$, and

$$\operatorname{Re}(\lambda(Az)_{J,s-1}(Az)_{J,s}^*) = \operatorname{Re}(\lambda z_{J,s-1} z_{J,s}^*) + |z_{J,s}|^2 \geq M + |z_{J,s}|^2 \geq M$$

so $Az \in \mathcal{I}_{i+1}$ if $i < m$, and $Az \in \mathcal{I}_0$ if $i = m$ (since $\lambda^m = 1$). Hence, \mathcal{I} is invariant under A . \square

We conclude with Theorem 3.10 by combining Theorems 5.7, 5.10, and 5.11.

5.3 From Non-simple Real Instances to Core Pairs

We now present our sequence of reductions moving from non-simple real Orbit instances to core pairs. The goal of this section is to establish the following result, which assumes that the only semilinear invariant for a core instances of dimension d is \mathbb{C}^d :

THEOREM 3.12. *Assuming Theorem 3.11, there is a polynomial time algorithm deciding whether a non-simple conjugated Orbit instance admits a semilinear invariant.*

Figure 5 recalls our pipeline of reductions; we now include detailed definitions of the classes of instances (and pairs) that we consider.

This section is broken into four parts that correspond to the four downwards arrows in Figure 5.

Ensuring aperiodicity. We say that two complex numbers are equivalent if their quotient or their product is a root of unity. We say that a matrix A is aperiodic if any two equivalent eigenvalues are in fact equal and if any eigenvalue that is a root of unity is in fact 1. Note that choosing m to be a common multiple to all orders of roots of unity that occur as eigenvalues or quotients or products of eigenvalues, we get that A^m is aperiodic.

We obtain the following reduction:

LEMMA 5.12. *There is a polynomial time reduction from a non-simple conjugated Orbit instance to many non-simple conjugated Orbit instances that are aperiodic.*

PROOF. Let (x, A, y) be a non-simple conjugated Orbit instance, and let m be such that A^m is aperiodic. For each $k \in \{0, \dots, m-1\}$, set $(x'_k, A'_k, y'_k) = (A^k x, A^m, y)$. By Lemma A.2, proved in Appendix A, m is indeed polynomial. We now prove that (x, A, y) admits a semilinear invariant if and only if for all k , (x'_k, A'_k, y'_k) does.

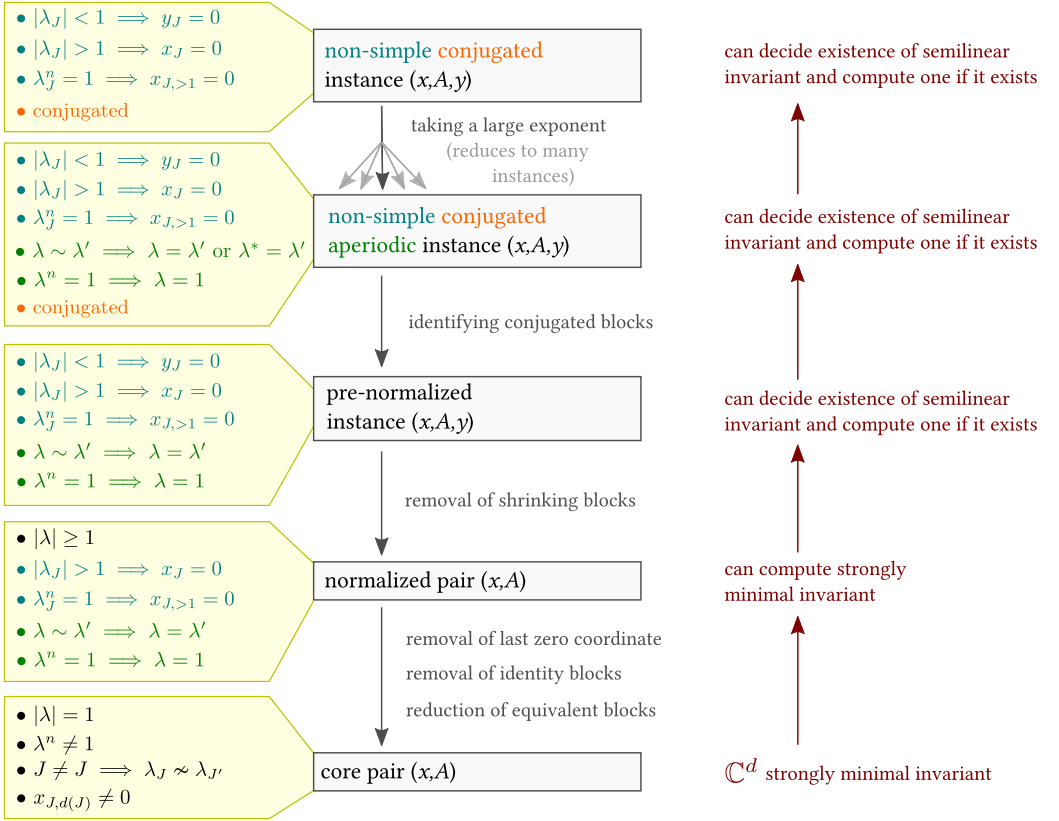


Fig. 5. Pipeline of reductions with detailed definitions.

Let \mathcal{I} be a semilinear invariant for (x, A, y) . Then, by an easy induction, $A^m \mathcal{I} \subseteq \mathcal{I}$, and clearly $A^k x \in \mathcal{I}$ for all k and $y \notin \mathcal{I}$ by assumption. Hence, for all $0 \leq k \leq m-1$, it holds that \mathcal{I} defines a semilinear invariant for (x'_k, A'_k, y'_k) .

Conversely, consider a family of respective semilinear invariants $(\mathcal{I}'_k)_{0 \leq k \leq m-1}$ for $(x'_k, A'_k, y'_k)_{0 \leq k \leq m-1}$. Then, for all k , let

$$\tilde{\mathcal{I}}_k = \{z \mid A^k z \in \mathcal{I}'_k\}.$$

Note that for any element $z \in \tilde{\mathcal{I}}_k$, we have $A^k z \in \mathcal{I}'_k$, thus $A^{m+k} z \in A^m \mathcal{I}'_k \subseteq \mathcal{I}'_k$ and therefore, $A^m z \in \tilde{\mathcal{I}}_k$. Hence, $\tilde{\mathcal{I}}_k$ is stable under A^m , and so the same holds for

$$\tilde{\mathcal{I}} = \bigcap_{k=0}^{m-1} \tilde{\mathcal{I}}_k.$$

Finally, we let

$$\mathcal{I} = \bigcup_{k=0}^{m-1} A^k \tilde{\mathcal{I}},$$

which we claim to be a semilinear invariant for ℓ . First, we have $x \in \tilde{\mathcal{I}}_k$ for each k and thus $x \in \mathcal{I}$. Second, the fact that \mathcal{I} is invariant under A follows directly from the fact that $\tilde{\mathcal{I}}$ is invariant under

A^m . Third, assume for contradiction that $y \in \mathcal{I}$: There is k such that $y \in A^k \tilde{\mathcal{I}} \subseteq A^k \tilde{\mathcal{I}}_k$. Then, $y = A^k z$ for some z such that $A^k z \in \mathcal{I}'_k$ and thus $y \in \mathcal{I}'_k$, a contradiction.

To conclude the proof of the lemma, we should argue that the resulting instances $(A^k x, A^m, y)$ are non-simple and conjugated. This requires an additional base change, since A^m is no longer in JNF; however, this base change preserves respective Jordan blocks and does not affect being non-simple or conjugated. We conclude by applying Lemma 5.2. \square

5.3.1 From Aperiodic Conjugated to Pre-normalized. We say that an orbit instance is pre-normalized if it is non-simple, equivalent eigenvalues are in fact equal, and eigenvalues that are roots of unity are in fact 1.

We proceed with the following reduction, which identifies synchronised conjugated Jordan blocks:

LEMMA 5.13. *There exists a polynomial time reduction from non-simple aperiodic conjugated Orbit instances to pre-normalized ones.*

PROOF. Consider a non-simple aperiodic conjugated Orbit instance $\ell = (x, A, y)$. Let S be the union of all Jordan blocks whose eigenvalues are either real or have positive imaginary part, and put $\ell' = (x', A', y') = (x_S, A_S, y_S)$. Note that ℓ' is non-simple, aperiodic, and such that no two different eigenvalues are conjugate; it is thus pre-normalized.

Let \mathcal{I} be a semilinear invariant for ℓ . Consider the semilinear set

$$\mathcal{I}' = \{z' \in \mathbb{C}^S \mid z \in \mathcal{I} \text{ where } z \text{ is such that for all block } J, \pi_{J^*}(z) = \pi_J(z)^*\}.$$

Since (x, A, y) is conjugated, we have $x' \in \mathcal{I}'$, and $y' \notin \mathcal{I}'$ otherwise, we would have $y \in \mathcal{I}$. Now, if $z' \in \mathcal{I}'$, then the vector z as in the definition belongs to \mathcal{I} , therefore, $Az \in \mathcal{I}$, and, since A is conjugated it follows that $(Az)_S = A'z'$ belongs to \mathcal{I}' .

Conversely, let \mathcal{I}' be a semilinear invariant for ℓ' and consider the semilinear set

$$\mathcal{I} = \{z \in \mathbb{C}^d \mid z_S \in \mathcal{I}' \text{ and for all block } J, \pi_{J^*}(z) = \pi_J(z)^*\}.$$

Since ℓ is conjugated, we have $x \in \mathcal{I}$, $y \notin \mathcal{I}$ and $A\mathcal{I} \subseteq \mathcal{I}$. \square

From pre-normalized instances to normalized pairs. We say that a pair (x, A) is normalized if

- all eigenvalues have modulus ≥ 1 ;
- blocks J such that $|\lambda_J| > 1$ satisfy $x_J = 0$;
- blocks J such that λ_J is a root of unity satisfy $x_{J, >1} = 0$;
- equivalent eigenvalues are equal; and
- eigenvalues that are roots of unity are in fact 1.

Thus, turning a pre-normalized Orbit instance (x, A, y) to a normalized pair (x', A') amounts to removing blocks whose eigenvalues have modulus < 1 ; we call these blocks shrinking. It turns out that normalized pairs admit strongly minimal invariants, which we now define.

Say that \mathcal{J} is a weak invariant for a pair $\ell = (x, A)$ if $A\mathcal{J} \subseteq \mathcal{J}$ and there exists n such that $A^n x \in \mathcal{J}$. A strongly minimal invariant \mathcal{I} for a pair $\ell = (x, A)$ is a semilinear invariant for ℓ (that is, $x \in \mathcal{I}$ and $A\mathcal{I} \subseteq \mathcal{I}$), which is contained in any semilinear weak invariants \mathcal{J} for ℓ . Note that strongly minimal invariants are always assumed to be semilinear. Note also that when such an invariant exists, it is unique.

The following lemma states the existence of a weak kind of reductions, which will turn out to be sufficient for our needs:

LEMMA 5.14. *Let (x, A, y) be a pre-normalized Orbit instance in jNF. There exists a normalized pair (x', A') , computable in polynomial time from (x, A, y) , such that given a strongly minimal invariant*

I' for (x', A') , one may decide whether (x, A, y) has a semilinear invariant in polynomial time, and in this case, compute one in polynomial time.

Note that Lemma 5.14 does not assert existence of strongly minimal invariants for normalized pairs; this will, however, follow from the rest of the proof. The proof makes use of Lemma 5.9 from the previous section.

PROOF. Let S be the union of all coordinates from shrinking blocks of A . Since (x, A, y) is non-simple, it holds that $y_S = 0$. We distinguish two cases. First, if $A_{S^c}^n x = y$ for infinitely many n 's, then, since $y_S = 0$, it follows that y belongs to the topological closure of the orbit $\{\overline{A^n x}, n \in \mathbb{N}\}$ and therefore there exists no closed semilinear invariant. Note that this can be tested in polynomial time: First test if y belongs to the orbit of x_{S^c} under A_{S^c} , and then test if some power of A_{S^c} is the identity matrix, which amounts to testing whether all eigenvalues of A_{S^c} are roots of unity.

So, we now assume that there is n_0 such that $y_{S^c} \notin \{A_{S^c}^n x_{S^c}, n \geq n_0\}$. First, we claim that $A_{S^c}^{n_0} x_{S^c}$, for some n_0 as above, can be computed in polynomial time. Indeed, if y_{S^c} does not belong to the orbit of x_{S^c} under A_{S^c} , then we may pick $n_0 = 0$, and if $y_{S^c} = A_{S^c}^n x_{S^c}$, then we take $n_0 = n + 1$ and compute $A_{S^c}^{n_0} x_{S^c} = A_{S^c} x_{S^c}$.

We let $(x', A') = (A_{S^c}^{n_0} x_{S^c}, A_{S^c})$; it is a pre-normalized pair. Let I' be a strongly minimal invariant for (x', A') .

If $y_{S^c} \notin I'$, it is a direct check that $\{x, Ax, \dots, A^{n_0-1}x\} \cup \pi_{S^c}^{-1}(I')$ is a semilinear invariant for (x, A, y) .

Otherwise, $y_{S^c} \in I'$, and we claim that in this case there exist no semilinear invariant for (x, A, y) . Towards a contradiction, consider such an invariant $I: x \in I, y \notin I$ and $AI \subseteq I$. Let $\epsilon = \frac{1}{2} \text{dist}(y, I)$ (recall that we compute distances with respect to the infinity norm).

For each shrinking block J , apply Lemma 5.9 to obtain a closed semilinear P_J satisfying $A_J P_J \subseteq P_J$ and $B(0, \frac{\epsilon}{2}) \subseteq P_J \subseteq B(0, \epsilon) \subseteq \mathbb{C}^J$. Let P_S be the Cartesian product of the P_J 's over shrinking blocks; we have $A_S P_S \subseteq P_S$ and $B(0, \frac{\epsilon}{2}) \subseteq P_S \subseteq B(0, \epsilon)$. Take $n \geq n_0$ large enough so $\|A_S^n x_S\| \leq \frac{\epsilon}{2}$, and hence $A_S^n x_S \in P_S$.

