




The Pseudo-Reachability Problem for Diagonalisable Linear Dynamical Systems

Julian D’Costa  

Department of Computer Science, University of Oxford, UK

Toghrul Karimov  

Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany

Rupak Majumdar  

Max Planck Institute for Software Systems, Kaiserslautern, Germany

Joël Ouaknine  

Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany

Mahmoud Salamati  

Max Planck Institute for Software Systems, Kaiserslautern, Germany

James Worrell  

Department of Computer Science, University of Oxford, UK

Abstract

We study fundamental reachability problems on pseudo-orbits of linear dynamical systems. Pseudo-orbits can be viewed as a model of computation with limited precision and pseudo-reachability can be thought of as a robust version of classical reachability. Using an approach based on \mathcal{o} -minimality of \mathbb{R}_{exp} we prove decidability of the discrete-time pseudo-reachability problem with arbitrary semialgebraic targets for diagonalisable linear dynamical systems. We also show that our method can be used to reduce the continuous-time pseudo-reachability problem to the (classical) time-bounded reachability problem, which is known to be conditionally decidable.

2012 ACM Subject Classification Theory of computation \rightarrow Logic and verification

Keywords and phrases pseudo-orbits, Orbit problem, Skolem problem, linear dynamical systems, reachability

Digital Object Identifier 10.4230/LIPIcs.MFCS.2022.40

Related Version *Full Version*: <https://arxiv.org/abs/2204.12253>

Funding *Julian D’Costa*: emmy.network foundation under the aegis of the Fondation de Luxembourg. *Rupak Majumdar*: DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>).

Joël Ouaknine: 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.

1 Introduction

A *discrete-time linear dynamical system* (LDS) is given by an update matrix $M \in \mathbb{Q}^{d \times d}$ and a starting point $s \in \mathbb{Q}^d$. An LDS describes a system whose state contains d rational numbers and evolves linearly. The *orbit* of such a system is the infinite sequence $\langle s, Ms, M^2s, \dots \rangle$ of points in \mathbb{Q}^d . Orbits of LDS arise in many areas of computer science and mathematics, including verification of linear loops [10], automata theory [3], and the theory of linear recurrence sequences [17].

A fundamental problem about LDS is the question of deciding, given a system $\langle M, s \rangle$ and a semialgebraic target set $S \subseteq \mathbb{R}^d$, whether there exists $n \in \mathbb{N}$ such that $M^n s \in S$. This problem is known as the *reachability problem for LDS* and has been studied extensively over



© Julian D’Costa, Toghrul Karimov, Rupak Majumdar, Joël Ouaknine, Mahmoud Salamati, and James Worrell;

licensed under Creative Commons License CC-BY 4.0

47th International Symposium on Mathematical Foundations of Computer Science (MFCS 2022).

Editors: Stefan Szeider, Robert Ganian, and Alexandra Silva; Article No. 40; pp. 40:1–40:13



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the last few decades. In their seminal work, Kannan and Lipton showed that the point-to-point reachability problem, i.e., the case in which S is a singleton, is decidable in polynomial time. At the same time they observed that the case in which S is a $(d-1)$ -dimensional subspace of \mathbb{R}^d (i.e. a hyperplane) is equivalent to the famous *Skolem problem* whose decidability remains open to this day. The Skolem problem asks, given a linear recurrence sequence defined by a recurrence relation $u_{n+d} = a_1 u_n + \dots + a_d u_{n+d-1}$ and initial values u_0, \dots, u_{d-1} , to decide whether there exists n such that $u_n = 0$. In addition to this Skolem-hardness, the difficulty of settling the reachability problem was further demonstrated by the results of [16], which show that solving the reachability problem with halfspace targets, known as the *positivity problem*, would entail major mathematical breakthroughs in the field of Diophantine approximation.

The reachability problem is defined with reference to the exact dynamics of an LDS. Since computational systems typically operate with finite precision, it is natural to consider an alternate notion of reachability involving so-called *pseudo-orbits*. The notion of pseudo-orbit is an important conceptual tool in dynamical systems that was studied by Anosov [2], Bowen [4], and Conley [6], and was used by the latter to prove what is sometimes called the fundamental theorem of dynamical systems. Given an LDS $\langle M, s \rangle$, a sequence $\langle x_n \mid n \in \mathbb{N} \rangle$ is an ϵ -pseudo-orbit of s under M if $x_0 = s$ and $\|Mx_n - x_{n+1}\| < \epsilon$ for all $n \in \mathbb{N}$. In other words, in a pseudo-orbit one considers an enlarged transition relation that is obtained by considering the dynamical system up to precision ϵ . Given $\epsilon > 0$, a set S is said to be ϵ -pseudo-reachable if there exists an ϵ -pseudo-orbit $\langle x_0 = s, x_1, x_2, \dots \rangle$ of s under M that reaches S . We further say that S is pseudo-reachable if S is ϵ -pseudo-reachable for every $\epsilon > 0$. If a set S of error states is not pseudo-reachable then we can consider the system as being safe if implemented with sufficient precision, while if S is pseudo-reachable, it means that no finite amount of precision suffices to make the system reliably safe.

Recently, D’Costa et al. [7] considered the pseudo-reachability problem and, somewhat surprisingly, showed decidability in cases where S is a point (the pseudo-orbit problem), a hyperplane (the pseudo-Skolem problem) or a halfspace (the pseudo-positivity problem). Their proof of the first result relies on an exact characterisation of ϵ -pseudo-orbits. Their solution to the latter two problems, however, depends heavily on the fact that a hyperplane (a halfspace) can be defined using a single equality (inequality), an approach which unfortunately cannot be generalised to arbitrary semialgebraic targets. **In this work, we develop a novel logical approach to show the decidability of the pseudo-reachability problem for diagonalisable systems with arbitrary semialgebraic targets.**

1.1 High-level proof sketch of our approach

Our solution to the diagonalisable pseudo-reachability problem can be summarised as follows. Let $\tilde{\mathcal{O}}_\epsilon(n)$ denote the set of all points that are reachable exactly at time n via an ϵ -pseudo-orbit. The pseudo-reachability problem then consists in checking whether the sentence $\Phi := \forall \epsilon. \exists n \in \mathbb{N} : \tilde{\mathcal{O}}_\epsilon(n) \cap S \neq \emptyset$ is true. In this form, Φ is not amenable to application of logical methods as it involves both integer and real-valued variables, in addition to exponentiation with a complex base (coming from non-real eigenvalues of M). We therefore first move to the continuous domain and construct an abstraction $\mathcal{A}_\epsilon(t)$ for $t \in \mathbb{R}^{\geq 0}$ that is definable in \mathbb{R}_{exp} such that $\mathcal{A}_\epsilon(n) \supseteq \tilde{\mathcal{O}}_\epsilon(n)$ for all $n \in \mathbb{N}$. We then investigate the values of ϵ and t that make $\Psi(\epsilon, t) := \mathcal{A}_\epsilon(t) \cap S \neq \emptyset$ true. We show that by the o -minimality of \mathbb{R}_{exp} , either for every $\epsilon > 0$ there exists T such that for all $t > T$, $\Psi(\epsilon, t)$ holds, or the pseudo-reachability problem is equivalent to a finite-horizon reachability problem that is easily solvable. In the former case, it follows that for every $\epsilon > 0$, $\Psi(\epsilon, n)$ holds for all sufficiently large integer values n , thus establishing a bridge back to the discrete setting. We

conclude by showing that in this case, S is pseudo-reachable. Intuitively, the idea is to use the universal quantification over ϵ to argue that if S can be reached using an $\epsilon/2$ -abstraction at all but finitely many time steps, then it can be reached by an ϵ -pseudo-orbit, in fact at infinitely many possible time steps. The importance of the universal quantification is also illustrated by the following hardness result. For any fixed $\epsilon > 0$, it is decidable whether $\exists n \in \mathbb{N} : \mathcal{A}_\epsilon(n) \cap S \neq \emptyset$ holds, whereas the ϵ -pseudo-reachability problem of determining whether $\exists n \in \mathbb{N} : \tilde{\mathcal{O}}_\epsilon(n) \cap S \neq \emptyset$ holds is hard with respect to (a hard subclass of) the Skolem problem, as shown in Section 4.

