

ON ROBUSTNESS FOR THE SKOLEM, POSITIVITY AND ULTIMATE POSITIVITY PROBLEMS

S. AKSHAY ^a, HUGO BAZILLE ^b, BLAISE GENEST ^c, AND MIHIR VAHANWALA ^d

^a Indian Institute of Technology Bombay, Mumbai, India
e-mail address: akshayss@cse.iitb.ac.in

^b Laboratoire de Recherche de l’Epita (LRE), Rennes, France
e-mail address: hugo@lrde.epita.fr

^c CNRS, IRL 2955 IPAL, Singapore
e-mail address: blaise.genest@cnrs.fr

^d Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany
e-mail address: mvahanwa@mpi-sws.org

ABSTRACT. The Skolem problem is a long-standing open problem in linear dynamical systems: can a linear recurrence sequence (LRS) ever reach 0 from a given initial configuration? Similarly, the positivity problem asks whether the LRS stays positive from an initial configuration. Deciding Skolem (or positivity) has been open for half a century: the best known decidability results are for LRS with special properties (e.g., low order recurrences). On the other hand, these problems are much easier for “uninitialised” variants, where the initial configuration is not fixed but can vary arbitrarily: checking if there is an initial configuration from which the LRS stays positive can be decided by polynomial time algorithms (Tiwari in 2004, Braverman in 2006).

In this paper, we consider problems that lie between the initialised and uninitialised variants. More precisely, we ask if 0 (resp. negative numbers) can be avoided from every initial configuration in a neighbourhood of a given initial configuration. This can be considered as a robust variant of the Skolem (resp. positivity) problem. We show that these problems lie at the frontier of decidability: if the neighbourhood is given as part of the input, then robust Skolem and robust positivity are Diophantine hard, i.e., solving either would entail major breakthroughs in Diophantine approximations, as happens for (non-robust) positivity. Interestingly, this is the first Diophantine hardness result on a variant of the Skolem problem. On the other hand, if one asks whether such a neighbourhood exists, then the problems turn out to be decidable in their full generality, with PSPACE complexity. Our analysis is based on the set of initial configurations such that positivity holds, which leads to new insights into these difficult problems, and interesting geometrical interpretations.

Key words and phrases: Skolem problem, verification, dynamical systems, robustness.

Extended version of the conference paper that appeared in STACS’22.

This work was partly supported by the DST/CEFIPRA/INRIA associated team EQuaVE, DST/SERB Matrices grant MTR/2018/00074, ANR-20-CE25-0012 MAVeriQ, and DFG grant 389792660 as part of TRR 248 (see <https://perspicuous-computing.science>).

Our techniques also allow us to tackle robustness for ultimate positivity, which asks whether there is a bound on the number of steps after which the LRS remains positive. There are two natural robust variants depending on whether we ask for a “uniform” bound on this number of steps, independent of the starting configuration in the neighbourhood. We show that for the uniform variant, results are similar to positivity. On the other hand, for the non-uniform variant, robust ultimate positivity has different properties when the neighbourhood is open and when it is closed. When it is open, the problem turns out to be tractable, even when the neighbourhood is given as part of the input.

1. INTRODUCTION

A rational linear recurrence relation (LRR) of order κ is a relation $u_{n+\kappa} = \sum_{j=0}^{\kappa-1} a_j \cdot u_{n+j}$ defined by a tuple of coefficients $(a_0, \dots, a_{\kappa-1}) \in \mathbb{Q}^\kappa$, $a_0 \neq 0$. Given the initial configuration $(u_0, \dots, u_{\kappa-1})$, i.e. the first κ entries of the recurrence, which could be rationals or real algebraic numbers, there is a unique infinite sequence $(u_n)_{n \in \mathbb{N}}$ that satisfies the relation. This is called a Linear Recurrence Sequence (LRS). The Skolem problem asks, given an LRS, i.e., a recurrence relation and an initial configuration, whether the sequence ever hits 0, i.e. does there exist $n \in \mathbb{N}$ with $u_n = 0$. The positivity problem is a variant where the question asked is whether for all $n \in \mathbb{N}$, $u_n \geq 0$. Another variant is the ultimate positivity problem which asks whether there exists an integer $N_0 \in \mathbb{N}$, such that for all $n \geq N_0$, $u_n \geq 0$. All these problems have applications in software verification, probabilistic model checking, discrete dynamic systems, theoretical biology, economics.