Let

$$\mathcal{J} = \{s \in I \mid z_S \in P_S\}.$$

By construction, $A^n x \in \mathcal{J}, A\mathcal{J} \subseteq \mathcal{J}$ and $y \notin \mathcal{J}$. Now, let $\mathcal{J}' = \pi_{S^c}(\mathcal{J})$. It holds that $A_{S^c}^n x_{S^c} \in \mathcal{J}'$ and $A_{S^c} \mathcal{J}' \subseteq \mathcal{J}'$: \mathcal{J}' is a weak invariant for (x', A') . Hence, since I' is strongly minimal, it holds that $I' \subseteq \mathcal{J}'$, and thus $y_{S^c} \in \mathcal{J}'$.

By definition, this means that there exists $z_S \in \mathbb{C}^S$ such that $z = (z_S, y_{S^c}) \in \mathcal{J}$, meaning that $z \in I$ and $z_S \in P$, which implies $\|z_S\| \leq \epsilon$. But then, since $y_S = 0$, we get

$$2\epsilon = \text{dist}(I, y) \leq \|z - y\| = \|z_S\| \leq \epsilon,$$

a contradiction. □

Before reducing normalized pairs to core instances, we introduce structured reductions.

Structured reductions. Define complex affine maps to be function $f: \mathbb{C}^d \rightarrow \mathbb{C}^{d'}$ of the form $f: z \mapsto Az + u$, where $A \in \mathbb{C}^{d' \times d}$ and $u \in \mathbb{C}^{d'}$. Let C be a class of pairs (x, A) in dimension d and C' be a class of pairs (x', A') in dimension d' . A structured reduction from C to C' is given by a function $R: C \rightarrow C'$ mapping a pair (x, A) to a pair $R(x, A) = (x', A')$, and for each pair $(x, A) \in C$, two complex affine maps $f: \mathbb{C}^{d'} \rightarrow \mathbb{C}^d$ and $g: \mathbb{C}^d \rightarrow \mathbb{C}^{d'}$ satisfying that

- $f(x') = x$;
- for all $z' \in \mathbb{C}^{d'}$, it holds that $g(f(z')) = z'$;
- for all $z' \in \mathbb{C}^{d'}$, it holds that $f(A'z') = Af(z')$.

It is easy to see that the compositions of structured reductions are structured reductions.

Note that for a complex affine map f and a semilinear \mathcal{I} , both $f(\mathcal{I})$ and $f^{-1}(\mathcal{I})$ are semilinear. We now prove that structured reductions indeed give rise to reductions (as defined in Section 5.1), and that, moreover, they reflect strongly minimal invariants.

LEMMA 5.15. *Consider a structured reduction between classes of pairs C and C' . Then, the following properties hold for all pair $(x, A) \in C$:*

- For all semilinear invariants \mathcal{I} for (x, A) , it holds that $f^{-1}(\mathcal{I})$ is a semilinear invariant for $R(x, A)$.
- For all semilinear invariants \mathcal{I}' for $R(x, A)$, it holds that $f(\mathcal{I}')$ is a semilinear invariant for (x, A) .
- If \mathcal{I}' is a strongly minimal invariant for $R(x, A)$, then $f(\mathcal{I}')$ is a strongly minimal invariant for (x, A) .

Stated differently, if there is a structured reduction from C to C' , then for all $(x, A) \in C$, (x, A) admits a semilinear invariant if and only if $R(x, A)$ does. Moreover, if $R(x, A)$ admits a strongly minimal invariant, then so does (x, A) . We now prove Lemma 5.15.

PROOF. Let us consider a pair $(x, A) \in C$ and let us write $(x', A') = R(x, A) \in C'$.

Let \mathcal{I} be a semilinear invariant for (x, A) , we show that $f^{-1}(\mathcal{I})$ is a semilinear invariant for (x', A') . First, $f(x') = x \in \mathcal{I}$ so $x' \in f^{-1}(\mathcal{I})$. Second, let $z' \in f^{-1}(\mathcal{I})$, meaning $f(z') \in \mathcal{I}$. We want to show that $A'z' \in f^{-1}(\mathcal{I})$. By assumption $f(A'z') = Af(z')$, and because \mathcal{I} is stable under A , this implies that $f(A'z') \in \mathcal{I}$, thus $A'z' \in f^{-1}(\mathcal{I})$.

Let \mathcal{I}' be a semilinear invariant for (x', A') , we show that $f(\mathcal{I}')$ is a semilinear invariant for (x, A) . First, $f(x') = x$ and $x' \in \mathcal{I}'$, so $x \in f(\mathcal{I}')$. Second, let $z \in f(\mathcal{I}')$, meaning $f(z') = z$ for some $z' \in \mathcal{I}'$. We want to show that $Az \in f(\mathcal{I}')$. By assumption, $f(A'z') = Af(z') = Az$, and because \mathcal{I}' is stable under A' , we have $A'z' \in \mathcal{I}'$, implying that $Az \in f(\mathcal{I}')$.

Let \mathcal{I}' be a strongly minimal semilinear invariant for (x', A') , we claim that $f(\mathcal{I}')$ is a strongly minimal semilinear invariant for (x, A) . Let \mathcal{J} be a semilinear weak invariant for (x, A) . As we proved above $f^{-1}(\mathcal{J})$ is stable under A' , and we obtain by an easy induction on n that $f(A'^n x') = A^n x$. Thus, since $A^n x \in \mathcal{J}$ for some n , it holds that $A'^n x' \in f^{-1}(\mathcal{J})$ for the same n . Hence, $f^{-1}(\mathcal{J})$ is a semilinear weak invariant for (x', A') . By minimality of \mathcal{I}' , we have $\mathcal{I}' \subseteq f^{-1}(\mathcal{J})$. This implies that $f(\mathcal{I}') \subseteq f(f^{-1}(\mathcal{J})) = \mathcal{J}$, as required. \square

From normalized to core pairs. Recall that (x, A) is a core pair if:

- all eigenvalues have modulus 1;
- no eigenvalue is a root of unity;
- two different blocks have non-equivalent eigenvalues;
- the last coordinate $x_{J, d(J)}$ of x on each block is $\neq 0$.

To reduce from normalized pairs to core pairs, we will use three structured reductions:

- (1) removal of the last coordinate $(J, d(J))$ of a block J assuming $x_{J, d(J)} = 0$;
- (2) removal of diagonal blocks with eigenvalue one (called identity blocks);
- (3) removal of a coordinate assuming there exists two different block J_1, J_2 associated with the same eigenvalue and $x_{J_2, d(J_2)} \neq 0$.

These reductions are detailed below. To obtain a core pair, we proceed as follows:

- We start by applying the first reduction repeatedly to remove all blocks J associated with eigenvalues of modulus > 1 ; this is possible, since in normalized pairs, such blocks satisfy $x_J = 0$.

Calculating $Af(z')$, we obtain the same results as for $f(A'z')$:

$$\begin{array}{llll} \text{for } i < d_1 - 1 & (f(A'z'))_{J_1, i} & = \lambda z'_{J_1, i} + z'_{J_1, i+1} & = (Af(z'))_{J_1, i}, \\ & (f(A'z'))_{J_1, d_1-1} & = \lambda z'_{J_1, d_1-1} + z'_{J_2, d_2} & = (Af(z'))_{J_1, d_1-1}, \\ & (f(A'z'))_{J_1, d_1} & = \mu \lambda z'_{J_2, d_2} & = (Af(z'))_{J_1, d_1}, \\ \text{for } i < d_2 & (f(A'z'))_{J_2, i} & = \lambda z'_{J_2, i} + z'_{J_2, i+1} & = (Af(z'))_{J_2, i}, \\ & (f(A'z'))_{J_2, d_2} & = \lambda z'_{J_2, d_2} & = (Af(z'))_{J_2, d_2}. \end{array}$$

Note that the obtained matrix A' is indeed one dimension lower, but it is not in JNF. To prove the lemma, there remains to apply a base change and reduce once more, to $(x'', A'') = (Q^{-1}x', Q^{-1}A'Q)$. It is a direct check that this defines a structural reduction, with affine maps $f : z \mapsto Qz$ and $g : z \mapsto Q^{-1}z$. Last, we should argue that (x'', A'') is strongly normalized: This simply follows from the facts that strong normality depends only on the set of eigenvalues of A (without multiplicity) and that A, A' , and A'' have the same set of eigenvalues. \square

Combining Lemmas 5.16, 5.17, and 5.18, as explained above, we obtain the following result:

LEMMA 5.19. *There exists a structured reduction from normalized pairs to core pairs.*

5.3.2 *Proof of Theorem 3.12.* We are finally ready to prove Theorem 3.12, which we first recall for convenience.

THEOREM 3.12. *Assuming Theorem 3.11, there is a polynomial time algorithm deciding whether a non-simple conjugated Orbit instance admits a semilinear invariant.*

PROOF. Consider given a non-simple conjugate Orbit instance (x, A, y) . Apply the following reductions:

$$(x, A, y) \xrightarrow{\text{Lemma 5.12}} (x_1, A_1, y_1) \xrightarrow{\text{Lemma 5.13}} (x_2, A_2, y_2) \xrightarrow{\text{Lemma 5.14}} (x_3, A_3) \xrightarrow{\text{Lemma 5.19}} (x_4, A_4).$$

By Theorem 3.11, the only semilinear invariant for the core pair (x_4, A_4) is C^{d_4} ; it is a strongly minimal invariant. Since Lemma 5.19 provides a structured reduction, this gives a strongly minimal invariant I_3 for the normalized pair (x_3, y_3) . At this point, Lemma 5.14 allows to decide whether (x_2, A_2, y_2) admits a semilinear invariant, and if it does, construct such an invariant I_2 . Then, Lemma 5.13 gives the same conclusion for (x_1, A_1, y_1) , and then Lemma 5.12 concludes. \square

6 Decidability Proof: Core Instances

A matrix $A \in \mathbb{C}^{d \times d}$ is a core matrix if it is in Jordan normal form, eigenvalues of A have modulus 1, are not roots of unity, and eigenvalues associated to different Jordan blocks are non-equivalent. Recall that a pair (x, A) is a core pair if A is a core matrix and the last coordinate of x on each Jordan block is non-zero.

Fix a core matrix A . Say that a set is a basic invariant if it is of the form

$$\prod_{J \in \mathcal{J}} \mathbb{C}^{p_J} \times \{0\}^{d(J)-p_J},$$

where \mathcal{J} stands for the set of Jordan blocks of A , and for each $J \in \mathcal{J}$, p_J is an integer in $[0, d(J)]$. Note that basic invariants are indeed invariant under A , closed, and semilinear.

The goal of this section is to prove that they are the only possible invariants.

THEOREM 6.1. *Let A be a core matrix. Then, all closed semilinear sets that are invariant under A are unions of basic invariants.*

Observe that if (x, A) is a core pair, then the last coordinate of x on each block is non-zero, hence the only basic invariant containing x is \mathbb{C}^d . Thus, Theorem 3.11 follows from Theorem 6.1.

6.1 Dimension of Invariants for Core Pairs

We start by establishing that invariants for core pairs have full dimension (Lemma 6.8 below). First, some definitions and some technical results.

The dimension of a set X of \mathbb{R}^d , which we denote by $\dim(X)$, is the minimal k in \mathbb{N} such that X is included in a finite union of affine subspaces of dimension at most k . The dimension of a set X in \mathbb{C}^d , which we denote by $\dim_{\mathbb{R}}(X)$, is the dimension of $X_{\mathbb{R}}$. The following lemma is a standard result about semilinear sets:

LEMMA 6.2 (DIMENSION OF SEMILINEAR SETS). *Let \mathcal{I} be a semilinear set in \mathbb{R}^d . If it has empty interior, meaning $\mathcal{I}^o = \emptyset$, then \mathcal{I} has dimension at most $d - 1$.*

COROLLARY 6.3. *For any semilinear set \mathcal{I} in \mathbb{R}^d , $\partial\mathcal{I}$ has dimension at most $d - 1$.*

We will make use of the following powerful theorem about linear recurrence sequences. This result is due to Skolem [34], and more general versions were subsequently obtained by Mahler [25, 26] and Lech [24]. This result can also be found in the recent monograph of Everest et al. [11, Theorem 2.1 and subsequent discussion]. Recall that \mathbb{U} denotes the set of roots of unity and S^1 the complex unit circle. Recall that a linear recurrence sequence is *degenerate* if some quotient of two distinct roots of its characteristic polynomial is a root of unity.

THEOREM 6.4 (SKOLEM, MAHLER, LECH). *Let $(u_n)_{n \in \mathbb{N}}$ be a real non-degenerate linear recurrence sequence. Then, $\{n \in \mathbb{N} : u_n = 0\}$ is either finite or all of \mathbb{N} .*

We will also require the following technical lemmas:

LEMMA 6.5. *Let $A \in \mathbb{C}^{d \times d}$ be in Jordan normal form, \mathcal{J} its Jordan blocks, and let λ_J denote the eigenvalue of each Jordan block $J \in \mathcal{J}$. Let $x \in \mathbb{C}^d$ and $M = [x_{\mathbb{R}} \quad (Ax)_{\mathbb{R}} \quad \cdots \quad (A^{2d-1}x)_{\mathbb{R}}]$. If all the eigenvalues of $A_{\mathbb{R}}$ are distinct and $x_{J, d(J)} \neq 0$ for all $J \in \mathcal{J}$, then $\det(M) \neq 0$.*

PROOF. See Appendix B. □

LEMMA 6.6. *Let $A \in \mathbb{C}^{d \times d}$ be a core matrix and \mathcal{J} range over its Jordan blocks. There exists a change of basis P that stabilizes any basic invariant, and such that $PA^{-1}P^{-1} = \text{Diag}(\mathcal{J}_{d(J)}(\lambda_J^{-1}), J \in \mathcal{J})$.*

PROOF. See Appendix C. □

LEMMA 6.7. *Let (x, A) be a core pair. Then, for any vector $v \in \mathbb{R}^{2d} \setminus \{0\}$, $v^T(A^n x)_{\mathbb{R}}$ is zero for finitely many n .*

PROOF. Let $u_n = v^T(A^n x)_{\mathbb{R}}$, which is a real linear recurrence sequence. Furthermore, the roots of the characteristic polynomial of $(u_n)_n$ are the eigenvalues of $A_{\mathbb{R}}$. It is not hard to see that the eigenvalues of $A_{\mathbb{R}}$ are $\lambda_1, \dots, \lambda_s, \lambda_1^*, \dots, \lambda_s^*$ so in particular the quotients of any two distinct such eigenvalues are of the forms

$$\frac{\lambda_i}{\lambda_j}, \quad \frac{\lambda_i}{\lambda_j^*} = \lambda_i \lambda_j, \quad \frac{\lambda_i^*}{\lambda_j} = \frac{1}{\lambda_i \lambda_j}, \quad \frac{\lambda_i^*}{\lambda_j^*} = \frac{\lambda_j}{\lambda_i^*},$$

none of which are roots of unity by our assumptions (recall that for complex number z of modulus 1, $z^* = z^{-1}$). We can now apply Theorem 6.4 to conclude that either $u_n = 0$ for all n or there are only finitely many n such that $u_n = 0$. We will show that the former case implies that $v = 0$, which is excluded.