The approach outlined above can be adapted to solve a few other related problems about linear dynamical systems. An example would be the *robust reachability problem* recently considered by Akshay et al. in [1]: given an LDS $\langle M, s \rangle$ and a semialgebraic target S , decide whether for all $\epsilon > 0$ there exists a point s' in the ϵ -neighbourhood of s whose orbit reaches S . This problem can be thought of as a modification of the pseudo-reachability problem where only one perturbation is allowed at the very beginning. Due to this simplification, we are able to show, in the full version of this paper, full decidability (that is, without the restriction to diagonalisable systems) of the robust reachability problem. Finally, as the first step of our solution is to translate the problem into the continuous domain, the continuous versions of both the pseudo-reachability problem (Section 5) and the robust reachability problem (discussed in the full version) can be handled using the same approach, arguably more naturally. For the former, because we proceed by reducing the pseudo-reachability problem to bounded-time reachability problem, the decidability result assumes Schanuel’s conjecture.

2 Mathematical tools

We write $B(c, r)$ for the closed ℓ_2 -ball of radius r centred around $c \in \mathbb{R}^d$ and $\mathbf{0} \in \mathbb{R}^d$ for the d -dimensional zero vector. We denote by $\mathbb{T} \subseteq \mathbb{C}$ the unit circle in the complex plane and by $\|x\|$ the ℓ_2 -norm of a vector $x \in \mathbb{R}^d$.

2.1 First-order logic

We denote by \mathbb{R}_0 the (structure of) real numbers with addition and multiplication, by \mathbb{R}_{exp} the real numbers with addition, multiplication and (unbounded) exponentiation and by $\mathbb{R}_{\text{exp,cos}} \upharpoonright [0, T]$ the real numbers with exponentiation and bounded (in input, by some $T > 0$) trigonometric functions. By the Tarski-Seidenberg theorem, the theory of \mathbb{R}_0 admits effective quantifier elimination and is therefore decidable. The theories of \mathbb{R}_{exp} and $\mathbb{R}_{\text{exp,cos}} \upharpoonright [0, T]$ are known to be decidable subject to Schanuel’s conjecture (see, e.g., [11]) in transcendental number theory [13, 19]. However $\mathbb{R}_{\text{exp,cos}} \upharpoonright [0, T]$ (and hence \mathbb{R}_{exp} and \mathbb{R}_0) are unconditionally known to be o -minimal [18]. That is, any subset of \mathbb{R} definable using arithmetic operations, real exponentiation and bounded trigonometric functions is a finite union of intervals. In particular, any subset of $\mathbb{R}^{\geq 0}$ definable in this way is either bounded or contains all sufficiently large real numbers. For the discrete-time problems considered in this paper we will only need to work with \mathbb{R}_{exp} . We will need $\mathbb{R}_{\text{exp,cos}} \upharpoonright [0, T]$ only when considering the classical bounded-time reachability problem for continuous-time linear dynamical systems.

A *semialgebraic* set is a subset of \mathbb{R}^d definable (without parameters) in \mathbb{R}_0 . We say that a function $\varphi : \mathbb{R}^l \rightarrow \mathbb{R}^m$ is semialgebraic if its graph is a semialgebraic subset of \mathbb{R}^{l+m} . Intuitively, semialgebraic functions are exactly the functions that can be specified using arithmetic and logical operations over the real numbers.

2.2 Kronecker's theorem and its applications

The analysis of problems about linear dynamical systems often reduces to that of the orbit $\langle \Gamma^n \mid n \in \mathbb{N} \rangle$ where $\Gamma^n = (\gamma_1^n, \dots, \gamma_k^n)$ for $\gamma_1, \dots, \gamma_k \in \mathbb{T}$. Let $\mathcal{T} = \text{cl}(\{\Gamma^n : n \in \mathbb{N}\})$ be the topological closure of this discrete orbit. The set \mathcal{T} is semialgebraic and well-understood with the help of Kronecker's theorem in simultaneous Diophantine approximation [9].

► **Theorem 1** (Kronecker). *Let $\theta_1, \dots, \theta_k, \varphi_1, \dots, \varphi_k \in \mathbb{R}$ be such that for any $a_1, \dots, a_k \in \mathbb{Z}$,*

$$\sum_{i=1}^k a_i \theta_i \in \mathbb{Z} \Rightarrow \sum_{i=1}^k a_i \varphi_i \in \mathbb{Z}.$$

For any $\epsilon > 0$ there exist infinitely many $n \in \mathbb{N}$ such that $\{n\theta_i - \varphi_i\} < \epsilon$ for all $1 \leq i \leq k$, where $\{x\}$ denotes the distance from $x \in \mathbb{R}$ to the nearest integer.

To apply this theorem to our situation, let

$$\mathcal{T} = \{(z_1, \dots, z_k) : \forall a_1, \dots, a_k \in \mathbb{Z} : \gamma_1^{a_1} \cdots \gamma_k^{a_k} = 1 \Rightarrow z_1^{a_1} \cdots z_k^{a_k} = 1\}.$$

For $z = (z_1, \dots, z_k) \in \mathcal{T}$, by considering $\theta_i = \frac{\arg(\gamma_i)}{2\pi}$ and $\varphi_i = \frac{\arg(z_i)}{2\pi}$ for $1 \leq i \leq k$ we can deduce that for each $\epsilon > 0$ there exists n such that $\|z - \Gamma^n\| < \epsilon$ and hence the orbit $\langle \Gamma^n \mid n \in \mathbb{N} \rangle$ is dense in \mathcal{T} . On the other hand, using Masser's deep results [14] about multiplicative relations between algebraic numbers one can compute, in polynomial time, a finite basis for $\{(a_1, \dots, a_k) \in \mathbb{Z}^k : \gamma_1^{a_1} \cdots \gamma_k^{a_k} = 1\}$. Hence \mathcal{T} is closed, semialgebraic and effectively computable. It then follows that $\mathcal{T} = \text{cl}(\{\Gamma^n : n \in \mathbb{N}\})$.

We will also need the following lemma which is a consequence of the effective computability of $\mathcal{T} = \text{cl}(\{\Gamma^n : n \in \mathbb{N}\})$ as a semialgebraic set.