While the statements seem innocuous, the decidability of all these problems remains open since their introduction in the 1930’s. Only partial decidability results are known, e.g., when the dimension is less than 5 [Ver85]. For a subclass of the so-called *simple* LRS, positivity is known to be decidable for order up to 9 [OW14a], while ultimate positivity is decidable for all orders for that class [OW14c]. On the other hand, the authors of [OW14b] prove an important hardness result: solving positivity or ultimate positivity would entail major breakthroughs in Diophantine approximations. More precisely, one would be able to approximate the (Lagrange) type of many transcendental numbers, which deals with how close one can approximate the transcendental number using rational numbers having small denominators.

This hardness result contrasts with positive results obtained for relaxations of the problems: instead of considering a fixed initial configuration, [Tiw04, Bra06] consider every possible configuration as initial, i.e., they ask if there exists an initial configuration starting from which ensures that all entries of the sequence remain positive (this is sometimes called the uninitialised positivity problem). Surprisingly they show that this problem can be decided in PTIME. More recently, this result has been extended to processes with choices [AGV18].

In this article, we consider natural variants that lie between the hard question of fixed initial configuration [OW14b], and the easy question when the initial configuration is completely unconstrained [Tiw04, Bra06]. Our goal is to undertake a comprehensive study of what happens when starting from a neighbourhood (ball) around the initial configuration. An immediate question that arises is whether the neighbourhood is part of the input or not and it turns out that this has a significant impact on decidability. Hence, we consider two sub-variants, first by fixing the neighbourhood and the second by asking if there exists a neighbourhood around the initial configuration (the existential variant). In both these cases, starting from any initial configuration in this neighbourhood we ask if

	Exact	Robust	\exists Robust
Skolem	? (NP hard [BP02])	Diophantine hard	PSPACE*
Positivity	Diophantine hard [OW14b]	Diophantine hard	PSPACE*
Ult. Pos	Lagrange hard [OW14b]	Unif Lagrange hard	PSPACE
		Non-unif Lagrange hard (closed balls)	PSPACE
		PSPACE (open balls)	PSPACE

Table 1: Complexity results for different variants of Skolem/Positivity/Uniform Positivity. The most surprising results of our paper are in **bold**. \exists robustness is tractable in all cases. On the other hand, Robust Skolem is Diophantine hard while it is not the case for usual Skolem. Further, non-uniform robust ultimate positivity is tractable for open balls while every other variant of robust (ultimate) positivity shows some form of number theoretic hardness. The asterisks (*) denote that these results hold for LRS of *a priori* bounded order.

- all entries of the recurrence sequence remain positive. We call this the *robust positivity problem*.
- all entries of the recurrence sequence remain away from zero. We call this the *robust Skolem problem*.
- all the entries of the recurrence sequence eventually (i.e., after a certain number of steps) become and remain positive. We call this the *robust ultimate positivity problem*. In this last case, it is also natural to consider a *uniform variant*, namely whether there is a uniform bound N_0 on the number of steps, such that all starting configurations within the neighbourhood are positive after N_0 .

Our motivation to look at these problems stems from their role in capturing a powerful and natural notion of robustness, where the exact initial configuration cannot be fixed with arbitrarily high precision (which is often the case with real systems).