Assume that $u_n = 0$ for all n . In particular, $v^T(A^n x)_{\mathbb{R}} = 0$ for all $n \in \{0, 1, \dots, 2d - 1\}$. Hence, v is in the kernel of $M = [x_{\mathbb{R}} \quad (Ax)_{\mathbb{R}} \quad \cdots \quad (A^{2d-1}x)_{\mathbb{R}}]$. But, by our assumptions, the eigenvalues of A and their conjugates are all distinct (if two were equal, then their product or quotient would

be 1, hence a root of unity). Furthermore, $x_{J,d(J)} \neq 0$ for all J . It follows by Lemma 6.5 that M is invertible so $v = 0$. \square

We now prove that invariants for core pairs have full dimension.

LEMMA 6.8. *Let \mathcal{I} be a non-empty closed semilinear invariant for a core pair (x, A) . Then, \mathcal{I} has full-dimension, i.e., $\mathcal{I}_{\mathbb{R}}$ has dimension $2d$.*

PROOF. Let $m = \dim_{\mathbb{R}}(\mathcal{I})$ and assume, toward contradiction, that $m < 2d$. Then, \mathcal{I} is contained into the union of finitely many affine subspaces of dimension m :

$$\mathcal{I}_{\mathbb{R}} \subseteq \bigcup_{j=1}^p F_j,$$

where for all j , $F_j \subseteq \mathbb{R}^{2d}$ is a real affine subspace of dimension m . For all $n \in \mathbb{N}$, $(A^n x)_{\mathbb{R}} \in \mathcal{I}_{\mathbb{R}}$, since \mathcal{I} is invariant under A , so there exists $j_n \in [1, p]$ such that $(A^n x)_{\mathbb{R}} \in F_{j_n}$. Hence, there must be some j_{∞} such that $F_{j_{\infty}}$ contains $(A^n x)_{\mathbb{R}}$ for infinitely many values of n . Since $F_{j_{\infty}}$ has dimension $m < 2d$, it is contained in some hyperplane $H = \{y \in \mathbb{R}^{2d} : y^T v = 0\}$ of normal $v \in \mathbb{R}^{2d} \setminus \{0\}$. Therefore, $v^T (A^n x)_{\mathbb{R}} = 0$ for infinitely many n 's, contradicting Lemma 6.7. \square

6.2 The Diagonal Case

We now deal with the case where the matrix A is diagonal. This case is important for two reasons. First, it plays the role of the base case in our general induction. Second, it is also used as a technical tool in the general case to rule out certain scenarios.

LEMMA 6.9. *Let A be a diagonal core matrix, and let \mathcal{I} be a non-empty closed semilinear set invariant under A , which, moreover, contains a point $x \in \mathcal{I}$ that is nonzero on each coordinate. Then, $\mathcal{I} = \mathbb{C}^d$.*

PROOF. We show a few facts:

- (i) \mathcal{I} must have full dimension (i.e., real dimension $2d$),
- (ii) $\partial\mathcal{I}$ is invariant under A ,
- (iii) if $\partial\mathcal{I}$ is non-empty (that is, if $\mathcal{I} \neq \mathbb{C}^d$), then it contains a point that is nonzero on each coordinate.

This implies the desired result: If towards contradiction we had that $\mathcal{I} \neq \mathbb{C}^d$, then $\mathcal{I}' := \partial\mathcal{I}$ would be a non-empty closed semilinear set invariant under A , thanks to (ii), and it would contain a point that is nonzero on each coordinate, thanks to (iii). Therefore, we could apply the same reasoning to \mathcal{I}' , which would satisfy the above points as well and have full dimension, thanks to (i), which contradicts Corollary 6.3.

(i) This is proved by Lemma 6.8.

- (ii) Since multiplication is continuous, $A\mathcal{I}^c \subseteq \mathcal{I}^c$ implies $A\overline{\mathcal{I}^c} \subseteq \overline{\mathcal{I}^c}$. Now, since \mathcal{I} is closed, we have $\partial\mathcal{I} = \mathcal{I} \cap \overline{\mathcal{I}^c}$ and therefore, since \mathcal{I} is invariant under A , it suffices to show that \mathcal{I}^c is invariant under A .

We now show that \mathcal{I}^c is invariant under A . This amounts to proving that \mathcal{I} is invariant under A^{-1} . Let x in \mathcal{I} and

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

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

$$T_A = \left\{ \text{Diag}(\mu_1, \dots, \mu_d) : \mu \in S^d \text{ and } \forall v \in L_A (\mu_1^{v_1} \cdots \mu_d^{v_d} = 1) \right\}$$

whose diagonal entries satisfy the multiplicative relations in L_A . Using Kronecker's Theorem on inhomogeneous simultaneous Diophantine approximation [5], it is shown in Reference [32, Proposition 3.5] that $\{A^n : n \in \mathbb{N}\}$ is a dense subset of T_A . This implies that

$$\overline{\{A^n x : n \in \mathbb{N}\}} = \{Mx : M \in T_A\} .$$

Since x is in \mathcal{I} and \mathcal{I} is invariant under A , we have that $\overline{\{A^n x : n \in \mathbb{N}\}} \subseteq \overline{\mathcal{I}} = \mathcal{I}$. Now, observe that $A^{-1} = \text{Diag}(\lambda_1^{-1}, \dots, \lambda_d^{-1})$ is in T_A , and thus $A^{-1}x$ is in \mathcal{I} .

- (iii) Assume that $\partial\mathcal{I} \neq \emptyset$. Let $\mathcal{Q} = \bigcup_{i=1}^d \mathbb{C}^{i-1} \times \{0\} \times \mathbb{C}^{d-i}$ be the set of points with at least one zero coordinate. Note that \mathcal{Q} is closed. Observe that \mathcal{Q}^c is path-connected (this follows from applying coordinate-wise the fact that $\mathbb{C} \setminus \{0\}$ is path-connected). Note that, since \mathcal{I} is closed, $\mathbb{C}^d = \mathcal{I}^o \cup \partial\mathcal{I} \cup \mathcal{I}^c$ where the union is disjoint. Assume for contradiction that $\partial\mathcal{I} \subseteq \mathcal{Q}$. Then,

$$\mathcal{Q}^c = \mathbb{C}^d \setminus \mathcal{Q} = (\mathbb{C}^d \setminus \partial\mathcal{I}) \setminus \mathcal{Q} = (\mathcal{I}^o \cup \mathcal{I}^c) \setminus \mathcal{Q} = (\mathcal{I}^o \setminus \mathcal{Q}) \cup (\mathcal{I}^c \setminus \mathcal{Q}).$$

Now, $\mathcal{I}^o \setminus \mathcal{Q}$ is open and it is non-empty, because \mathcal{I} contains a point $x \notin \mathcal{Q}$ by assumption so $x \in \mathcal{I} \setminus \mathcal{Q} = \mathcal{I}^o \setminus \mathcal{Q}$, since $\partial\mathcal{I} \subseteq \mathcal{Q}$. Similarly, $\mathcal{I}^c \setminus \mathcal{Q}$ is open (\mathcal{I} is closed) and non-empty, because otherwise $\mathcal{I}^c \subseteq \mathcal{Q}$ so $\mathcal{Q}^c \subseteq \mathcal{I}$, hence $\mathbb{C}^d = \overline{\mathcal{Q}^c} \subseteq \overline{\mathcal{I}} = \mathcal{I}$, which implies that $\partial\mathcal{I} = \emptyset$, contrary to our assumption. Therefore, \mathcal{Q}^c is the disjoint union of two non-empty open sets, hence disconnected, a contradiction. \square

We may easily deduce that Theorem 6.1 holds in the diagonal case (which corresponds to having Jordan blocks of size 1).

THEOREM 6.10. *Let $A \in \mathbb{C}^{d \times d}$ be a diagonal core matrix. Closed semilinear invariant sets that are invariant for A are of the form $\prod_{i=1}^d \varepsilon_i$, where $\varepsilon_i \in \{\{0\}, \mathbb{C}\}$.*

PROOF. We show that for any $x \in \mathcal{I}$, \mathcal{I} must contain $\prod_i \varepsilon_i$, with $\varepsilon_i = \begin{cases} \{0\} & \text{if } x_i = 0 \\ \mathbb{C} & \text{otherwise} \end{cases}$, which implies the result. This follows directly from applying Lemma 6.9 to the projection of $\mathcal{I} \cap \prod_i \varepsilon_i$ on coordinates $\{i \in [1, d] : x_i \neq 0\}$. \square

6.3 General Case

We now work with a general (not necessarily diagonal) core matrix A ; as usual, we let \mathcal{J} range over the Jordan blocks of A and let $s = |\mathcal{J}|$. The proof of Theorem 6.1 will proceed by induction on d . Since it involves several nontrivial steps, we explicitly spell out the induction hypothesis.

(HR_d) The only closed semilinear invariants for core matrices of dimension d are unions of basic invariants.

We let last denote the set of last coordinates of Jordan blocks of A :

$$\text{last} = \{(J, d(J)) \mid J \in \mathcal{J}\}.$$

Recall that $\pi_S : \mathbb{C}^d \rightarrow \mathbb{C}^S$ denotes the projection on a given set of coordinates S .

We start with the intuition. Let \mathcal{I} be a semilinear set that is invariant under A . We will project \mathcal{I} on the last coordinate of each block (using π_{last}). Since A acts diagonally on these coordinates, this projection is invariant under a diagonal matrix so we may apply Theorem 6.10 and decompose $\pi_{\text{last}}(\mathcal{I})$ as above. Assuming that $\pi_{\text{last}}(\mathcal{I})$ is not the whole set \mathbb{C}^S , then some of its components are identically zero, which allows us to reduce the dimension and conclude by induction.

LEMMA 6.11. *Let \mathcal{I} be a closed semilinear set that is invariant under a core matrix A of dimension d . If ($HR_{d'}$) holds for all $d' < d$, then either \mathcal{I} is a union of basic invariants or $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^S$.*

Note that the only basic invariant that has full dimension is \mathbb{C}^d ; therefore, the conclusion of the lemma implies that if \mathcal{I} has full dimension, then $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^s$.

PROOF. Let $\lambda_1, \dots, \lambda_s$ be the eigenvalues of A associated with the Jordan blocks J_1, \dots, J_s . Let \mathcal{I} be a semilinear set invariant under A and consider $\mathcal{I}' = \pi_{\text{last}}(\mathcal{I}) \subseteq \mathbb{C}^{\text{last}}$, the projection of \mathcal{I} on the last coordinate of each block. We identify \mathbb{C}^{last} with \mathbb{C}^s by identifying the coordinate $(J_i, d(J_i))$ with i . Since $(Ax)_{J_i, d(J_i)} = \lambda_i x_{J_i, d(J_i)}$ and \mathcal{I} is invariant under A , \mathcal{I}' is invariant under $B = \text{Diag}(\lambda_1, \dots, \lambda_s)$.

Observe that B is a core matrix, therefore, we may apply Theorem 6.10 to \mathcal{I}' and B ; it follows that $\mathcal{I}' = \bigcup_{\ell=1}^k \mathcal{I}'_{\ell}$ for some k , where $\mathcal{I}'_{\ell} = \prod_{J \in \mathcal{J}} \varepsilon_{\ell, J}$ and $\varepsilon_{\ell, J} \in \{\{0\}, \mathbb{C}\}$. Therefore, we have $\mathcal{I} = \bigcup_{\ell=1}^k \mathcal{I}_{\ell}$, where $\mathcal{I}_{\ell} = \mathcal{I} \cap \pi_{\text{last}}^{-1}(\mathcal{I}'_{\ell})$. Furthermore, it is not hard to check that $A\pi_{\text{last}}^{-1}(\mathcal{I}'_{\ell}) = \pi_{\text{last}}^{-1}(\mathcal{I}'_{\ell})$ given the special form of \mathcal{I}'_{ℓ} . It follows that \mathcal{I} and $\pi_{\text{last}}^{-1}(\mathcal{I}'_{\ell})$ are invariant under A so \mathcal{I}_{ℓ} is invariant under A . Also note that, since $\mathcal{I}'_{\ell} \subseteq \mathcal{I}'$, we have that $\pi_{\text{last}}(\mathcal{I}_{\ell}) = \mathcal{I}'_{\ell}$. Therefore, it now suffices to prove the result for each \mathcal{I}_{ℓ} to prove (HR_d) . Hence, we now assume that $\pi_{\text{last}}(\mathcal{I}) = X$ for some set $X = \prod_{J \in \mathcal{J}} \varepsilon_{\ell, J}$ as above. If $X = \mathbb{C}^s$, then the lemma holds.