► **Lemma 2.** *Let $R = \text{diag}(\Lambda_1, \dots, \Lambda_k) \in \mathbb{R}^{2k \times 2k}$ be a block diagonal matrix where Λ_i is an algebraic rotation matrix for $1 \leq i \leq k$. The closure of the set $\{R^n x : n \in \mathbb{N}\}$, for x with algebraic entries, is semialgebraic and effectively computable.*

The proof follows immediately from diagonalising R^n and observing that all eigenvalues of R are algebraic numbers in \mathbb{T} .

3 Decidability for discrete-time diagonalisable systems

In this section we prove our main result: the decidability of the pseudo-reachability problem for discrete-time diagonalisable *affine dynamical systems*, which are a generalisation of LDS. The reason we consider affine systems is that the well-known homogenisation trick (increasing the dimension by one and adding a coordinate that is always equal to 1) used for reducing the classical reachability problem for affine systems to the reachability problem for LDS doesn't work for the pseudo-reachability problem: when perturbations are allowed, one cannot force a coordinate to remain constant. Hence affine systems require separate treatment.

► **Problem 3** (pseudo-reachability). *Let $M \in \mathbb{Q}^{L \times L}$ be an update matrix, $s \in \mathbb{Q}^L$ be a starting point, $b \in \mathbb{Q}^L$ be an affine term and $S \subseteq \mathbb{R}^L$ be a semialgebraic target set. A sequence $\langle x_0 = s, x_1, x_2 \dots \rangle$ is an ϵ -pseudo-orbit of s if $\|Mx_n + b - x_{n+1}\| \leq \epsilon$ for all n . The pseudo-reachability problem asks: given M, b, s and S , decide whether for each $\epsilon > 0$ there exists an ϵ -pseudo-orbit of s that reaches the set S .*

Let $\tilde{\mathcal{O}}_\epsilon(n)$ denote the set of all points that are reachable via an ϵ -pseudo-orbit of s under the map $x \mapsto Mx + b$ at time n . Since $\tilde{\mathcal{O}}_\epsilon(0) = \{s\}$ and $\tilde{\mathcal{O}}_\epsilon(n+1) = M\tilde{\mathcal{O}}_\epsilon(n) + b + \epsilon B(\mathbf{0}, 1)$, by induction we can show that $\tilde{\mathcal{O}}_\epsilon(n) = M^n s + \sum_{i=0}^{n-1} M^i b + \epsilon \sum_{i=0}^{n-1} M^i B(\mathbf{0}, 1)$. The pseudo-reachability problem is then equivalent to determining the truth of $\forall \epsilon. \exists n : \tilde{\mathcal{O}}_\epsilon(n) \cap S \neq \emptyset$. Here $B(\mathbf{0}, 1)$ can be viewed as a set of ‘‘control inputs’’, and the pseudo-reachability problem can be viewed as the problem of determining whether S can be reached using arbitrarily small control inputs. The next lemma shows that we can in fact, choose any reasonable control set.

► **Lemma 4** (Invariance under change of the control set). *Let $\mathcal{B} \subseteq \mathbb{R}^L$ be a bounded set containing an open ball around the origin.*

1. *The pseudo-reachability problem as defined above is equivalent to the problem of determining whether*

$$\forall \epsilon. \exists n : (M^n s + \sum_{i=0}^{n-1} M^i b + \epsilon \sum_{i=0}^{n-1} M^i \mathcal{B}) \cap S \neq \emptyset.$$

2. *We may assume the matrix M is in real Jordan form.*

Proof. Since \mathcal{B} is assumed to be bounded and to contain an open neighbourhood around the origin, there must exist constants C_1, C_2 such that $C_1 B(\mathbf{0}, 1) \subseteq \mathcal{B} \subseteq C_2 B(\mathbf{0}, 1)$. Hence

$$C_1 \epsilon \sum_{i=0}^{n-1} M^i B(\mathbf{0}, 1) \subseteq \epsilon \sum_{i=0}^{n-1} M^i \mathcal{B} \subseteq C_2 \epsilon \sum_{i=0}^{n-1} M^i B(\mathbf{0}, 1).$$

The proof of (1) then follows from the fact that ϵ is universally quantified: one can simulate (i) an ϵ -pseudo-orbit with control set \mathcal{B} using a $C_2 \epsilon$ -pseudo-orbit with control set $B(\mathbf{0}, 1)$ and (ii) an ϵ -pseudo-orbit with control set $B(\mathbf{0}, 1)$ using a ϵ/C_1 -pseudo-orbit with control set \mathcal{B} . Proof of (2) follows from observing that multiplying $B(\mathbf{0}, 1)$ by an invertible change of basis matrix results in a bounded control set containing a neighbourhood around $\mathbf{0}$. ◀

Observe that the change of the control set described above is not applicable when ϵ is fixed, as in the ϵ -pseudo-reachability problem discussed in Section 4.

3.1 A closed form for $\tilde{\mathcal{O}}_\epsilon(n)$

We now use Lemma 4 to choose a control set that results in $\tilde{\mathcal{O}}_\epsilon(n)$ with a convenient first-order closed form: observe that the naïve formulation above involves the term $\sum_{i=0}^{n-1} M^i B(\mathbf{0}, 1)$ which is not ‘‘first-order’’.

Assume M is diagonalisable and in real Jordan form: $M = \text{diag}(\Lambda_1, \dots, \Lambda_k, \rho_{k+1}, \dots, \rho_d)$. That is, M consists of d block, the first k of which have dimension 2×2 and a pair of non-real conjugate eigenvalues, whereas the remaining blocks are 1×1 and real. Write ρ_j for the spectral radius of the j th block. We can factor M into a ‘‘scaling’’ and a ‘‘rotation’’ as $M = DR$ where $D = \text{diag}(\rho_1, \rho_1, \dots, \rho_k, \rho_k, \rho_{k+1}, \rho_{k+2}, \dots, \rho_d)$ is diagonal and R is a block-diagonal matrix that consists of blocks that are either 2×2 rotation matrices or 1×1 and equal to $[\pm 1]$. Hereafter we will be using the convenient ‘‘rotation-invariant’’ control set

$$\mathcal{B} = \prod_{j=1}^k B((0, 0), 1) \times \prod_{j=k+1}^d [-1, 1] = \prod_{j=1}^d B(\mathbf{0}, 1)$$

where $B((0,0),1)$ is the unit disc. Observe that \mathcal{B} is a product of ℓ_2 -balls that matches the block structure of M . It follows that $R\mathcal{B} = \mathcal{B}$ and hence

$$\tilde{\mathcal{O}}_\epsilon(n) = D^n R^n s + \sum_{i=0}^{n-1} M^i b + \epsilon \sum_{i=0}^{n-1} D^i R^i \mathcal{B} = D^n R^n s + \sum_{i=0}^{n-1} M^i b + \epsilon \mathcal{B}(n)$$

where $\mathcal{B}(n) = \sum_{i=0}^{n-1} D^i \mathcal{B}$. We then have

$$\mathcal{B}(n) = \sum_{i=0}^{n-1} D^i \prod_{j=1}^d B(\mathbf{0}, 1) = \sum_{i=0}^{n-1} \prod_{j=1}^d B(\mathbf{0}, \rho_j^i) = \prod_{j=0}^d B(\mathbf{0}, \sum_{i=0}^{n-1} \rho_j^i).$$

Geometrically, the idea is that a 2×2 or a 1×1 block of D maps an origin-centred disc (which corresponds to a symmetric interval in 1D) to an origin-centred disc, and a set-sum of such discs is again an origin-centred disc. Note that our ability to reason in this way crucially depends on the fact that M is diagonalisable. Finally, since $\sum_{i=0}^{n-1} \rho_j^i$ is either $\frac{\rho_j^n - 1}{\rho_j - 1}$ or $n\rho_j$, we can write $\mathcal{B}(n) = \{z : \varphi(z, n, \rho_1^n, \dots, \rho_d^n)\}$, where φ is a semialgebraic predicate.