We start by observing that as we need to tackle multiple initial configurations, we reason about the set of initial configurations from which positivity holds, which is sufficient to answer robustness questions. For this, we revisit the usual algebraic equations in a more graphical manner, which forms the crux of our approach. This allows us to reinterpret and generalise the hardness result of [OW14a], giving our first main contribution: if the neighbourhood is given as a fixed ball with real algebraic centre, then both robust Skolem and robust positivity are Diophantine hard, while robust ultimate positivity is Lagrange hard (these notions are defined formally in Section 2) in all cases except in the non-uniform case when the given ball is open. In this last case, it turns out that the problem can be solved in PSPACE. Note in particular that the Diophantine hardness of robust Skolem is somewhat surprising, since it has recently been shown that the Skolem problem itself, at least in the case of simple LRS, can be solved assuming the Skolem conjecture and the p-adic Schanuel conjecture [BLN⁺22], and hence is perhaps not expected to be Diophantine hard.

We then turn to the problems where the ball is not fixed, and ask if there exists a radius $\psi > 0$ such that 0 or negative numbers can be avoided from every initial configuration in the ψ ball around a given initial configuration. Our main contribution here is to show that the robust variants of the Skolem, positivity and ultimate positivity problems are all decidable in full generality, with PSPACE complexities. We summarise our results in Table 1, with the precise statements in Theorems 3.2, 3.3 in Section 3.

Related work. As mentioned earlier, the Skolem problem and its variants have received a lot of attention. Given the hardness of these problems, ε -approximate solutions have been considered, e.g., in [BHM14, AAGT15] with different definitions of approximations. In comparison with our work, these are designed towards allowing approximate model checking. More recently, the notion of imprecision in Skolem and related problems was considered in [BFJ⁺20, DKM⁺21]. In [BFJ⁺20], the authors consider rounding functions at every step of the trajectory. In [DKM⁺21], the so called Pseudo-Skolem problem is defined, where imprecisions up to ε are allowed at every step of the trajectory, which is shown to be decidable in PTIME. These are quite different from our notion of robustness, which faithfully considers the trajectories generated from a ball representing ε -perturbations around the initial configuration. In a very recent extension [DKM⁺22] of [DKM⁺21], it is shown that the existential robust Skolem question and a special case of the Pseudo-Skolem problem can be solved using o-minimality of the theory of reals with exponentiation. Lastly, [Neu21] considers the problem when specified in a model of computation that takes real numbers, as opposed to rational numbers, as input. In this setting, perturbations are permitted in both the initialisation and the recurrence itself. This approach sidesteps several of the number-theoretic challenges, and decides problems on Linear Recurrence Sequences on all inputs except a set of measure 0.

The novelty of this paper is three-fold: first, we provide the first comprehensive study of robustness focussed with respect to the initial configuration, covering several possible cases and variants; second, we provide and critically use geometric insight combined with the underlying number theory to prove our results; third, we show hardness results, in particular Diophantine hardness for a variant of the Skolem problem. This paper is an extended version of the work [ABGV22] that featured in the proceedings of STACS 2022. The additional content can be summarised as follows:

- (1) We formalise the distinction between the notions of uniform and non-uniform Robust Ultimate Positivity.
- (2) We prove decidability results for non-uniform Robust Ultimate Positivity.
- (3) We leverage the underlying geometry to derive equivalences between robust Positivity, Uniform Ultimate Positivity and Skolem, in the cases of open and closed balls.
- (4) We prove Diophantine hardness of the robust problems at higher orders.

Structure of the paper. The structure of this paper is as follows: In Section 2 we define preliminaries, in particular the Skolem and (ultimate) positivity problem as well as known number-theoretic hardness results. In Section 3 we define the problem of our interest, namely robustness with respect to the initial configuration of Skolem as well as (ultimate) positivity and state all our main results. In Section 4 we provide a geometric interpretation of the behaviour of linear recurrences, which allows us to better understand and characterise the number-theoretic hardness results. In Section 5, we build upon the geometric insights to prove the hardness results claimed for robustness. In Section 6, we again use and generalise the geometric interpretation in Section 4 to prove our positive results stated in Section 3 both in terms of decidability and complexity upper bounds. Finally we end with a conclusion in Section 7.