Now, assume that $X \neq \mathbb{C}^s$. This means that there exists J such that $\varepsilon_{\ell, J} = \{0\}$. In particular, $X \subseteq \{z \in \mathbb{C}^s : z_{J, d(J)} = 0\}$ and therefore $\mathcal{I} \subseteq \pi_{\text{last}}^{-1}(X) \subseteq \{z \in \mathbb{C}^d : z_{J, d(J)} = 0\} =: P_J$. Let $p = \pi_{(J, d(J))^c}$ be the projection on all coordinates but $(J, d(J))$. Intuitively, \mathcal{I} is identically 0 on the coordinate $(J, d(J))$ so projecting it away (via p) and then pulling-back (via p^{-1}) and setting $(J, d(J))$ to zero (i.e., intersect with P_J) yields the same set. Formally, since $\mathcal{I} \subseteq P_J$, we have that $p^{-1}(p(\mathcal{I})) \cap P_J = \mathcal{I}$. Furthermore, P_J is invariant under A (since $(J, d(J))$ is the last coordinate of the block and it is zero) so for any set $X \subseteq P_J$, $p(AX) = Bp(X)$ where $B := A_{(J, d(J))^c}$. Hence, $Bp(\mathcal{I}) = p(A\mathcal{I}) = p(\mathcal{I})$ so $p(\mathcal{I})$ is invariant under B . But now, B has dimension $d - 1$, is a core matrix and $p(\mathcal{I})$ is a semilinear set invariant under B . Hence, by (HR_{d-1}) , $p(\mathcal{I})$ is a union of sets of the form $\prod_{J' \in \mathcal{J}'} \mathbb{C}^{p_{J'}} \times \{0\}^{d(J')-p_{J'}}$, where \mathcal{J}' is the set of Jordan blocks of A and $p_{J'}$ are some integers. By pulling back through p^{-1} as explained above, we get that \mathcal{I} is a union of sets of the form

$$p^{-1} \left(\prod_{J' \in \mathcal{J}'} \mathbb{C}^{p_{J'}} \times \{0\}^{d(J')-p_{J'}} \right) \cap P_J = \left(\prod_{J' \in \mathcal{J} \setminus \{J\}} \mathbb{C}^{p_{J'}} \times \{0\}^{d(J')-p_{J'}} \right) \times \left(\mathbb{C}^{p_J} \times \{0\}^{d(J)-1-p_J} \times \{0\} \right),$$

since $p^{-1}(\cdot) \cap P_J$ leaves all Jordan block unchanged except for J where it adds one component, which is 0. This shows that (HR_d) hold for \mathcal{I} . \square

Overview of the remainder of the proof. We now focus on the case where $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^{\text{last}}$, which is the difficult case. The remainder of the proof proceeds in two steps, which we now roughly describe.

- First, we establish that \mathcal{I} contains the set \mathcal{Q} of points that are zero on the last coordinate of each block (Lemma 6.11). This goes through a careful examination of the behavior of the second-to-last coordinate (Lemma 6.12).
- Then, we will describe the structure of \mathcal{I} in the close neighborhood of \mathcal{Q} . Assuming that $\mathcal{I} \neq \mathbb{C}^d$ and applying the previous point yields that $\mathcal{Q} \subseteq \mathcal{I}$ and also $\mathcal{Q} \subseteq \overline{\mathcal{I}^c}$. This will allow us to obtain a precise understanding of the shape of \mathcal{I} in the neighborhood of \mathcal{Q} , which eventually leads to a contradiction.

We now proceed with the first step. Lemma 6.12 shows that if each block has size 1 or 2, then \mathcal{I} contains an element that is 0 on the last coordinate of each block but nonzero on the *second last* coordinate (of each block of size 2). The intuition is as follows: Let J be a block such that $d(J) = 2$ and let λ be its eigenvalue. If $z \in \mathcal{I}$ is such that $z_{J,2} \neq 0$, then for all $k \in \mathbb{N}$,

$$(A^k z)_{J,1} = \lambda^k z_{J,1} + k\lambda^{k-1} z_{J,2}.$$

Now, recall that $|\lambda| = 1$ so $\lambda^k z_{J,1}$ has constant modulus while $k\lambda^{k-1}z_{J,2}$ diverges to infinity in norm, since we took $z_{J,2} \neq 0$. Essentially, this means that by carefully choosing k , we can ensure that $(A^k z)_{J,1}$ belongs to some “donut,” that is, bounded away from 0 but not too far away from the origin either. In other words, the orbit of z under A (which is contained in \mathcal{I}) always intersects a set K , which is essentially a donut on the second-to-last coordinates of each block. If we now consider a sequence of points z_n as above and make sure $z_n \rightarrow 0$ with nonzero last coordinates (which is possible, since $\pi_{\text{last}}(\mathcal{I})$ contains a ball around 0), then we can make sure that the orbit of each z_n intersects the *same* set K . Since the donut is compact, this means that we can find a converging subsequence and, since $z_n \rightarrow 0$, this limit will be 0 on the last coordinate of each block but nonzero on the second-to-last of each block because of the definition of the K . The technical aspects of the proof lies in how we choose z_n and how we ensure that the set K is the same for all n .

LEMMA 6.12. *Let A be a core matrix of dimension d with $d(J) \in \{1, 2\}$ for all $J \in \mathcal{J}$ and let \mathcal{I} be a closed semilinear set that is invariant under A and such that $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^s$. Then, there exists $z \in \mathcal{I}$ such that $\pi_{\text{last}}(z) = 0$ and for all $J \in \mathcal{J}$, if $d(J) = 2$ then $z_{J,1} \neq 0$.*

PROOF. Let x_n be a sequence of non-zero complex numbers of modulus at most 1, that is, a decreasing (in modulus) and converges to 0. Let $y^{(n)} = (x_n, \dots, x_n) \in \mathbb{C}^s$. Since $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^s$, $y^{(n)} \in \pi_{\text{last}}(\mathcal{I})$ for all n , so the section $\{x \in \mathcal{I} : \pi_{\text{last}}(x) = y^{(n)}\}$ is non-empty. By Lemma 2.3, since $y^{(n)}$ has norm less than 1, there is some B such that for all n , there exists $z^{(n)} \in \mathcal{I}$ of norm at most B such that $\pi_{\text{last}}(z^{(n)}) = y^{(n)}$ for all n . Since the $z^{(n)}$ are bounded in norm, without loss of generality, we can assume that they converge to some $z^{(\infty)}$ by extracting a subsequence. Since \mathcal{I} is closed, it is the case that $z^{(\infty)} \in \mathcal{I}$ and by continuity, $\pi_{\text{last}}(z^{(\infty)}) = \lim_{n \rightarrow \infty} \pi_{\text{last}}(z^{(n)}) = \lim_{n \rightarrow \infty} y^{(n)} = 0$. Let $\mathcal{J}_2 = \{J \in \mathcal{J} : d(J) = 2\}$ and $\mathcal{J}' = \{J \in \mathcal{J}_2 : z_{J,1}^{(\infty)} \neq 0\}$. Now, let

$$\delta = \min \left(1, \min \left\{ \left| z_{J,1}^{(\infty)} \right| : J \in \mathcal{J}' \right\} \right) > 0.$$

(In the case where $\mathcal{J}' = \emptyset$, we have $\delta = 1$.)

Let n be large enough so $\|z^{(n)} - z^{(\infty)}\| \leq \delta/4$. Since $\pi_{\text{last}}(z^{(\infty)}) = 0$, we have $|x_n| = \|y_n\| = \|\pi_{\text{last}}(z^{(n)})\| \leq \delta/4$. Then, for any $J \in \mathcal{J}$, and $k \in \mathbb{N}$, using that the eigenvalue λ_J of J has modulus 1,

$$\left| \left(A^k z^{(n)} \right)_{J, d(J)} \right| = \left| \lambda_J^k z_{J, d(J)}^{(n)} \right| = |x_n| \leq \frac{\delta}{4}. \quad (3)$$

Let $k \in \mathbb{N}$ and $J \in \mathcal{J}_2$, then

$$\left(A^k z^{(n)} \right)_{J,1} = \lambda_J^k \left(z_{J,1}^{(n)} + k\lambda_J^{-1}z_{J,2}^{(n)} \right) = \lambda_J^k \left(z_{J,1}^{(n)} + k\lambda_J^{-1}x_n \right).$$

Let $k(n) = \left\lceil \frac{\delta}{2|x_n|} \right\rceil$. Then, for all $J \in \mathcal{J}$,

$$\left| \left(A^{k(n)} z^{(n)} \right)_{J,1} \right| \leq \left| z_{J,1}^{(n)} \right| + k(n) |x_n| \leq \left| z_{J,1}^{(\infty)} \right| + \frac{\delta}{4} + \left(\frac{\delta}{2|x_n|} + 1 \right) |x_n| \leq \delta + \left| z_{J,1}^{(\infty)} \right|. \quad (4)$$

We now make a case analysis on $J \in \mathcal{J}_2$:

– If $J \in \mathcal{J}'$, then

$$\left| \left(A^{k(n)} z^{(n)} \right)_{J,1} \right| \geq \left| z_{J,1}^{(n)} \right| - k(n) |x_n| \geq \delta - \left(\frac{\delta}{2|x_n|} + 1 \right) |x_n| \geq \frac{\delta}{4} \quad (5)$$

by the definition of δ .

– If $J \in \mathcal{J}_2 \setminus \mathcal{J}'$, then $z_{J,1}^{(\infty)} = 0$ so $|z_{J,1}^{(n)}| \leq \delta/4$ and

$$\left| \left(A^{k(n)} z^{(n)} \right)_{J,1} \right| \geq k(n) |x_n| - |z_{J,1}^{(n)}| \geq \frac{\delta}{2|x_n|} |x_n| - \frac{\delta}{4} \geq \frac{\delta}{2}. \quad (6)$$

Now, \mathcal{I} being invariant under A , the sequence $(A^{k(n)} z^{(n)})_n$ has its elements in \mathcal{I} , and ultimately lies in the compact set

$$K = \left\{ u \in \mathbb{C}^d : \forall J \in \mathcal{J}, |u_{J,d(J)}| \leq \frac{\delta}{4} \quad \text{and} \quad \forall J \in \mathcal{J}_2, \frac{\delta}{4} \leq |u_{J,1}| \leq |z_{J,1}^{(\infty)}| + \delta \right\},$$

thanks to Equations (3),(4), (5), and (6). We may then extract a converging subsequence in K , with its limit $u^{(\infty)}$ in $\mathcal{I} \cap K$. Clearly, $\pi_{\text{last}}(u^{(\infty)}) = 0$, since $\lim_{n \rightarrow \infty} \pi_{\text{last}}(z^{(n)}) = 0$. Furthermore, for all $J \in \mathcal{J}_2$, $\frac{\delta}{4} \leq |u_{J,1}^{(\infty)}|$ so $u_{J,1}^{(\infty)} \neq 0$. This shows the result. \square

We now extend this result to the case where the blocks do not necessarily have size 2. We first project the invariant \mathcal{I} on the *last two coordinates of each block* to obtain \mathcal{I}' . It is easy to see that \mathcal{I}' is invariant under the suitable restriction of A to those coordinates. Hence, by Lemma 6.12, we can find a point that is zero on the last coordinate but nonzero on the second-to-last coordinate of each block. We can then pull back this point through the projection and obtain a point $x \in \mathcal{I}$ with the same property. We next argue that the existence of x implies that \mathcal{I} must contain $\mathcal{Q} := \pi_{\text{last}}^{-1}(\{0\}) = \prod_{J \in \mathcal{J}} \mathbb{C}^{d(J)-1} \times \{0\}$. Indeed, if we project $\mathcal{I} \cap \mathcal{Q}$ on all coordinates *except the last one of each block*, then we obtain an invariant set again and so by applying (HR_{d-s}) , we conclude that it is a union of basic invariants. Now, since $x \in \mathcal{I} \cap \mathcal{Q}$, it follows that $\mathcal{I} \cap \mathcal{Q} = \mathbb{C}^{d-s}$. We now formalize this proof.

LEMMA 6.13. *Let $A \in \mathbb{C}^{d \times d}$ be a core matrix and \mathcal{I} be a closed semilinear set that is invariant under A and such that $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^s$. If $(HR_{d'})$ holds for all $d' < d$, then $\mathcal{Q} \subseteq \mathcal{I}$ where $\mathcal{Q} := \pi_{\text{last}}^{-1}(\{0\}) = \prod_{J \in \mathcal{J}} \mathbb{C}^{d(J)-1} \times \{0\}$.*

PROOF. In this proof, we will need to refer to coordinates with respect to both the original matrix A and some sub-matrices A_S with S a subset of the coordinates. To avoid confusing notations, we view the coordinates of A_S as a subset of that of A , so $d(J)$ still refers to the size of the Jordan block J in A . We will also write projection π_X with different domains, i.e., $\pi_X : \mathbb{C}^d \rightarrow \mathbb{C}^X$ and $\pi_X : \mathbb{C}^S \rightarrow \mathbb{C}^X$, it should be clear from the context what the domain of each projection is.

First note that if $\text{last}^c = \emptyset$ (which corresponds to the diagonal case), then $\mathcal{I} = \pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^s$ so the result is trivially true. Hence, we now assume that $\text{last}^c \neq \emptyset$. In particular, A has at least one Jordan block of size at least 2.

Note that $\mathcal{Q} \subseteq \mathcal{I}$ is equivalent to $\pi_{\text{last}^c}(\mathcal{I}) = \mathbb{C}^{d-s}$. Let $p = \pi_{\text{last}^c}$, let $\mathcal{I}' = p(\mathcal{I} \cap \mathcal{Q})$, and let $A' = A_{\text{last}^c}$. The last coordinate on each block in $\mathcal{I} \cap \mathcal{Q}$ is zero, therefore, $\mathcal{I} \cap \mathcal{Q} = p^{-1}(p(\mathcal{I} \cap \mathcal{Q})) \cap \mathcal{Q}$ and $p(A(\mathcal{I} \cap \mathcal{Q})) = A'p(\mathcal{I} \cap \mathcal{Q})$. It follows that \mathcal{I}' is invariant under A' and of dimension $d-s < d$. By (HR_{d-s}) , it holds that \mathcal{I}' is a union of basic invariants. If $\mathcal{I}' = \mathbb{C}^{d-s}$, then the lemma holds. Therefore, we assume, toward a contradiction that $\mathcal{I}' \neq \mathbb{C}^{d-s}$.