We can apply the blockwise summation technique, distinguishing between the cases where the spectral radius of the block is 1 or different from 1, to the term $\sum_{i=0}^{n-1} M^i b$ to obtain the closed form $\sum_{i=0}^{n-1} M^i b = D^n R^n x' + cn + d$, where x' , c and d only depend on M and b . We then fold s and x' into a new, fictive starting point x to obtain the final closed form

$$\tilde{\mathcal{O}}_\epsilon(n) = D^n R^n x + cn + d + \epsilon \mathcal{B}(n).$$

In order to solve the pseudo-reachability problem, we henceforth consider the problem of determining the truth of the sentence $\forall \epsilon > 0. \exists n : (D^n R^n x + cn + d + \epsilon \mathcal{B}(n)) \cap S \neq \emptyset$, where all the input vectors and matrices have real algebraic entries.

3.2 Passing to the abstraction

The expression for $\tilde{\mathcal{O}}_\epsilon(n)$ contains the term $D^n R^n x$, which is the last obstacle to obtaining an expression which we can attack using known results about theories of real numbers. To address this issue we resort to abstracting $\tilde{\mathcal{O}}_\epsilon(n)$. Let

$$\mathcal{T} := \text{cl}(\{R^n x : n \in \mathbb{N}\}) \text{ and } \mathcal{A}_\epsilon(n) := D^n \mathcal{T} + cn + d + \epsilon \mathcal{B}(n)$$

where \mathcal{T} is the closure of the orbit of x under R , and is semialgebraic and effectively computable by the discussion in Subsection 2.2. Moreover, recall that by Kronecker's theorem for every $z \in \mathcal{T}$ and $\epsilon > 0$ there exist infinitely many integers $0 < n_1 < n_2 < \dots$ such that $\|R^{n_i} x - z\| < \epsilon$ for all i .

Here $\mathcal{A}_\epsilon(n)$ acts as an abstraction of $\tilde{\mathcal{O}}_\epsilon(n)$. In particular, for all $\epsilon > 0$ and $n \in \mathbb{N}$ we have $\mathcal{A}_\epsilon(n) \supseteq \tilde{\mathcal{O}}_\epsilon(n)$. Observe that $\mathcal{A}_\epsilon(n) = \{z : \varphi(z, \epsilon, n, \rho_1^n, \dots, \rho_d^n)\}$ for a semialgebraic predicate φ . Viewing $\mathcal{A}_\epsilon(n)$ as a proxy for $\tilde{\mathcal{O}}_\epsilon(n)$, we arrive at the following dichotomy.

► **Lemma 5.** *Either*

1. *for every $\epsilon > 0$ there exists N_ϵ such that for all $n > N_\epsilon$, $\mathcal{A}_\epsilon(n)$ intersects S , or*
2. *there exist N and $\epsilon > 0$, both effectively computable, such that $\mathcal{A}_\epsilon(n)$ does not intersect S for all $n > N$.*

Moreover, it can be effectively determined which case holds.

Proof. First we show that the dichotomy holds, putting the issues of effectiveness aside. Let

$$\Phi(\epsilon, n) = \bigvee_{\alpha \in A} \bigwedge_{\beta \in B} p_{\alpha, \beta}(\epsilon, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$$

be a quantifier-free formula equivalent to $\mathcal{A}_\epsilon(n) \cap S \neq \emptyset$. Such $\Phi(\epsilon, n)$ must exist because $\mathcal{A}_\epsilon(n)$ is semialgebraic with parameters from $\{\epsilon, n, \rho_1^n, \dots, \rho_d^n\}$ and by the Tarski-Seidenberg theorem, each such set can be described using a quantifier-free formula of the form given above. Suppose Case 1 does not hold. Then there exists a particular $\epsilon > 0$ such that $\Phi(\epsilon, n)$ does not hold for arbitrarily large n . Treating n as a continuous parameter, consider the set $\{n \in \mathbb{R}_{\geq 0} : \Phi(\epsilon, n) \text{ does not hold}\}$. By o -minimality of \mathbb{R}_{exp} this set is a finite union of intervals and by the assumption that Case 1 does not hold, it contains arbitrarily large integers. Hence it must contain all integers in (N, ∞) for some $N \in \mathbb{N}$. That is, for all $n > N$ the formula $\Phi(\epsilon, n)$ does not hold.

Effectiveness. We now address the issues of effectiveness. Consider the formula

$$\Psi(\epsilon) = \exists N_\epsilon. \forall n > N_\epsilon : \Phi(\epsilon, n).$$

We show that $\Psi(\epsilon)$ is equivalent to a formula $\psi(\epsilon)$ in the language of \mathbb{R}_0 . To determine which case holds it then remains to determine the truth value of the sentence $\forall \epsilon : \psi(\epsilon)$.

By the o -minimality argument above, given $\epsilon > 0$, each $p_{\alpha, \beta}(\epsilon, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$ either holds for finitely many integer values of n or holds for all sufficiently large integer values n . By elementary considerations it follows that $\Psi(\epsilon)$ is equivalent to

$$\bigvee_{\alpha \in A} \bigwedge_{\beta \in B} \exists N_\epsilon. \forall n > N_\epsilon : p_{\alpha, \beta}(\epsilon, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0.$$

Hence it suffices to show how to construct a formula $\psi(\epsilon)$ in the language of \mathbb{R}_0 that is equivalent to $\exists N_\epsilon. \forall n > N_\epsilon : p_{\alpha, \beta}(\epsilon, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$. For each $\epsilon > 0$, the formula first tests if $p_{\alpha, \beta}(\epsilon)$ (as a polynomial in $d+1$ remaining variables) is identically zero. If yes, then $\varphi(\epsilon)$ is true or false depending only on $\bowtie_{\alpha, \beta}$. Otherwise, write $p_{\alpha, \beta}(\epsilon, n, \rho_1^n, \dots, \rho_d^n) = \sum_{i=1}^k q_i(\epsilon, n) R_i^n$ where $q_i(\epsilon)$ is not identically zero for all i and $R_1 > \dots > R_k > 0$ are real algebraic numbers of the form $\rho_1^{p_1} \dots \rho_d^{p_d}$ for $p_1, \dots, p_d \in \mathbb{N}$. Since $|q_1(\epsilon, n) R_1^n| > \left| \sum_{i=2}^k q_i(\epsilon, n) R_i^n \right|$ for sufficiently large n , whether $p_{\alpha, \beta}(\epsilon, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$ holds for sufficiently large n depends only on $q_1(\epsilon, n)$. Hence we can choose $\psi(\epsilon)$ to be $\lim_{n \rightarrow \infty} q_1(\epsilon) \bowtie_{\alpha, \beta} 0$, which amounts to a sign condition on the coefficients of $q_1(\epsilon, n)$.