2. PRELIMINARIES

Let κ be any non-negative integer. We let \mathbb{Q}, \mathbb{R} denote the set of rationals and reals, respectively. Further, $\mathbb{Q}^\kappa, \mathbb{R}^\kappa$ denote κ -dimensional vectors over rationals, reals, respectively. Let \mathbf{c}, \mathbf{d} be two vectors of \mathbb{R}^κ that can be seen as one dimensional matrices of $\mathbb{R}^{\kappa \times 1}$. The distance between \mathbf{c}, \mathbf{d} is defined as $\|\mathbf{c} - \mathbf{d}\| = \sqrt{(\mathbf{c} - \mathbf{d})^T (\mathbf{c} - \mathbf{d})}$, the standard ℓ^2 -distance. In this paper, we will consider two ways of “measuring” vectors: the first is the standard ℓ^2 -norm $\|\mathbf{c}\|$ for Euclidean length. The second is $\text{size}(\mathbf{c})$, denoting the size of its bit representation i.e., number of bits needed to write down \mathbf{c} (for complexity). We use the same notation for scalar constants with $\text{size}(a)$ denoting the number of bits to represent an algebraic/rational constant a . An algebraic number α is a root of a polynomial p with integer coefficients. It can be represented [Mig82] by a 4-tuple $(p, a, b, r) \in \mathbb{Z}[X] \times \mathbb{Q}^3$ as the only root of p at distance $< r$ from $a + ib$ (also see the Appendix). We define $\text{size}(\alpha)$ as the size of the bit representation of (p, a, b, r) .

Given $\mathbf{c} \in \mathbb{R}^\kappa$, an *open (resp. closed) ball* with centre \mathbf{c} and radius $\psi > 0$, denoted \mathcal{B} , refers to the set of all vectors $\mathbf{c}' \in \mathbb{R}^\kappa$ such that $\|\mathbf{c}' - \mathbf{c}\| < \psi$ (resp. $\|\mathbf{c}' - \mathbf{c}\| \leq \psi$). For convenience we sometimes just say ball to mean an open or closed ball.

2.1. Linear Recurrence Sequences. We start by defining linear recurrence relations and sequences.

Definition 2.1. A linear recurrence relation, *LRR* for short, $(u_n)_{n \in \mathbb{N}}$ of order κ is specified by a tuple of coefficients $\mathbf{a} = (a_0, \dots, a_{\kappa-1})$ with $a_0 \neq 0$. Given an initial configuration $\mathbf{c} = (c_0, \dots, c_{\kappa-1})$, the LRR uniquely defines a linear recurrence sequence (LRS henceforth), which is the sequence $(u_n(\mathbf{c}))_{n \in \mathbb{N}}$, inductively defined as $u_j(\mathbf{c}) = c_j$ for $j \leq \kappa - 1$, and

$$u_{n+\kappa}(\mathbf{c}) = \sum_{j=0}^{\kappa-1} a_j u_{n+j}(\mathbf{c}) \text{ for all } n \in \mathbb{N}.$$

The companion matrix associated with the LRR/LRS (it does not depend upon the initial configuration \mathbf{c}) is:

$$\mathbf{M} = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ a_0 & a_1 & a_2 & \dots & a_{\kappa-1} \end{bmatrix}.$$

The characteristic polynomial of the LRR/LRS is $X^\kappa - \sum_{j=0}^{\kappa-1} a_j X^j$. The LRS is said to be simple if every root of the characteristic polynomial has multiplicity one. The size s of the LRS is the size of its bit representation and is given by $s = \sum_{j=0}^{\kappa-1} (\text{size}(a_j) + \text{size}(c_j))$.