Basic invariants corresponding to A' are those of the form

$$\prod_{J \in \mathcal{J}_{\geq 2}} \mathbb{C}^{p_J} \times \{0\}^{d(J)-1-p_J},$$

where $\mathcal{J}_{\geq 2} = \{J \in \mathcal{J} \mid d(J) \geq 2\}$, and note that each such set that is not \mathbb{C}^{d-s} is identically zero on at least one coordinate $(J, d(J) - 1)$ for some $J \in \mathcal{J}_{\geq 2}$. It follows that

$$\mathcal{I}' \subseteq \bigcup_{J \in \mathcal{J}_{\geq 2}} \pi_{(J, d(J)-1)}^{-1}(\{0\});$$

in words, for each $z' \in I'$, there is a block $J \in \mathcal{J}_{\geq 2}$ such that the last coordinate of z' on J is zero. But, since $I \cap Q = p^{-1}(I') \cap Q$, we have

$$I \cap Q \subseteq \bigcup_{J \in \mathcal{J}_{\geq 2}} \left(\pi_{(J, d(J)-1)}^{-1}(\{0\}) \cap Q \right) = \bigcup_{J \in \mathcal{J}_{\geq 2}} \pi_{\text{last} \cup \{(J, d(J)-1)\}}^{-1}(\{0\}); \quad (7)$$

in words, all vectors of $I \cap Q$ have a second-to-last coordinate that is zero on some block. We will now project on the last two coordinates of each block. Formally, let

$$\text{last-two} = \{(J, 1) : J \in \mathcal{J}, d(J) = 1\} \cup \bigcup_{J \in \mathcal{J}_{\geq 2}} \{(J, d(J) - 1), (J, d(J))\}$$

and consider $I'' = \pi_{\text{last-two}}(I)$. We claim that I'' is invariant under $A'' := A_{\text{last-two}}$, since the last two coordinates of each block do not depend on the other coordinates when applying A . Furthermore, since $\text{last} \subseteq \text{last-two}$, we have that $\pi_{\text{last}}(I'') = \pi_{\text{last}}(\pi_{\text{last-two}}(I)) = \pi_{\text{last}}(I) = \mathbb{C}^s$. Therefore, we may apply Lemma 6.12 to I'' and A'' and get that there exists $z \in I''$ such that $\pi_{\text{last}}(z) = 0$ and for all $J \in \mathcal{J}$, if $d(J) \geq 2$, then $z_{J, d(J)-1} \neq 0$. But this contradicts Equation (7), because $z' = \pi_{\text{last-two}}^{-1}(z)$ is now such that $z' \in Q$ but $z'_{J, d(J)-1} \neq 0$ for all $J \in \mathcal{J}_{\geq 2}$. \square

At this stage, we may thus assume that the invariant I satisfies $\pi_{\text{last}}(I) = \mathbb{C}^s$ and contains $Q := \pi_{\text{last}}^{-1}(\{0\})$. We now aim to show that this implies that $I = \mathbb{C}^d$, and in particular I is a basic set.

LEMMA 6.14. *Let A be a core matrix of dimension d and I be a closed semilinear set that is invariant under A and such that $\pi_{\text{last}}(I) = \mathbb{C}^s$. If $(HR_{d'})$ holds for all $d' < d$, then $I = \mathbb{C}^d$.*

The rest of the section establishes Lemma 6.14; together with Lemma 6.11, this concludes our inductive proof of Theorem 6.1.

Let A be a core matrix of dimension d and I be a semilinear set invariant under A such that $\pi_{\text{last}}(I) = \mathbb{C}^s$, and assume that $(HR_{d'})$ holds for all $d' < d$. Since I is a closed semilinear set, we have

$$I = \bigcup_{\mathcal{P} \in P} \mathcal{P}, \quad \mathcal{P} = \bigcap_{\mathcal{H} \in H_{\mathcal{P}}} \mathcal{H},$$

where P is a finite set of polyhedra \mathcal{P} , and each \mathcal{P} is the intersection of a set $H_{\mathcal{P}}$ of finitely many closed half-spaces. We let $H = \bigcup_{\mathcal{P} \in P} H_{\mathcal{P}}$ denote the set of all half-spaces that appear in the definition of I .

We let $P_f = \{\mathcal{P} \in P \mid \dim_{\mathbb{R}} \mathcal{P} = 2d\}$ denote the set of fully dimensional polyhedra appearing in the definition of I . By Lemma 6.8, I has full dimension $2d$, therefore, P_f is non-empty. We will now show that we may, without loss of generality, ignore polyhedra that do not have full dimension.

LEMMA 6.15. *The semilinear set $I' = \bigcup_{\mathcal{P} \in P_f} \mathcal{P}$ is invariant under A .*

PROOF. Since all polyhedra in I' are closed and of full dimension, we have $\overline{I'^0} = I'$. Thus, it suffices to prove that $A(I'^0) \subseteq I'$ and conclude by continuity.

Let $z \in I'^0$ and let $\varepsilon > 0$ such that $B(z, \varepsilon) \subseteq I'$. Then, $AB(z, \varepsilon) \subseteq AI \subseteq I$. Note that for some small enough ε' , we have $B(Az, \varepsilon') \subseteq AB(z, \varepsilon)$, thus $B(Az, \varepsilon') \subseteq I$. It follows that $Az \in I'$, since $I \setminus I'$ has empty interior. \square

Now, by Lemma 6.11, either $\pi_{\text{last}}(I') = \mathbb{C}^s$ or I' is a union of basic invariants $\subseteq \mathbb{C}^d$. But the latter case is excluded, since such basic invariants do not have full dimension. Thus, $\pi_{\text{last}}(I') = \mathbb{C}^s$. Therefore, in the remainder of the proof, we now assume without loss of generality that $P_f = P$: All polyhedra have full dimension.

By Lemma 6.13, \mathcal{I} contains $\mathcal{Q} := \pi_{\text{last}}^{-1}(\{0\}) = \prod_{J \in \mathcal{J}} \mathbb{C}^{d(J)-1} \times \{0\}$. Let P_{full} denote the set of polyhedra \mathcal{P} in P such that $\mathcal{P} \cap \mathcal{Q}$ has dimension $2(d-s)$. Since $\mathcal{Q} \subseteq \mathcal{I}$ and \mathcal{Q} has dimension $2(d-s)$, we have $P_{\text{full}} \neq \emptyset$. We now exclude the case where all polyhedra in P_{full} are included in \mathcal{Q} .

LEMMA 6.16. *There is $\mathcal{P}_0 \in P_{\text{full}}$ such that $\mathcal{P}_0 \not\subseteq \mathcal{Q}$.*

PROOF. Assume towards a contradiction that for all $\mathcal{P} \in P_{\text{full}}$ it holds that $\mathcal{P} \subseteq \mathcal{Q}$. Let \mathcal{I}' be the closed semilinear set defined by

$$\mathcal{I}' = \bigcup_{\mathcal{P} \in P \setminus P_{\text{full}}} \mathcal{P}.$$

First, note that \mathcal{I}' is non-empty, otherwise, we would have $\mathcal{I} \subseteq \mathcal{Q}$, which implies $\pi_{\text{last}}(\mathcal{I}) = \{0\}$, however, $\pi_{\text{last}}(\mathcal{I}) = \mathbb{C}^s$.

We now prove that \mathcal{I}' is stable under A . Again, since all polyhedra have full dimension, it suffices to prove that $A(\mathcal{I}'^o) \subseteq \mathcal{I}'$. Let $z \in \mathcal{I}'^o$. There exists $\mathcal{P} \notin P_{\text{full}}$ and $\varepsilon > 0$ such that $B(z, \varepsilon) \subseteq \mathcal{P} \subseteq \mathcal{I}'$. But then it cannot be that $z \in \mathcal{Q}$, otherwise, it would hold that $\mathcal{P} \in P_{\text{full}}$. Since $\mathcal{Q}^c \subseteq \mathcal{Q}^c$, we get $Az \in \mathcal{I} \cap \mathcal{Q}^c$, which is contained in \mathcal{I}' by our assumption. Hence, \mathcal{I}' is stable under A .

Now, since \mathcal{I}' is stable under A and full dimension, it contains \mathcal{Q} by Lemma 6.11. However, \mathcal{I}' is a finite union of polyhedra whose intersections with \mathcal{Q} have dimension $< 2(d-s) = \dim_{\mathbb{R}} \mathcal{Q}$; thus \mathcal{I} cannot contain \mathcal{Q} . \square

Therefore, there exists $\mathcal{P}_0 \in P_{\text{full}}$, which is not contained in \mathcal{Q} . Let H_{general} be the family of half-spaces in H that do not contain \mathcal{Q} in their boundary. Now, if $\mathcal{H} \in H_{\text{general}}$, then $\partial\mathcal{H} \cap \mathcal{Q}$ has dimension $< 2(d-s)$. It follows that the countable union

$$X := \bigcup_{\mathcal{H} \in H_{\text{general}}} \bigcup_{k \in \mathbb{N}} A^{-k}(\partial\mathcal{H} \cap \mathcal{Q})$$

has dimension $< 2(d-s)$, so it may not cover $\mathcal{P}_0 \cap \mathcal{Q}$. Let $z \in (\mathcal{P}_0 \cap \mathcal{Q}) \setminus X$.

LEMMA 6.17. *For all $\mathcal{H} \notin H_{\text{general}}$, it holds that $\pi_{\text{last}}(\mathcal{H})$ is a closed half-space of \mathbb{C}^s satisfying $\mathcal{H} = \pi_{\text{last}}^{-1}(\pi_{\text{last}}(\mathcal{H}))$ and $0 \in \partial\pi_{\text{last}}(\mathcal{H})$.*

PROOF. By definition, if $\mathcal{H} \notin H_{\text{general}}$, then $\mathcal{Q} \subseteq \partial\mathcal{H} \subseteq \mathcal{H}$. Since \mathcal{Q} is a linear subspace of \mathbb{C}^d , this implies $\mathcal{H} + \mathcal{Q} = \mathcal{H}$. Now, $\mathcal{Q} = \ker(\pi_{\text{last}})$, therefore, $\pi_{\text{last}}^{-1}(\pi_{\text{last}}(\mathcal{H})) = \mathcal{H} + \ker(\pi_{\text{last}}) = \mathcal{H}$. Finally, we have $0 \in F \subseteq \partial\mathcal{H}$ and thus $0 \in \pi_{\text{last}}(\mathcal{H})$. \square

We proceed with another technical lemma, which prepares for our final construction.

LEMMA 6.18. *For all $k \in \mathbb{N}$, there exists $\varepsilon_k > 0$ such that for all $\mathcal{H} \in H_{\text{general}}$, the set $B(A^k z, \varepsilon_k) \cap \mathcal{H}$ is either empty or the whole ball $B(A^k z, \varepsilon_k)$.*

PROOF. We distinguish three cases.

- If $A^k z \in \mathcal{H}^o$, then there exists $\varepsilon_{k, \mathcal{H}} > 0$ such that $B(A^k z, \varepsilon) \cap \mathcal{H} = B(A^k z, \varepsilon)$ for all $\varepsilon \leq \varepsilon_{k, \mathcal{H}}$.
- If $A^k z \in \partial\mathcal{H}$, then $z = A^{-k} A^k z \in A^{-k} \partial\mathcal{H}$ and $z \in \mathcal{Q} = A^{-k} \mathcal{Q}$. But, since $\mathcal{H} \in H_{\text{general}}$, this implies $z \in X$, which is not possible.
- If $A^k z \notin \mathcal{H}$, then there exists $\varepsilon_{k, \mathcal{H}} > 0$ such that $B(A^k z, \varepsilon) \cap \mathcal{H} = \emptyset$ for all $\varepsilon \leq \varepsilon_{k, \mathcal{H}}$, since \mathcal{H} is closed.

Since the set H_{general} is finite, we conclude by taking ε_k to be the smallest $\varepsilon_{k, \mathcal{H}}$ for all $\mathcal{H} \in H_{\text{general}}$. \square

We then pick a sequence $(\varepsilon_k)_{k \in \mathbb{N}}$ as above and take it to be non-increasing and converging to 0 without loss of generality. For each $\mathcal{P} \in P$, we then have

$$B(A^k z, \varepsilon_k) \cap \bigcap_{\mathcal{H} \in H_{\mathcal{P}} \cap H_{\text{general}}} \mathcal{H} = B(A^k z, \varepsilon_k) \cap E_{k, \mathcal{P}}, \quad (8)$$

where $E_{k, \mathcal{P}}$ is either empty or \mathbb{C}^d . Now that $z \in \mathcal{I}$, this implies that for all k , $A^k z \in \mathcal{I}$, so there exists $\mathcal{P} \in P$ such that $A^k z \in \mathcal{P}$, and then $A^k z \in \mathcal{H}$ for all $\mathcal{H} \in H_{\mathcal{P}}$, by definition. Hence, for all k , there exists $\mathcal{P} \in P$ such that $E_{k, \mathcal{P}} = \mathbb{C}^d$.

It follows that for all k and \mathcal{P} ,

$$\begin{aligned} B(A^k z, \varepsilon_k) \cap \mathcal{P} &= B(A^k z, \varepsilon_k) \cap E_{k, \mathcal{P}} \cap \bigcap_{\mathcal{H} \in H_{\mathcal{P}} \setminus H_{\text{general}}} \mathcal{H} \\ &= B(A^k z, \varepsilon_k) \cap E_{k, \mathcal{P}} \cap \bigcap_{\mathcal{H} \in H_{\mathcal{P}} \setminus H_{\text{general}}} \pi_{\text{last}}^{-1}(\pi_{\text{last}}(\mathcal{H})) \\ &= B(A^k z, \varepsilon_k) \cap E_{k, \mathcal{P}} \cap \pi_{\text{last}}^{-1} \left(\bigcap_{\mathcal{H} \in H_{\mathcal{P}} \setminus H_{\text{general}}} \pi_{\text{last}}(\mathcal{H}) \right) \\ &= B(A^k z, \varepsilon_k) \cap \pi_{\text{last}}^{-1}(C_{k, \mathcal{P}}), \end{aligned}$$

where $C_{k, \mathcal{P}} = \bigcap_{\mathcal{H}' \in H_{k, \mathcal{P}}} \mathcal{H}'$ with

$$H_{k, \mathcal{P}} = \begin{cases} \{\pi_{\text{last}}(\mathcal{H}) : \mathcal{H} \in H_{\mathcal{P}} \setminus H_{\text{general}}\} & \text{if } E_{k, \mathcal{P}} = \mathbb{C}^d \\ \emptyset & \text{if } E_{k, \mathcal{P}} = \emptyset \end{cases}. \quad (9)$$

Note that by Lemma 6.17, $H_{k, \mathcal{P}}$ is a finite set of closed half-spaces \mathcal{H}' of \mathbb{C}^s such that $0 \in \partial \mathcal{H}'$. We further let $C_k = \bigcup_{\mathcal{P} \in P} C_{k, \mathcal{P}}$. Since $\mathcal{I} = \bigcup_{\mathcal{P} \in P} \mathcal{P}$, it follows that for all $k \in \mathbb{N}$,

$$B(A^k z, \varepsilon_k) \cap \mathcal{I} = B(A^k z, \varepsilon_k) \cap \pi_{\text{last}}^{-1}(C_k). \quad (10)$$

We now establish further properties of the C_k 's.