Computing N . Finally, we show that in Case 2, the value N can be effectively computed. To this end, by repeatedly trying smaller and smaller values of ϵ first compute a rational $e > 0$ such that $\Psi(e)$ (equivalently, $\psi(e)$) does not hold. To be able to compute N it then suffices to compute, for a particular (α, β) , a value $N_{\alpha, \beta}$ such that $p_{\alpha, \beta}(e, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$ does not hold for all $n > N_{\alpha, \beta}$, assuming that it does not hold for sufficiently large n . We can then take N to be the maximum of $N_{\alpha, \beta}$ over $(\alpha, \beta) \in A \times B$.

To compute $N_{\alpha, \beta}$, consider $p := p_{\alpha, \beta}(e)$. Assuming it is not identically zero (otherwise we can choose $N_{\alpha, \beta}$ to be any positive integer), write $p(n, \rho_1^n, \dots, \rho_d^n) = \sum_{i=1}^k q_i(n) R_i^n$ where q_i is not identically zero for all i and $R_1 > \dots > R_k > 0$ are real algebraic. Since $p_{\alpha, \beta}(e, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$ and hence $p(n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} 0$ do not hold for sufficiently large n , it must be the case that $q_1(n) \bowtie_{\alpha, \beta} 0$ does not hold for sufficiently large n . Hence it remains to choose $N_{\alpha, \beta}$ large enough so that for all $n > N_{\alpha, \beta}$, $|q_1(n) R_1^n|$ dominates $\left| \sum_{i=2}^k q_i(n) R_i^n \right|$. \blacktriangleleft

3.3 From the abstraction back to ϵ -pseudo-orbits

In this section we consider the relationship between the two cases of Lemma 5 and our original pseudo-reachability problem. We start with Case 2. Observe that $\mathcal{O}_\epsilon(n) \subseteq \mathcal{A}_\epsilon(n)$ for every $n \in \mathbb{N}$ and $\epsilon > 0$. Therefore, when Case 2 holds, for any $n > N$ and $\epsilon > 0$ the target set cannot be reached by \mathcal{O}_ϵ . It remains to check pseudo-reachability at time steps $n \leq N$. We claim that S is pseudo-reachable if and only

$$\forall \epsilon. \exists n \leq N : (M^n x + cn + d + \epsilon \mathcal{B}(n)) \cap S \neq \emptyset.$$

Let \bar{S} denote the topological closure of S . We show that the statement above is equivalent to $\exists n \leq N : M^n x + cn + d \in \bar{S}$. Observe that if for all $n \leq N$ the point $M^n x + cn + d$ is not in \bar{S} , then by compactness the smallest distance from $\{M^n x + cn + d \mid n \leq N\}$ to \bar{S} is positive and hence for sufficiently small ϵ the target S cannot be ϵ -pseudo-reached within the first N steps. Conversely, if $M^n x + cn + d \in \bar{S}$ for some $n \leq N$, then because $\mathcal{B}(n)$ is full dimensional and contains $\mathbf{0}$ in its interior, it follows that $(M^n x + cn + d + \epsilon \mathcal{B}(n)) \cap S \neq \emptyset$ for all $\epsilon > 0$. Therefore, in Case 2 pseudo-reachability can be decided by simply checking if $\{M^n x + cn + d \mid n \leq N\}$ reaches \bar{S} .

Next we will show that S is pseudo-reachable if Case 1 holds. Given $z \in \mathcal{T}$, we define a “localisation” of the abstraction at the point z as $\mathcal{A}_\epsilon(n)(z) := D^n z + cn + d + \epsilon \mathcal{B}(n)$. Observe that $\mathcal{A}_\epsilon(n) = \{\mathcal{A}_\epsilon(n)(z) : z \in \mathcal{T}\}$. This definition of a localisation will allow us to select a “concrete trajectory” from the set of all possible (abstract) trajectories.

Fix $\epsilon > 0$ and let $T_n := \{z \in \mathcal{T} : \mathcal{A}_\epsilon(n)(z) \text{ intersects } S\}$. The next lemma implies that the sequence T_n must tend towards a limiting shape; i.e. it cannot “jump around” forever.

► **Lemma 6.** *Let $T_n = \{z : \varphi(z, n, \rho_1^n, \dots, \rho_d^n)\}$, where φ is a semialgebraic predicate and ρ_1, \dots, ρ_d are real algebraic, be a family of non-empty sets contained in a compact set \mathcal{T} . There exists a non-empty limiting set $L \subseteq \mathcal{T}$ to which the sequence T_n converges as $n \rightarrow \infty$, in the following sense.*

a *For every $\epsilon > 0$, there exists N such that for all $n > N$, $T_n \subseteq L + B(\mathbf{0}, \epsilon)$.*

b *For all $z \in L$ and $\epsilon > 0$ there exists N such that for all $n > N$, $z + B(\mathbf{0}, \epsilon)$ intersects T_n .*

Proof. Write $\varphi(z, n, \rho_1^n, \dots, \rho_d^n) = \bigvee_{\alpha \in A} \bigwedge_{\beta \in B} p_{\alpha, \beta}(z, n, \rho_1^n, \dots, \rho_d^n) \bowtie_{\alpha, \beta} \mathbf{0}$. We can define the sequence $\langle T_t \mid t \in \mathbb{R} \rangle$ as $T_t = \{z : \varphi(z, t, \rho_1^t, \dots, \rho_d^t)\}$. Let $L = \{x : \liminf d(x, T_t) = 0\}$ where $d(x, T_t)$ denotes the shortest Euclidean distance from x to a point in T_t .

We prove the first claim by contradiction. Suppose there exists $\epsilon > 0$ such that at infinitely many unbounded time steps $t_1 < t_2 < \dots$ there are points z_1, z_2, \dots such that $z_i \in T_{t_i}$ but $z_i \notin L + B(\mathbf{0}, \epsilon)$. Then the sequence z_i must have an accumulation point z in $\mathcal{T} \setminus L$. But z will also satisfy $\liminf d(z, T_t) = 0$ and hence $z \in L$, a contradiction.

We prove the second claim using \mathcal{o} -minimality of \mathbb{R}_{exp} . Fix $z \in L$ and $\epsilon > 0$ and consider the set $Z = \{t \in \mathbb{R} : z + B(\mathbf{0}, \epsilon) \text{ intersects } T_t\}$. The set Z is \mathcal{o} -minimal, and since $z \in L$, it is unbounded from above. Hence it must contain an interval of the form (N, ∞) , which implies the desired result. ◀

One can also show that the set L described above is in fact semialgebraic, but this is not necessary for our arguments. We are now ready to show that S is pseudo-reachable if Case 1 of Lemma 5 holds.