LEMMA 6.19. *For all $k \in \mathbb{N}$, it holds that C_k has full dimension $2s$, that $C_k \neq \mathbb{C}^s$, and that C_k is a union of convex cones. Moreover, there are finitely many different sets C_k when k ranges over \mathbb{N} .*

PROOF. We prove the properties one-by-one.

- C_k has full dimension $2s$. For this, we argue that $B(A^k z, \varepsilon_k) \cap \mathcal{I}$ has full dimension and therefore by Equation (10), $\pi_{\text{last}}^{-1}(C_k)$ also has full dimension that concludes. To show that $B(A^k z, \varepsilon_k) \cap \mathcal{I}$ has full dimension, it suffices to observe that $A^k z \in \mathcal{I}$ so $A^k z \in \mathcal{P}$ for some $\mathcal{P} \in P$, all of which are fully dimensional.
- $C_k \neq \mathbb{C}^s$. For this, let us consider $\overline{\mathcal{I}^c}$, a closed semilinear set that is invariant under A^{-1} . Using Lemma 6.6, A^{-1} rewrites as $\text{Diag}(\mathcal{J}_{d(J)}(\lambda_J^{-1}), J \in \mathcal{J})$ under an appropriate change of basis that preserves basic invariants. Hence, Lemma 6.11 applies to $\overline{\mathcal{I}^c}$, which yields that either $\overline{\mathcal{I}^c}$ is a basic invariant or $\pi_{\text{last}}(\overline{\mathcal{I}^c}) = \mathbb{C}^s$. But, since \mathcal{I} is closed and $\neq \mathbb{C}^d$, it holds that $\overline{\mathcal{I}^c}$ is fully dimensional, hence $\pi_{\text{last}}(\overline{\mathcal{I}^c}) = \mathbb{C}^s$. It follows from Lemma 6.13 that $Q \subseteq \overline{\mathcal{I}^c}$.

Now, recall that $A^k z \in Q$ so there exists $x \in \overline{\mathcal{I}^c} \cap B(A^k z, \varepsilon_k)$. In particular, it follows that $B(A^k z, \varepsilon_k) \cap \mathcal{I} \neq B(A^k z, \varepsilon_k)$. By Equation (10), this implies that $B(A^k z, \varepsilon_k) \cap \pi_{\text{last}}^{-1}(C_k) \neq B(A^k z, \varepsilon_k)$ and therefore $C_k \neq \mathbb{C}^s$.

- C_k is a union of convex cones: It suffices to show that each $C_{k,\mathcal{P}}$ is a convex cone. To show this, recall that $C_{k,\mathcal{P}} = \bigcap_{\mathcal{H}' \in H_{k,\mathcal{P}}} \mathcal{H}'$, which are such that $0 \in \partial \mathcal{H}'$ and contain a line, if the intersection consists of just one half-space.
- There are finitely many different sets C_k for k in \mathbb{N} . Indeed, by Equation (9), $H_{k,\mathcal{P}}$, and thus C_k , is determined only by whether $E_{k,\mathcal{P}} = \mathbb{R}^{2d}$ or \emptyset , so there are only $2^{|P|}$ possible values. Note that, however, ε_k does depend on k and may take arbitrarily small values if $A^k z$ gets arbitrarily close to some \mathcal{H} in H when k ranges over \mathbb{N} . \square

We are now finally in a position to present the final step of our proof.

As we have seen in the proof of Lemma 6.9, $\{(\lambda_1^k, \dots, \lambda_s^k), k \in \mathbb{N}\}$ is dense in $\{(\lambda_1^t, \dots, \lambda_s^t), t \in \mathbb{R}\}$. Hence, there exists an increasing sequence $\varphi : \mathbb{N} \rightarrow \mathbb{N}$ and $\varepsilon_k/2 \leq \mu_k \leq \varepsilon_k$ such that for all k , $(\lambda_1^{\varphi(k)}, \dots, \lambda_s^{\varphi(k)}) = (\lambda_1^{\mu_k}, \dots, \lambda_s^{\mu_k})$. Let C be such that $C = C_{\varphi(k)}$ for infinitely many k . Since C has full dimension by Lemma 6.19 it contains a point that is nonzero on every coordinate. Therefore, C cannot be stable under $M = \text{Diag}(\lambda_1, \dots, \lambda_s)$, otherwise Lemma 6.9 would conclude that $C = \mathbb{C}^s$, contradicting Lemma 6.19. In particular, there is $\tilde{u} \in C$ such that $M\tilde{u} \notin C$.

Let $t_0 = \sup\{t \in [0, 1] \mid M^t \tilde{u} \in C\}$ where $M^t := \text{Diag}(\lambda_1^t, \dots, \lambda_s^t)$. Since $t \mapsto M^t \tilde{u}$ is continuous and C is closed, it holds that $M^{t_0} \tilde{u} \in C$, hence $t_0 < 1$. Now, let $u = M^{t_0} \tilde{u}$. Then, for all sufficiently small $\varepsilon > 0$, we have $M^\varepsilon u \notin C$ by definition of the supremum.

We let $N \in \mathbb{N}$ be such that for $n \geq N$, ε_n is small enough in this sense, and N' be such that $\varphi(N') - \varphi(0) \geq N$. Recall that C is a cone, so we may re-scale u so $\|u\| \leq 2^{-\varphi(N')\varepsilon_{\varphi(N')}}$ and everything proved above about u remains true. Let $v = A^{\varphi(0)}z + u$. Recall that $z \in \mathcal{Q}$, that $\pi_{\text{last}}(\mathcal{Q}) = \{0\}$, and \mathcal{Q} is stable under A so $A^{\varphi(0)}z \in \mathcal{Q}$, hence $\pi_{\text{last}}(v) = u$. Therefore,

$$v \in \left[B \left(A^{\varphi(0)}z, 2^{-\varphi(N')\varepsilon_{\varphi(N')}} \right) \cap \pi_{\text{last}}^{-1}(\{u\}) \right] \subseteq \left[B \left(A^{\varphi(0)}z, \varepsilon_{\varphi(0)} \right) \cap \pi_{\text{last}}^{-1}(C) \right] \subseteq \mathcal{I},$$

where the last inclusion holds by Equation (10). We argue that $A^{\varphi(N')-\varphi(0)}v \in B(A^{\varphi(N')z}, \varepsilon_{\varphi(N')})$. Indeed, A is 2-lipschitzian, so $A^{\varphi(N')-\varphi(0)}$ is $2^{\varphi(N')}$ -lipschitzian (for the infinity norm $\|\cdot\|$), so

$$\|A^{\varphi(N')-\varphi(0)}v - A^{\varphi(N')z}\| \leq 2^{\varphi(N')} \|v - A^{\varphi(0)}z\| \leq \varepsilon_{\varphi(N')}.$$

Hence, since \mathcal{I} is stable under A , and by Equation (10),

$$A^{\varphi(N')-\varphi(0)}v \in B \left(A^{\varphi(N')z}, \varepsilon_{\varphi(N')} \right) \cap \mathcal{I} = B \left(A^{\varphi(N')z}, \varepsilon_{\varphi(N')} \right) \cap \pi_{\text{last}}^{-1}(C). \quad (11)$$

However,

$$\pi_{\text{last}} \left(A^{\varphi(N')-\varphi(0)}v \right) = \text{Diag} \left(\lambda_1^{\mu_{N'}}, \dots, \lambda_s^{\mu_{N'} } \right) u = \text{Diag} \left(\lambda_1^{\mu_{N'}}, \dots, \lambda_s^{\mu_{N'} } \right) u = M^{\mu_{N'} } u.$$

Since $\mu_{N'} \leq \varepsilon_{N'} \leq \varepsilon_N$ and we chose N such that $M^\alpha u \notin C$ for any $0 < \alpha \leq \varepsilon_N$, this shows that $\pi_{\text{last}}(A^{\varphi(N')-\varphi(0)}v) \notin C$, contradicting Equation (11) and concluding the proof.

7 Conclusions

In this article, we have proved that the Monniaux problem is undecidable already in a very restricted setting: using semilinear invariants for affine programs (without guards), and in fact using only a single control location and two transitions. This very foundational undecidability result shows that there is little hope for decidability for the Monniaux problem, as most natural classes will include them. What we leave as an open question is whether convex invariants can help recover decidability. This is a very exciting perspective, since, as pointed out in the introduction, convex invariants appear naturally in many practical scenarios.

Our decidability result considers the case of a single transition. On a technical level, the proof helps us understand what exactly semilinear invariants can be used for in the context of affine

programs. This surprising positive result opens several perspectives. First, going beyond semilinear invariants: It is already known that the Monniaux problem is decidable for semialgebraic invariants [13, 14], but it remains open for other natural classes of invariants. Second, this decidability result implies a complexity result, but not yet an efficient algorithm. We leave open whether the problem can be efficiently solved and what consequences are there for static analysis of programs.

Acknowledgments

We thank the reviewers for their extremely detailed comments, which greatly helped improve and clarify the article.

Appendices

A Upper Bound on Aperiodicity Index

In this section, we give an polynomial (when the dimension d is fixed) upper bound on the smallest number n such that A^n is aperiodic. We start with the following lemma:

LEMMA A.1. *If λ is an m th root of unity and an algebraic number of degree d , then $m \leq 2d^2$.*

PROOF. Let $m > 1$ be the smallest integer such that $\lambda^m = 1$ and let Φ_m be the m th cyclotomic polynomial. Then, Φ_m is the minimal polynomial of λ by definition of m . Furthermore, recall that the degree of Φ_m is $\phi(m)$ where ϕ is Euler's totient function. But, since Φ_m is the minimal polynomial of λ , its degree is exactly d , hence $d = \phi(m)$. It is well-known that $\phi(m) \geq \sqrt{m}/2$, therefore, $m \leq 2d^2$. \square

LEMMA A.2. *Let A be a rational matrix whose eigenvalues $\lambda_1, \dots, \lambda_s$ have modulus 1. Any number of the form λ_i , λ_i/λ_j , or $\lambda_i\lambda_j$, if it is a root of unity, has order bounded by a polynomial in the dimension of A .*

It follows that A^n is aperiodic for some n that is at most polynomial.

PROOF. We first consider the case where λ is an eigenvalue of A that is a root of unity. Since A has rational coefficients, its minimal polynomial p_A also has rational coefficients. Furthermore, $p_A(\lambda) = 0$, since λ is an eigenvalue. Therefore, λ is an algebraic number of degree at most n , the dimension of A . It follows by Lemma A.1 that the order of λ is polynomial in n . Now, assume that λ and μ are two eigenvalues of A such that $\lambda\mu$ is a root of unity. Since λ and μ are algebraic numbers of degree at most n , $\lambda\mu$ and λ/μ are algebraic of degree at most n^2 . Therefore, by the same argument, the order of $\lambda\mu$ and λ/μ is polynomial in n^2 , hence in n . \square

B Computing Some Determinants

In this Appendix, we prove Lemma 6.5, which we first restate for convenience.

LEMMA 6.5. *Let $A \in \mathbb{C}^{d \times d}$ be in Jordan normal form, \mathcal{J} its Jordan blocks, and let λ_J denote the eigenvalue of each Jordan block $J \in \mathcal{J}$. Let $x \in \mathbb{C}^d$ and $M = [x_{\mathbb{R}} \quad (Ax)_{\mathbb{R}} \quad \dots \quad (A^{2d-1}x)_{\mathbb{R}}]$. If all the eigenvalues of $A_{\mathbb{R}}$ are distinct and $x_{J, d(J)} \neq 0$ for all $J \in \mathcal{J}$, then $\det(M) \neq 0$.*

Towards proving Lemma 6.5, we first compute a similar determinant in a complex setting.

LEMMA B.1. *Let $A \in \mathbb{C}^{d \times d}$ be in Jordan normal form, \mathcal{J} its Jordan blocks, and let λ_J denote the eigenvalue of each Jordan block $J \in \mathcal{J}$. Let $x \in \mathbb{C}^d$ and $M = [x \quad A^1x \quad \dots \quad A^{d-1}x]$. Then,*

$$\det(M) = \prod_{J \in \mathcal{J}} (-x_{J, d(J)})^{d(J)} \cdot \prod_{J, H \in \mathcal{J}, J \neq H} (\lambda_J - \lambda_H)^{d(J)d(H)}.$$

PROOF OF LEMMA B.1. We let $M(A, x, n) = [A^0x \ A^1x \ \cdots \ A^{n-1}x]$ for any $A \in \mathbb{C}^{d \times d}$, which is a rectangular matrix, and $M(A, x) = M(A, x, d)$ which is a square matrix. First check, by an easy induction, that for any integer n , Jordan block $J \in \mathcal{J}$ and $i \in [1, d(J)]$,

$$(A^n x)_{J,i} = \sum_{j=i}^{d(J)} \binom{n}{j-i} \lambda_J^{n+i-j} x_{J,j}.$$

We now proceed by induction on the number of blocks. If $d = 0$, then $\det(M(A)) = 1$ so the formula is true. Let $A \in \mathbb{C}^{d \times d}$ with $s > 0$ blocks. Fix a Jordan block $J_0 \in \mathcal{J}$. To avoid any confusion, let $\mu = \lambda_{J_0}$. By performing linear combination of the columns, we can transform $M(A, x)$ into

$$B = [C^0 \ \cdots \ C^{d-1}] \quad \text{where} \quad C^n = A^n x + \sum_{k=0}^{n-1} (-\mu)^{n-k} \binom{n}{k} A^k x.$$

Note that these linear transformations are all of the form “add a multiple of a column to another one,” hence it does not affect the determinant. Let $J \in \mathcal{J}$ and $i \in [1, d(J)]$, then

$$\begin{aligned} C_{J,i}^n &= (A^n x)_{J,i} + \sum_{k=0}^{n-1} (-\mu)^{n-k} \binom{n}{k} (A^k x)_{J,i} \\ &= \sum_{j=i}^{d(J)} \binom{n}{j-i} \lambda_J^{n+i-j} x_{J,j} + \sum_{k=0}^{n-1} (-\mu)^{n-k} \binom{n}{k} \sum_{j=i}^{d(J)} \binom{k}{j-i} \lambda_J^{k+i-j} x_{J,j} \\ &= \sum_{j=i}^{d(J)} \left[\binom{n}{j-i} \lambda_J^{n+i-j} + \sum_{k=0}^{n-1} (-\mu)^{n-k} \lambda_J^{k+i-j} \binom{n}{k} \binom{k}{j-i} \right] x_{J,j} \\ &= \sum_{j=i}^{d(J)} \left[\sum_{k=0}^n (-\mu)^{n-k} \lambda_J^{k+i-j} \binom{n}{k} \binom{k}{j-i} \right] x_{J,j}. \end{aligned}$$

Now, observe that

$$\begin{aligned} \sum_{k=0}^n (-\mu)^{n-k} \lambda_J^{k+i-j} \binom{n}{k} \binom{k}{j-i} &= \sum_{k=j-i}^n (-\mu)^{n-k} \lambda_J^{k+i-j} \binom{n}{k} \binom{k}{j-i} && \text{since } \binom{k}{j-i} = 0 \text{ for } k \leq j-i \\ &= \sum_{k=0}^{n-j+i} (-\mu)^{n-k-j+i} \lambda_J^k \binom{n}{k+j-i} \binom{k+j-i}{j-i} && \text{by re-indexing} \\ &= \sum_{k=0}^{n-j+i} (-\mu)^{n-j+i-k} \lambda_J^k \binom{n}{j-i} \binom{n-j+i}{k} && \text{by the identity } \binom{n}{k+h} \binom{k+h}{h} = \binom{n}{h} \binom{n-h}{k} \\ &= \binom{n}{j-i} \sum_{k=0}^{n-j+i} (-\mu)^{n-j+i-k} \lambda_J^k \binom{n-j+i}{k} \\ &= \binom{n}{j-i} (\lambda - \mu)^{n-j+i} && \text{by the binomial theorem.} \end{aligned}$$

Therefore,

$$C_{J,i}^n = \sum_{j=i}^{d(J)} \binom{n}{j-i} (\lambda - \mu)^{n-j+i} x_{J,i}.$$

Note that this is exactly the expression for a Jordan block with eigenvalue $\lambda - \mu$. In other words, $\det(M(A, x)) = \det(M(\tilde{A}, x))$, where $\tilde{A} = A - \mu I_d$ has the same Jordan blocks as A but different

eigenvalues. In particular, the block J_0 has eigenvalue $\lambda_{J_0} - \mu = 0$ in \tilde{A} so it is a nilpotent block. Thus, $\tilde{A}_{J_0}^{d(J_0)} = 0$ so

$$M(\tilde{A}, x) = \begin{bmatrix} M(\tilde{A}_{J_0}, x_{J_0}, d(J_0)) & \tilde{A}_{J_0}^{d(J_0)} M(\tilde{A}_{J_0}, x_{J_0}, d - d(J_0)) \\ M(\tilde{A}_{J_0^c}, x_{J_0^c}, d(J_0)) & \tilde{A}_{J_0^c}^{d(J_0)} M(\tilde{A}_{J_0^c}, x_{J_0^c}, d - d(J_0)) \end{bmatrix} = \begin{bmatrix} M(\tilde{A}_{J_0}, x_{J_0}) & 0 \\ * & \tilde{A}_{J_0^c}^{d(J_0)} M(\tilde{A}_{J_0^c}, x_{J_0^c}) \end{bmatrix}.$$

In particular,

$$\det(M(A, x)) = \det(M(\tilde{A}, x)) = \det\left(M\left(\tilde{A}_{J_0}, x_{J_0}\right)\right) \det\left(\tilde{A}_{J_0^c}^{d(J_0)}\right) \det\left(M\left(\tilde{A}_{J_0^c}, x_{J_0^c}\right)\right).$$

It is not hard to see that

$$M(\tilde{A}_{J_0}, x_{J_0}) = \begin{bmatrix} x_{J_0,1} & \cdots & x_{J_0,d(J_0)} \\ \vdots & \ddots & \\ x_{J_0,d(J_0)} & & \end{bmatrix},$$

so its determinant is $(-x_{J_0,d(J_0)})^{d(J_0)}$. Furthermore, since \tilde{A} is in JNF, its determinant is the product of its eigenvalues (with multiplicities). Now, let $\tilde{\mathcal{J}} = \mathcal{J} \setminus \{J_0\}$ denote the Jordan blocks of \tilde{A} , $\tilde{\lambda}_J = \lambda_J - \lambda_{J_0}$ denote the eigenvalue of the block J in \tilde{A} . Then,

$$\begin{aligned} \det(M(A, x)) &= (-x_{J_0,d(J_0)})^{d(J_0)} \left(\prod_{J \in \tilde{\mathcal{J}}} \tilde{\lambda}_J^{d(J)} \right)^{d(J_0)} \det\left(M\left(\tilde{A}_{J_0^c}, x_{J_0^c}\right)\right) \\ &= (-x_{J_0,d(J_0)})^{d(J_0)} \cdot \prod_{J \in \tilde{\mathcal{J}}} (\lambda_J - \lambda_{J_0})^{d(J)d(J_0)} \cdot \det\left(M\left(\tilde{A}_{J_0^c}, x_{J_0^c}\right)\right). \end{aligned}$$

By the induction hypothesis applied to $\tilde{A}_{J_0^c}$ and $x_{J_0^c}$, we get that (note that $(x_{J_0^c})_{J,d(J)} = x_{J,d(J)}$ for $J \neq J_0$)

$$\begin{aligned} \det\left(M\left(\tilde{A}_{J_0^c}, x_{J_0^c}\right)\right) &= \prod_{J \in \tilde{\mathcal{J}}} (-x_{J,d(J)})^{d(J)} \cdot \prod_{J,H \in \tilde{\mathcal{J}}, J \neq H} (\tilde{\lambda}_J - \tilde{\lambda}_H)^{d(J)d(H)} \\ &= \prod_{J \in \tilde{\mathcal{J}}} (-x_{J,d(J)})^{d(J)} \cdot \prod_{J,H \in \tilde{\mathcal{J}}, J \neq H} (\lambda_J - \lambda_H)^{d(J)d(H)}. \end{aligned}$$

And we get the result by putting everything together. \square

We are now ready to show Lemma 6.5.

PROOF OF LEMMA 6.5. It is not hard to check that $A_{\mathbb{R}}$ is the block matrix $(R(A_{ij}))_{i,j}$ where

$$R(z) = \begin{bmatrix} \operatorname{Re}(z) & -\operatorname{Im}(z) \\ \operatorname{Im}(z) & \operatorname{Re}(z) \end{bmatrix}$$

for all $z \in \mathbb{C}$. Furthermore, there is a change of basis Q (independent of z) such that

$$Q^{-1}R(z)Q = \begin{bmatrix} z & 0 \\ 0 & z^* \end{bmatrix}.$$

Therefore, by applying Q block-wise and permuting rows and columns, we can build a change of basis P such that $A_{\mathbb{R}} = P^{-1}BP$ and $x_{\mathbb{R}} = P^{-1}y$ where

$$B = \begin{bmatrix} A & 0 \\ 0 & A^* \end{bmatrix}, \quad y = \begin{bmatrix} x \\ x^* \end{bmatrix}.$$

It follows that $PM = [y \quad By \quad \cdots \quad B^{2d-1}y]$. One can check that Q and P have determinant 1 so $\det(M) = \det(PM)$, and we have reduced the problem to computing the determinant of

$$N = [y \quad By \quad \cdots \quad B^{2d-1}y],$$

where B is in Jordan normal form. Specifically, the Jordan blocks of B are the Jordan blocks of A and their conjugates. By Lemma B.1,

$$\det(N) = \prod_{J \in \mathcal{J}_B} (-y_{J,d(J)})^{d(J)} \cdot \prod_{J,H \in \mathcal{J}_B, J \neq H} (\lambda_J^B - \lambda_H^B)^{d(J)d(H)}.$$

Since we assume all the eigenvalues of $A_{\mathbb{R}}$ to be distinct, the second product in this expression is nonzero. The first product is equal to

$$\prod_{J \in \mathcal{J}_B} (-y_{J,d(J)})^{d(J)} = \prod_{J \in \mathcal{J}_A} (-x_{J,d(J)} x_{J,d(J)}^*)^{d(J)} = \prod_{J \in \mathcal{J}_A} (-\operatorname{Re}(x_{J,d(J)})^2 - \operatorname{Im}(x_{J,d(J)})^2)^{d(J)}$$

and therefore is nonzero if $x_{J,d(J)} \neq 0$ for all J . \square

C Inverse of a Core Matrix

We now prove Lemma 6.6, which we first restate for convenience.

LEMMA C.1. *Let $A \in \mathbb{C}^{d \times d}$ be a core matrix and \mathcal{J} range over its Jordan blocks. There exists a change of basis P that stabilizes any basic invariant, and such that $PA^{-1}P^{-1} = \operatorname{Diag}(\mathcal{J}_{d(J)}(\lambda_J^{-1}), J \in \mathcal{J})$.*

PROOF. Recall that a core matrix A is a diagonal block matrix, where blocks are of the form

$$A_J = \begin{bmatrix} \lambda & 1 & & & \\ & \lambda & \ddots & & \\ & & \ddots & \ddots & \\ & & & \ddots & 1 \\ & & & & \lambda \end{bmatrix},$$

where J ranges over the set of Jordan blocks \mathcal{J} . Then, we have

$$A_J^{-1} = \begin{bmatrix} \lambda^{-1} & \lambda^{-2} & \lambda^{-3} & \cdots & \lambda^{-d(J)} \\ & \lambda^{-1} & \lambda^{-2} & & \lambda^{-d(J)+1} \\ & & & \ddots & \\ & & & & \lambda^{-2} \\ & & & & \lambda^{-1} \end{bmatrix}.$$

Then, we get $P_J A_J^{-1} P_J^{-1} = \mathcal{J}_{d(J)}(\lambda^{-1})$ with P_J upper triangular, and the lemma follows. \square

D Special JNF for Real Matrices

We now prove the following statement about Jordan normal form of real matrices:

however, that this implies no relationship between J'_i and J_j , they could have different sizes). Furthermore, a consequence of the equation $AP = MP$ is that P is of the form

$$P = \begin{bmatrix} S_1 & \cdots & S_r & T_1 & \cdots & T_s & T'_1 & \cdots & T'_s \end{bmatrix},$$

where the columns of each S_i , T_i , and T'_i form a Jordan chain, which means that for each i ,

$$T_i = \begin{bmatrix} x^{i,1} & \cdots & x^{i,d(J_i)} \end{bmatrix},$$

where $x^{i,1}, \dots, x^{i,d(J_i)}$ are generalized eigenvectors of A that satisfy

$$Ax^{i,1} = \mu_i x^{i,1}, \quad Ax^{i,j} = \mu_i x^{i,j} + x^{i,j-1},$$

for $j = 2, \dots, d(J_i)$. Finally, we introduce a few notations for vector spaces: Given a (not necessarily square) matrix X , we let \mathcal{V}_X be the span of the columns of X . A consequence of the JNF (i.e., P is invertible) is that

$$\mathbb{R}^d = \mathcal{V}_R \oplus \mathcal{V}_{I^+} \oplus \mathcal{V}_{I^-},$$

where

$$\mathcal{V}_R = \mathcal{V}_{R_1} \oplus \cdots \oplus \mathcal{V}_{R_r} \quad \mathcal{V}_{I^+} = \mathcal{V}_{J_1} \oplus \cdots \oplus \mathcal{V}_{J_s} \quad \mathcal{V}_{I^-} = \mathcal{V}_{J'_1} \oplus \cdots \oplus \mathcal{V}_{J'_s}$$

are the generalized eigenspaces of the real and complex eigenvalues. We now let

$$\tilde{P} = \begin{bmatrix} S_1 & \cdots & S_r & T_1 & \cdots & T_s & T_1^* & \cdots & T_s^* \end{bmatrix}, \quad \tilde{M} = \text{Diag}(R_1, \dots, R_r, J_1, \dots, J_s, J_1^*, \dots, J_s^*),$$

which are essentially P and M where we have replaced the T'_i and J'_i (which can be anything and possibly be unrelated to T_i and J_i) by T_i^* and J_i^* (the conjugates). Note that at this point, it is not clear that \tilde{P} and \tilde{M} are square matrices. We claim that

$$A\tilde{P} = \tilde{P}\tilde{M}.$$

For sub-matrices S_i and T_i , this follows directly from the equation $AP = PM$ but we need to verify that it holds for T_i^* . For that, we simply note that $A = A^*$ (A is real) and $AP = PM$ so

$$AT_i^* = (A^*T_i)^* = (AT_i)^* = (T_iJ_i)^* = T_i^*J_i^*,$$

which is what we wanted. Therefore, it only remains to see that \tilde{P} is square and invertible. The submatrix $X := [R_1 \ \cdots \ R_r \ T_1 \ \cdots \ T_s]$ has linearly independent columns, since it is a subset of the columns of P that is invertible. Furthermore, $\mathcal{V}_X = \mathcal{V}_R \oplus \mathcal{V}_{I^+}$. The columns of $Y := [T_1^* \ \cdots \ T_s^*]$ are also linearly independent, because $\mathcal{V}_{T_i^*} = \mathcal{V}_{T_i}^*$ and therefore

$$\mathcal{V}_Y = \mathcal{V}_{T_1^*} + \cdots + \mathcal{V}_{T_s^*} = \mathcal{V}_{T_1}^* + \cdots + \mathcal{V}_{T_s}^* = (\mathcal{V}_{T_1} + \cdots + \mathcal{V}_{T_s})^* = (\mathcal{V}_{T_1} \oplus \cdots \oplus \mathcal{V}_{T_s})^* = \mathcal{V}_{I^+}^*.$$

Therefore, we need to show that $\mathcal{V}_X \oplus \mathcal{V}_Y = \mathbb{R}^d$, that is to say $\mathbb{R}^d = \mathcal{V}_R \oplus \mathcal{V}_{I^+} \oplus \mathcal{V}_{I^+}^*$.