► **Lemma 7.** *If for every $\epsilon > 0$ there exists N_ϵ such that for all $n > N_\epsilon$, $\mathcal{A}_\epsilon(n)$ intersects S then S is pseudo-reachable.*

The main idea of the proof is to use the assumption that $\mathcal{A}_{\epsilon/2}(n)$ intersects S for sufficiently large n to construct an ϵ -pseudo-orbit that hits S . Intuitively, in order to simulate $\mathcal{A}_{\epsilon/2}(n)$ using an ϵ -pseudo-orbit, $\epsilon/2$ of the total control allowance is used to replicate the effect of the control inputs (of size at most $\epsilon/2$, corresponding to the $\frac{\epsilon}{2}\mathcal{B}(n)$ term in the definition of $\mathcal{A}_{\epsilon/2}(n)$) and the remaining $\epsilon/2$ is used to compensate for the abstraction from the starting point a to the set \mathcal{T} . In fact, we do not know if one can deduce that S is ϵ -pseudo-reachable from knowing that $\mathcal{A}_{\epsilon}(n)$ intersects S for sufficiently large n . This illustrates the reason why the pseudo-reachability problem is easier than the ϵ -pseudo-reachability problem; see Section 4 for a more concrete argument.

Proof. Fix $\epsilon > 0$. We show how to construct an ϵ -pseudo-orbit that hits S . Consider $\mathcal{A}_{\epsilon/2}(n)$. By assumption, there exists N_1 such that for all $n > N_1$, $\mathcal{A}_{\epsilon/2}(n)$ intersects S . We now investigate which localisations of the abstraction are responsible for intersecting S . Apply Lemma 6 to the sequence of sets $T_n = \{z \in \mathcal{T} : \mathcal{A}_{\epsilon/2}(n)(z) \text{ intersects } S\}$ to obtain their “limit” L . Fix any $p \in L$.

Let ϵ' be small enough so that $\epsilon'D^n\mathcal{B} \subseteq \frac{\epsilon}{2}\mathcal{B}(n)$ for all $n > 0$. Intuitively, such ϵ' must exist because $D^n\mathcal{B}$ and $D^{n-1}\mathcal{B}$ only differ by at most a constant factor that only depends on the magnitudes ρ_1, \dots, ρ_d of eigenvalues of M , and we have that $D^{n-1}\mathcal{B} \subseteq \sum_{i=0}^{n-1} D^i\mathcal{B} = \mathcal{B}(n)$. By Lemma 6 (b), there exists $N > N_1$ such that for all $n > N$, $p + B(\mathbf{0}, \epsilon'/2)$ intersects T_n . That is, for all $n > N$ there exists $p_n \in \mathcal{T}$ such that $\|p - p_n\| < \epsilon'/2$ and $p_n \in T_n$. Equivalently,

$$\|p - p_n\| < \epsilon'/2 \text{ and } \mathcal{A}_{\epsilon/2}(n)(p_n) \text{ intersects } S.$$

By Kronecker’s theorem there must exist $m > N$ such that $\|R^m x - p\| < \epsilon'/2$. Hence we have $\|R^m x - p_m\| < \epsilon'$ which implies $R^m x - p_m \in \epsilon'\mathcal{B}$ and hence $D^m(R^m x - p_m) \in \epsilon'D^m\mathcal{B}$. Since by construction of ϵ' we have $\epsilon'D^m\mathcal{B} \subseteq \frac{\epsilon}{2}\mathcal{B}(m)$, it follows that $D^m(R^m x - p_m) \in \frac{\epsilon}{2}\mathcal{B}(m)$ and hence $D^m p_m \in D^m R^m x + \frac{\epsilon}{2}\mathcal{B}(m)$. Therefore,

$$\tilde{\mathcal{O}}_\epsilon(m) = (D^m R^m x + \frac{\epsilon}{2}\mathcal{B}(m)) + cm + d + \frac{\epsilon}{2}\mathcal{B}(m) \supseteq D^m p_m + cm + d + \frac{\epsilon}{2}\mathcal{B}(m) = \mathcal{A}_{\epsilon/2}(m)(p_m).$$

Since $\mathcal{A}_{\epsilon/2}(m)(p_m)$ intersects S , it then follows that $\tilde{\mathcal{O}}_\epsilon(m)$ too must intersect S . \blacktriangleleft

3.4 The algorithm

To summarise, the analysis above gives us the following algorithm for determining if S is pseudo-reachable, i.e. if $\forall \epsilon > 0. \exists n : \tilde{\mathcal{O}}_\epsilon(n) \cap S \neq \emptyset$. Let $\varphi(n, \epsilon)$ be a quantifier-free formula in \mathbb{R}_{exp} defining the abstraction $\mathcal{A}_\epsilon(n)$. First determine, using the algorithm described in the proof of Lemma 5, whether Case 1 or Case 2 holds. If the former holds, then conclude that S is pseudo-reachable. If Case 2 holds, then compute the value of N effectively and check if there exists $n < N$ such that $(M^n x + cn + d) \cap \bar{S} \neq \emptyset$.

4 Skolem-hardness of the ϵ -pseudo-reachability problem

In this section we consider the ϵ -pseudo-reachability problem for discrete diagonalisable systems: given diagonalisable M , starting point s , a target set S and $\epsilon > 0$, decide whether there exists n such that $M^n s + \sum_{k=0}^{n-1} M^k B(\mathbf{0}, \epsilon) \cap S \neq \emptyset$. This problem is also known as the reachability problem for linear time-invariant systems [8] with the control set $B(\mathbf{0}, \epsilon)$. We will reduce a hard subclass of the Skolem problem to our ϵ -pseudo-reachability problem.

40:10 The Pseudo-Reachability Problem for Diagonalisable Linear Dynamical Systems

The Skolem problem is not known to be decidable for orders $d \geq 5$, even for diagonalisable recurrences. The largest class of sequences for which decidability is known is the MSTV (Mignotte-Shorey-Tijdeman-Vereschagin) class, which consists of all linear recurrence sequences over integers that (i) have at most three dominant roots with respect to the usual (Archimedean) absolute or (ii) have at most two dominant roots with respect to a p -adic absolute value [12]. We consider the Skolem problem for integer sequences whose roots $\rho, \lambda_1, \dots, \lambda_d$ satisfy $\rho = |\lambda_1| = \dots = |\lambda_d|$. This class of sequences contains infinitely many instances that are not in the MSTV class and hence is a hard subclass of the Skolem problem.

Recall that any linear recurrence sequence can be written as $u_n = c^\top M^n s$ where M is the companion matrix of u_n whose eigenvalues are the roots of u_n . Let u_n be a diagonalisable sequence that belongs to the hard subclass described above, i.e. $u_n = c^\top M^n s$ where $M = \text{diag}(\Lambda_1, \dots, \Lambda_d, \rho)$ and Λ_i is a 2×2 real Jordan block with $\rho(\Lambda_i) = \rho$ for $1 \leq i \leq d$. We reduce the problem “does u_n have a zero?” to an ϵ -pseudo-reachability problem.

Consider the sequence $v_n = u_n^2$. Observe that $v_n = \sum_{i=1}^L c_i \Gamma_i^n s_i + C r^n$ where

- $r = \rho^2$,
- Γ_i is a 2×2 real Jordan block with $\rho(\Gamma_i) = r$ for $1 \leq i \leq L$,
- $c_i, s_i \in \mathbb{R}^2$ for $1 \leq i \leq L$, and
- $C > 0$.

The first two statements follow from the fact that the eigenvalues of v_n are products of eigenvalues of u_n . That $C > 0$ can be deduced as follows. Consider $w_n = \sum_{i=1}^L c_i \Gamma_i^n s_i$. It only has non-real roots and hence by [15] is infinitely often positive and negative. Hence if C is not positive, then $v_n < 0$ for infinitely many n , which contradicts the fact that $v_n \geq 0$.

Next observe that u_n has a zero iff there exists n such that $v_n \leq 0$. Since we are interested only in the sign of v_n , by scaling v_n by $C(2r)^n$ if necessary we assume that $r \in (0, 1)$ and $C = 1$. We will construct an instance of the ϵ -pseudo-reachability problem that is positive if and only if there exists n such that $v_n \leq 0$.

Define

- $A = \text{diag}(\Gamma_1, \dots, \Gamma_L)$,
- $s = (s_1, \dots, s_L)$ and $c = (c_1, \dots, c_L)$,
- $\epsilon = \frac{1-r}{\|c\|}$, and
- $H = \{z : c^\top \cdot z + 1 \leq 0\}$.

Observe that H is ϵ -pseudo-reachable if and only if $A^n s + \sum_{i=0}^{n-1} A^i B(\mathbf{0}, \epsilon) \cap H \neq \emptyset$ for some n . Since $A^i B(\mathbf{0}, \epsilon) = B(\mathbf{0}, r^i \epsilon)$, we have $\sum_{i=0}^{n-1} A^i B(\mathbf{0}, \epsilon) = B(\mathbf{0}, \frac{1-r^n}{1-r} \epsilon) := \mathcal{B}(n)$ and

$$H \text{ is } \epsilon\text{-pseudo-reachable} \iff \min_{z \in \mathcal{B}(n)} c \cdot (A^n s + z) + 1 \leq 0 \text{ for some } n.$$

We will show that in fact $\min_{z \in \mathcal{B}(n)} c \cdot (A^n s + z) + 1 = v_n$, which will conclude the proof.

$$\begin{aligned} \min_{z \in \mathcal{B}(n)} c \cdot (A^n s + z) + 1 &= \sum_{i=1}^L c_i \Gamma_i^n s_i + 1 + \min_{z \in \mathcal{B}(n)} c \cdot z \\ &= \sum_{i=1}^L c_i \Gamma_i^n s_i + 1 - \|c\| \frac{1-r^n}{1-r} \epsilon \\ &= \sum_{i=1}^L c_i \Gamma_i^n s_i + r^n \\ &= v_n. \end{aligned}$$

5 The continuous-time pseudo-reachability problem

In this section we show that the approach we described in Section 3 for deciding the discrete-time pseudo-reachability problem for diagonalisable systems also works in the continuous setting with one important difference: to handle Case 2 of the dichotomy lemma (exactly the same as Lemma 5) we need to solve the *bounded-time reachability problem for continuous-time affine dynamical systems*, which is only known to be decidable assuming Schanuel's conjecture [5]. For detailed proofs see the full version of the paper.

Let $M = \text{diag}(\Lambda_1, \dots, \Lambda_k, \rho_{k+1}, \dots, \rho_d) \in (\mathbb{R} \cap \overline{\mathbb{Q}})^{L \times L}$ be a diagonalisable matrix in real Jordan form, $s \in \mathbb{Q}^L$ be a starting point, $b \in \mathbb{Q}^L$ be an affine term and $S \subseteq \mathbb{R}^L$ be a semialgebraic target set. The trajectory of the system (in the absence of additional control inputs) is given by

$$x(t) = e^{Mt}s + \int_0^t e^{Mh}b \, dh.$$

Intuitively, while in the discrete setting control inputs are applied after each unit of time and thus are represented by a sequence $\langle d_n \mid n \in \mathbb{N} \rangle$, in the continuous setting they are represented by a continuous function $\Delta : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^L$. Hence an ϵ -pseudo-orbit is defined as a trajectory

$$x(t) = e^{Mt}s + \int_0^t e^{Mh}b \, dh + \int_0^t e^{Mh}\Delta(t-h) \, dh.$$

for some control signal $\Delta : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}^L$ satisfying $\|\Delta_\epsilon(t)\| \leq \epsilon$ for all $t \geq 0$. The pseudo-reachability problem is then defined in the same way as before: decide whether for every $\epsilon > 0$ there exists an ϵ -pseudo-orbit that reaches S .

Let \mathcal{B} be the same control set as defined in Subsection 3.1. For $1 \leq j \leq k$, let $r_j = \text{Re}(\lambda_j)$ and $\omega_j = \text{Im}(\lambda_j)$ where λ_j is a non-real eigenvalue of the block Λ_j . For $k < j \leq d$ let $r_j = \rho_j$ and $\omega_j = 0$. By using essentially the same arguments as in Subsection 3.1, we can show that the pseudo-reachability problem is equivalent to determining the truth of

$$\forall \epsilon > 0. \exists t : (e^{Mt}x + ct + d + \epsilon\mathcal{B}(t)) \cap S \neq \emptyset$$

where x, c, d are L -dimensional vectors and $\mathcal{B}(t) = \{z : \varphi(z, t, e^{r_1 t}, \dots, e^{r_d t})\}$ for semialgebraic predicate φ . We denote the term $e^{Mt}x + ct + d + \epsilon\mathcal{B}(t)$ by $\tilde{\mathcal{O}}_\epsilon(t)$.

To define a convenient abstraction, we again write $e^{Mt} = D(t)R(t)$ where $D(t) := \text{diag}(e^{r_1 t}, e^{r_1 t}, \dots, e^{r_k t}, e^{r_k t}, e^{r_{k+1} t}, e^{r_{k+1} t}, \dots, e^{r_{k+d} t})$ is diagonal and $R(t)$ is a block diagonal matrix whose blocks are rotation matrices of the form $\begin{bmatrix} \cos(\omega_j t) & -\sin(\omega_j t) \\ \sin(\omega_j t) & \cos(\omega_j t) \end{bmatrix}$ for $1 \leq j \leq k$ and are of the form $\Omega_i = [1]$ for $k+1 \leq j \leq d$. Just as in the discrete case, we next define

$$\mathcal{T} := \text{cl}(\{R(t)x : t \in \mathbb{R}_{\geq 0}\}) \text{ and } \mathcal{A}_\epsilon(t) := D(t)\mathcal{T} + ct + d + \epsilon\mathcal{B}(t),$$

where \mathcal{T} is again semialgebraic and effectively computable [5] and \mathcal{A}_ϵ acts as an abstraction of $\tilde{\mathcal{O}}_\epsilon(t)$. In particular, for all $\epsilon > 0$ and $t \in \mathbb{R}_{\geq 0}$, we have $\tilde{\mathcal{O}}_\epsilon(t) \subseteq \mathcal{A}_\epsilon(t)$. Moreover, observe that $\mathcal{A}_\epsilon(t) = \varphi(t, e^{r_1 t}, \dots, e^{r_d t})$ for a semialgebraic function φ , which is the most important property we need. In the full version, using the same approach based on o -minimality we show the following.

► **Theorem 8.** *The continuous-time pseudo-reachability problem reduces to the bounded-time reachability problem for continuous-time affine dynamical systems.*

Intuitively, the dichotomy lemma (Lemma 5) holds verbatim for the continuous systems, and in Case 1 again S is always pseudo-reachable. It then remains to handle Case 2. Since bounded-time reachability problem for continuous-time affine dynamical systems can be encoded in $\mathbb{R}_{\text{exp,cos}} \upharpoonright_{[0,T]}$, we have the following (conditional) decidability result.