First, we make an observation: For any i , since we have $AT_i^* = T_i^*J_i^*$ and J_i^* is a Jordan block for μ_i^* , then T_i^* is a Jordan chain for μ_i^* and therefore $\mathcal{V}_{T_i^*}$ is included in the generalized eigenspace of A for μ_i^* .

We claim that this implies that $\mathcal{V}_R \oplus \mathcal{V}_{I^+}$ and $\mathcal{V}_{I^+}^*$ are in direct sum. To see that, we need to show that $\mathcal{V}_{J_i^*}$ and $\mathcal{V}_R \oplus \mathcal{V}_{I^+}$ are in direct sum for all i . On the one hand, $\mathcal{V}_{J_i^*}$ is included in the generalized eigenspace of A for the eigenvalues μ_i^* . On the other hand, $\mathcal{V}_R \oplus \mathcal{V}_{I^+}$ is the direct sum of all generalized eigenvalues δ that satisfy $\text{Im}(\delta) \geq 0$ and therefore are all distinct from μ_i^* because $\text{Im}(\mu_i^*) = -\text{Im}(\mu_i) < 0$ (a consequence of $\text{Im}(\lambda_j) = 0$ and $\text{Im}(\mu_j) > 0$ for all j), as just shown. By standard facts, generalized eigenspaces for distinct eigenvalues are in direct sum which shows the claim.

Finally, we claim that $\mathbb{R}^d = \mathcal{V}_{\mathcal{R}} \oplus \mathcal{V}_{\mathcal{I}^+} \oplus \mathcal{V}_{\mathcal{I}^*}$. To do that, we will show that $\mathcal{V}_{\mathcal{I}^-} \subseteq \mathcal{V}_{\mathcal{I}^*}$ which will conclude, since $\mathbb{R}^d = \mathcal{V}_{\mathcal{R}} \oplus \mathcal{V}_{\mathcal{I}^+} \oplus \mathcal{V}_{\mathcal{I}^-}$. Let $1 \leq i \leq s'$, we will show that $\mathcal{V}_{T'_i} \subseteq \mathcal{V}_{\mathcal{I}^*}$. Recall that T'_i is a Jordan chain and $AT'_i = T'_i J'_i$, hence

$$AT'_i{}^* = (A^* T'_i)^* = (AT'_i)^* = (T'_i J'_i)^* = T'_i{}^* J'_i{}^*,$$

since A is real. Since $J'_i{}^*$ is Jordan block, $T'_i{}^*$ is a Jordan chain for $\mu'_i{}^*$ and $\mathcal{V}_{T'_i{}^*}$ is included in the generalized eigenspace of $\mu'_i{}^*$. But $\text{Im}(\mu'_i{}^*) = -\text{Im}(\mu'_i) > 0$ so $\mu'_i{}^* = \mu_j$ for some j and then $\mathcal{V}_{T'_i{}^*}$ is included in the generalized eigenspace of μ_j and therefore in $\mathcal{V}_{\mathcal{I}^+}$. At the same time, $\mathcal{V}_{T'_i{}^*} = \mathcal{V}_{J'_i{}^*}$ so $\mathcal{V}_{T'_i} \subseteq \mathcal{V}_{\mathcal{I}^*}$.

In summary, we have shown that $A\tilde{P} = \tilde{P}\tilde{M}$, where \tilde{M} is in JNF with conjugated blocks and \tilde{P} has correspondingly “conjugated” columns.

It remains to see that \tilde{P}^{-1} has a similar structure. In what follows, we rename \tilde{P} to P and \tilde{M} to M so $AP = PM$, and P and M are “conjugated.” Assume for the moment that A is invertible. We claim that we can compute an invertible matrix Q , with the same conjugation pattern as P , such that $QA^{-1} = M^{-1}Q$. The details on how to compute such a Q can be found in Reference [3, Theorem 4.1 and Appendix 1] but, in short, this is exactly the same algorithm used to compute P but in “row form,” or equivalently, since $A^{-T}Q^T = Q^T M^{-T}$, can be seen as another version of the JNF applied to A^{-T} where the “ones” are below the diagonal. Intuitively, this works because the generalized eigenspaces of A^{-1} have exactly the same structure as that of A . Having found such a Q , we note that

$$MQP = M(QA^{-1})(AP) = M(M^{-1}Q)(PM) = QPM,$$

so M and QP commute. But M is block-diagonal, so it follows that QP must also be block-diagonal, i.e., $QP = X := \text{Diag}(X_1, \dots, X_{r+2s})$ with the same block structure as M . Furthermore, since both P and Q have the same conjugated structure, it follows that X is conjugated with the same structure. Finally, we observe that $P^{-1} = X^{-1}Q$, which preserves again the conjugated structure and shows the result. In the case where A is not invertible, we instead replace A with $A' = A + \delta I_n$ for some very large δ so A' is invertible. It is then not hard to see that the JNF of A' is $M' = M + \delta I_n$. We can then compute P and P^{-1} with the conjugated structure from A' as above and those will be acceptable for A as well.

References

- [1] Shaull Almagor, Dmitry Chistikov, Joël Ouaknine, and James Worrell. 2022. O-minimal invariants for discrete-time dynamical systems. *ACM Trans. Computat. Logic* 23, 2 (2022), 9:1–9:20. DOI : <https://doi.org/10.1145/3501299>
- [2] Alexey Bakhirkin and David Monniaux. 2018. Extending constraint-only representation of polyhedra with Boolean constraints. In *Proceedings of the International Symposium on Static Analysis (SAS'18) (Lecture Notes in Computer Science, Vol. 11002)*, Andreas Podelski (Ed.). Springer, 127–145. DOI : https://doi.org/10.1007/978-3-319-99725-4_10
- [3] Jin-yi Cai. 2000. *Computing Jordan Normal Forms Exactly for Commuting Matrices in Polynomial Time*. Technical Report. SUNY at Buffalo.
- [4] Jin-yi Cai, Richard J. Lipton, and Yechezkel Zalcstein. 2000. The complexity of the A B C problem. *SIAM J. Comput.* 29, 6 (2000), 1878–1888. DOI : <https://doi.org/10.1137/S0097539794276853>
- [5] John William Scott Cassels. 1965. *An Introduction to Diophantine Approximation*. Cambridge University Press. DOI : <https://doi.org/10.1017/S0008439500024693>
- [6] Robert Clarisó and Jordi Cortadella. 2007. The octahedron abstract domain. *Sci. Comput. Program.* 64, 1 (2007), 115–139. DOI : <https://doi.org/10.1016/J.SCICO.2006.03.009>
- [7] Henri Cohen. 1993. *A Course in Computational Algebraic Number Theory*. Springer-Verlag. DOI : <https://doi.org/10.1007/978-3-662-02945-9>
- [8] Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, and Xavier Rival. 2009. Why does Astrée scale up? *Form. Meth. Syst. Des.* 35, 3 (2009), 229–264. DOI : <https://doi.org/10.1007/S10703-009-0089-6>

- [9] Patrick Cousot and Nicolas Halbwachs. 1978. Automatic discovery of linear restraints among variables of a program. In *Proceedings of the Conference on Principles of Programming Languages (POPL '78)*. ACM Press. DOI: <https://doi.org/10.1145/512760.512770>
- [10] Jing Dong and Qinghui Liu. 2012. Undecidability of infinite post correspondence problem for instances of size 8. *RAIRO - Theoret. Inform. Applic.* 46, 3 (2012). DOI: <https://doi.org/10.1051/ITA/2012015>
- [11] Graham Everest, Alf van der Poorten, Igor Shparlinski, and Thomas Ward. 2003. *Recurrence Sequences (Mathematical Surveys and Monographs, Vol. 104)*. American Mathematical Society, United States.
- [12] Nathanaël Fijalkow, Engel Lefaucheux, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2019. On the Monniaux problem in abstract interpretation. In *Proceedings of the International Symposium on Static Analysis (SAS'19)*. DOI: https://doi.org/10.1007/978-3-030-32304-2_9
- [13] Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2017. Semialgebraic invariant synthesis for the Kannan-Lipton orbit problem. In *Proceedings of the Symposium on Theoretical Aspects of Computer Science (STACS'17)*. DOI: <https://doi.org/10.4230/LIPICS.STACS.2017.29>
- [14] Nathanaël Fijalkow, Pierre Ohlmann, Joël Ouaknine, Amaury Pouly, and James Worrell. 2019. Complete semialgebraic invariant synthesis for the Kannan-Lipton orbit problem. *Theor. Comput. Syst.* 63, 5 (2019). DOI: <https://doi.org/10.1007/S00224-019-09913-3>
- [15] Thomas Gawlitza and Helmut Seidl. 2007. Precise relational invariants through strategy iteration. In *Proceedings of the Conference on Computer Science Logic (CSL'07) (Lecture Notes in Computer Science, Vol. 4646)*, Jacques Duparc and Thomas A. Henzinger (Eds.). Springer, 23–40. DOI: https://doi.org/10.1007/978-3-540-74915-8_6
- [16] Khalil Ghorbal, Franjo Ivancic, Gogul Balakrishnan, Naoto Maeda, and Aarti Gupta. 2012. Donut domains: Efficient non-convex domains for abstract interpretation. In *Proceedings of the International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI'12)*. DOI: https://doi.org/10.1007/978-3-642-27940-9_16
- [17] Roberto Giacobazzi, Francesco Logozzo, and Francesco Ranzato. 2015. Analyzing program analyses. In *Proceedings of the Conference on Principles of Programming Languages (POPL '15)*. DOI: <https://doi.org/10.1145/2676726.2676987>
- [18] Roberto Giacobazzi, Francesco Ranzato, and Francesca Scozzari. 2000. Making abstract interpretations complete. *J. ACM* 47, 2 (2000), 361–416. DOI: <https://doi.org/10.1145/333979.333989>
- [19] Ehud Hrushovski, Joël Ouaknine, Amaury Pouly, and James Worrell. 2018. Polynomial invariants for affine programs. In *Proceedings of the Conference on Logic in Computer Science (LICS'18)*. DOI: <https://doi.org/10.1145/3209108.3209142>
- [20] Ravindran Kannan and Richard J. Lipton. 1980. The orbit problem is decidable. In *Proceedings of the Symposium on Theory of Computing (STOC'80)*. DOI: <https://doi.org/10.1145/800141.804673>
- [21] Ravindran Kannan and Richard J. Lipton. 1986. Polynomial-time algorithm for the orbit problem. *J. ACM* 33, 4 (1986), 808–821. DOI: <https://doi.org/10.1145/6490.6496>
- [22] Michael Karr. 1976. Affine relationships among variables of a program. *Acta Inform.* 6 (1976), 133–151. DOI: <https://doi.org/10.1007/BF00268497>
- [23] Zachary Kincaid, John Cyphert, Jason Breck, and Thomas W. Reps. 2018. Non-linear reasoning for invariant synthesis. *Proc. ACM Program. Lang.* 2 (2018). DOI: <https://doi.org/10.1145/3158142>
- [24] Christer Lech. 1953. A note on recurring series. *Arkiv Matem.* 2 (1953), 417–421.
- [25] Kurt Mahler. 1935. Eine arithmetische eigenschaft der Taylor Koeffizienten rationaler funktionen. *Proc. Konink. Nederl. Akad. Wetenschapp.* 38 (1935).
- [26] Kurt Mahler. 1956. On the Taylor coefficients of rational functions. In *Proceedings of the Cambridge Philosophical Society*. Cambridge University Press.
- [27] Antoine Miné. 2006. The octagon abstract domain. *High. Ord. Symbol. Comput.* 19, 1 (2006), 31–100. DOI: <https://doi.org/10.1007/s10990-006-8609-1>
- [28] David Monniaux. 2019. On the decidability of the existence of polyhedral invariants in transition systems. *Acta Inform.* 56, 4 (2019), 385–389. DOI: <https://doi.org/10.1007/S00236-018-0324-Y>
- [29] David Monniaux. 2023. Completeness in static analysis by abstract interpretation, a personal point of view. *Chall. Softw. Verif.* 238 (2023), 93–108. Retrieved from <https://hal.science/hal-03857312v2>
- [30] Markus Müller-Olm and Helmut Seidl. 2004. A note on Karr's algorithm. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP'04)*. DOI: https://doi.org/10.1007/978-3-540-27836-8_85
- [31] George C. Necula and Sumit Gulwani. 2005. Randomized algorithms for program analysis and verification. In *Proceedings of the Conference on Computer Aided Verification (CAV'05) (Lecture Notes in Computer Science, Vol. 3576)*, Kousha Etessami and Sriram K. Rajamani (Eds.). Springer, 1. DOI: https://doi.org/10.1007/11513988_1
- [32] Joël Ouaknine and James Worrell. 2014. Ultimate positivity is decidable for simple linear recurrence sequences. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP'14)*. DOI: https://doi.org/10.1007/978-3-662-43951-7_28

- [33] Sriram Sankaranarayanan, Henny B. Sipma, and Zohar Manna. 2005. Scalable analysis of linear systems using mathematical programming. In *Proceedings of the International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI'05)*. DOI:https://doi.org/10.1007/978-3-540-30579-8_2
- [34] Thoralf Skolem. 1934. Ein verfahren zur behandlung gewisser exponentialer gleichungen. In *Comptes Rendus du Congrès des Mathématiciens Scandinaves*.

Received 8 February 2024; revised 21 October 2024; accepted 6 November 2024