► **Corollary 9.** *Continuous-time pseudo-reachability problem for diagonalisable affine dynamical systems is decidable subject to Schanuel’s conjecture.*

6 Discussion

The main technical result of our paper is that it is decidable whether

$$\forall \epsilon > 0. \exists n : (M^n x + f(n) + \epsilon \mathcal{B}(n)) \cap S \neq \emptyset$$

where M is a diagonalisable matrix with algebraic entries, x is an algebraic starting point, f is a semialgebraic function, S is a semialgebraic target and $\mathcal{B}(n) = \varphi(n, \rho_1^n, \dots, \rho_d^n)$ for $\rho_1, \dots, \rho_d \in \mathbb{R} \cap \overline{\mathbb{Q}}$ and a semialgebraic function φ . We used this result to show decidability of the discrete-time pseudo-reachability problem for diagonalisable systems in the following way. We first observed that the pseudo-reachability problem can be cast as the problem of determining whether $\forall \epsilon > 0. \exists n : \tilde{\mathcal{O}}_\epsilon(n) \cap S \neq \emptyset$, where $\tilde{\mathcal{O}}_\epsilon(n)$ is the set of all points that are reachable exactly at the time n via an ϵ -pseudo-orbit. After choosing \mathcal{B} as the most convenient control set (see Lemma 4 and Subsection 3.1), we then showed that $\tilde{\mathcal{O}}_\epsilon(n)$ can be written as $M^n x + f(n) + \epsilon \mathcal{B}(n)$.

The reason we are unable to show decidability for non-diagonalisable systems in this fashion is that we are unable to write $\tilde{\mathcal{O}}_\epsilon(n)$ as $M^n x + f(n) + \epsilon \mathcal{B}(n)$. For example, already for the Jordan block $M = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$, and in general already for blocks with a single repeated real eigenvalue, we do not know whether it is even possible to eliminate the summation $\sum_{i=0}^{n-1} M^i B$ and express $\tilde{\mathcal{O}}_\epsilon(n) = M^n x + \epsilon \sum_{i=0}^{n-1} M^i B$, where B is any full dimensional shape containing $\mathbf{0}$ in its interior, in the required fashion.

Our approach, however, can be used to solve, in full generality, the robust reachability problem of [1]: given M, x and S , decide whether $\forall \epsilon > 0 : \exists n : (M^n x + \epsilon M^n B(\mathbf{0}, 1)) \cap S \neq \emptyset$. Intuitively, the reason is that in this version there is no summation of the form $\sum_{i=0}^{n-1} M^i B$. Detailed proofs (for both the discrete-time and the continuous-time versions) can be found in the full version. For diagonalisable systems in particular, decidability of the robust reachability problem is almost immediate. First, one can again show that the problem is equivalent to determining whether $\forall \epsilon > 0 : \exists n : (M^n x + \epsilon M^n \mathcal{B}) \cap S \neq \emptyset$. It then remains to observe that $M^n \mathcal{B} = \varphi(n, \rho_1^n, \dots, \rho_d^n)$ for a semialgebraic predicate φ and apply the technical result described above.

References

- 1 S. Akshay, Hugo Bazille, Blaise Genest, and Mihir Vahanwala. On Robustness for the Skolem and Positivity Problems. In *39th International Symposium on Theoretical Aspects of Computer Science (STACS 2022)*, volume 219 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:20, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi : 10.4230/LIPIcs.STACS.2022.5.
- 2 Dmitri V. Anosov. Geodesic flows on closed Riemannian manifolds of negative curvature. *Proc. Steklov Inst. Math.*, 90, 1967.

- 3 Corentin Barloy, Nathanaël Fijalkow, Nathan Lhote, and Filip Mazowiecki. A Robust Class of Linear Recurrence Sequences. In *28th EACSL Annual Conference on Computer Science Logic (CSL 2020)*, volume 152 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 9:1–9:16, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.CSL.2020.9.
- 4 Rufus Bowen. *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms*, volume 470 of *Lecture Notes in Mathematics*. Springer-Verlag, 1975.
- 5 Ventsislav Chonev, Joël Ouaknine, and James Worrell. On the Skolem Problem for Continuous Linear Dynamical Systems. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 100:1–100:13, Dagstuhl, Germany, 2016. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ICALP.2016.100.
- 6 Charles C. Conley. *Isolated invariant sets and the Morse index*, volume 25 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, 1978.
- 7 Julian D’Costa, Toghru Karimov, Rupak Majumdar, Joël Ouaknine, Mahmoud Salamati, Sadegh Soudjani, and James Worrell. The Pseudo-Skolem Problem is Decidable. In *46th International Symposium on Mathematical Foundations of Computer Science (MFCS 2021)*, volume 202 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 34:1–34:21, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.MFCS.2021.34.
- 8 Nathanaël Fijalkow, Joël Ouaknine, Amaury Pouly, João Sousa-Pinto, and James Worrell. On the decidability of reachability in linear time-invariant systems. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC ’19*, pages 77–86, New York, NY, USA, 2019. Association for Computing Machinery. doi:10.1145/3302504.3311796.
- 9 Godfrey H. Hardy and Edward M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, 1999.
- 10 Toghru Karimov, Engel Lefauchaux, Joël Ouaknine, David Purser, Anton Varonka, Markus A. Whiteland, and James Worrell. What’s decidable about linear loops? *Proc. ACM Program. Lang.*, 6(POPL), January 2022. doi:10.1145/3498727.
- 11 Serge Lang. *Introduction to transcendental numbers*. Addison-Wesley series in mathematics. Addison-Wesley Pub. Co., 1966.
- 12 Richard Lipton, Florian Luca, Joris Nieuwveld, Joël Ouaknine, David Purser, and James Worrell. On the skolem problem and the skolem conjecture. *To appear in LICS 2022*, 2022. URL: https://people.mpi-sws.org/~joel/publications/skolem_five22.pdf.
- 13 Angus Macintyre and Alex J. Wilkie. On the decidability of the real exponential field. In *Kreiseliana. About and Around Georg Kreisel*, pages 441–467. A K Peters, 1996.
- 14 David W. Masser. Linear relations on algebraic groups. In *New Advances in Transcendence Theory*. Camb. Univ. Press, 1988.
- 15 Kenji Nagasaka and Jau-Shyong Shiue. Asymptotic positiveness of linear recurrence sequences. *Fibonacci Quart*, 28(4):340–346, 1990.
- 16 Joël Ouaknine and James Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 366–379, 2014.
- 17 Joël Ouaknine and James Worrell. On linear recurrence sequences and loop termination. *ACM SIGLOG News*, 2(2):4–13, April 2015. doi:10.1145/2766189.2766191.
- 18 Laurentius Petrus Dignus “Lou” van den Dries. *Tame Topology and O-minimal Structures*. Cambridge University Press, 1998.
- 19 Alex J. Wilkie. Schanuel’s conjecture and the decidability of the real exponential field. In *Algebraic Model Theory*, pages 223–230. Springer Netherlands, Dordrecht, 1997.