When the coefficients, i.e., entries of \mathbf{a} , of an LRR are rational, we call it a *rational LRR*. **In this paper, we will mostly be concerned with rational LRR, but the initial configuration may have rational or real algebraic entries.** Note that arithmetic with algebraic numbers can indeed be performed with perfect precision: see the Appendix for references and a brief explanation.

Notice that given an initial configuration \mathbf{c} , we have that $\mathbf{M}^n \mathbf{c} = (u_n(\mathbf{c}), \dots, u_{n+\kappa-1}(\mathbf{c}))$. Reasoning in the κ dimensions $(u_n, \dots, u_{n+\kappa-1})$ is a very useful technique that we will use throughout the paper as it displays the LRR as a linear transformation \mathbf{M} .

The characteristic roots of an LRR/LRS are the roots of its characteristic polynomial, and also the eigenvalues of the companion matrix. Let $\gamma_1, \dots, \gamma_r \in \mathbb{C}$ be the characteristic roots of the LRR/LRS. An eigenvalue γ_i is called *dominant* if it has maximal modulus $|\gamma_i| = \max_{j \leq r} |\gamma_j|$, and *residual* otherwise. When \mathbf{a} has rational entries, for all $j \leq r$, γ_j is algebraic and $\text{size}(\gamma_j) = s^{\mathcal{O}(1)}$. We denote by m_j the multiplicity of γ_j . We have $\sum_{j=1}^r m_j = \kappa$.

Proposition 2.2 (Exponential polynomial solution [EvdPSW03]). *Given an initial configuration \mathbf{c} , there exists a unique tuple of coefficients $(\alpha_{ij}(\mathbf{c}))_{i \leq r, j < m_r}$ such that for all n ,*

$$u_n(\mathbf{c}) = \sum_{i=1}^r \left(\sum_{j=0}^{m_i-1} \alpha_{ij}(\mathbf{c}) n^j \right) \gamma_i^n.$$

The coefficients $\alpha_{ij}(\mathbf{c})$ can be solved for from the initial state \mathbf{c} [HHHK05]. When \mathbf{c} has algebraic entries, it is implicit in the solution that for all i, j , both α_{ij} and $\frac{1}{\alpha_{ij}}$ are algebraic with values and norms upper bounded by $2^{s^{\mathcal{O}(1)}}$. A formal proof of this claim can be found in [ABM⁺20, Lemmas 4, 5, 6].

If the LRS is simple, then by definition $m_i = 1$ for all i , and $u_n = \sum_{i=1}^r \alpha_i(\mathbf{c}) \gamma_i^n$, with $\alpha_i(\mathbf{c})$ linear in \mathbf{c} , ie $\alpha_i(\lambda \mathbf{c} + \lambda' \mathbf{c}') = \lambda \alpha_i(\mathbf{c}) + \lambda' \alpha_i(\mathbf{c}')$.

Example 2.3. As an example, consider the Linear Recurrence Relation of order 6 with $\mathbf{a} = (-1, 4, -8, 10, -8, 4)$, i.e. $u_{n+6} = 4u_{n+5} - 8u_{n+4} + 10u_{n+3} - 8u_{n+2} + 4u_{n+1} - u_n$. The roots of the characteristic polynomial are $1, e^{i2\pi\theta}, e^{-i2\pi\theta}$, with $\theta = \frac{1}{3}$, each with multiplicity 2, and all dominant (they have the same modulus 1). The exponential polynomial solution is of the form $u_n(\mathbf{c}) = z(\mathbf{c})n + z'(\mathbf{c}) + (x(\mathbf{c})n + x'(\mathbf{c}))e^{i2\pi n\theta} + (y(\mathbf{c})n + y'(\mathbf{c}))e^{-i2\pi n\theta}$. As $u_n(\mathbf{c})$ is real, we must have that $x(\mathbf{c}), y(\mathbf{c})$ are conjugates, as well as $x'(\mathbf{c}), y'(\mathbf{c})$, and thus: $u_n(\mathbf{c}) = z(\mathbf{c})n + z'(\mathbf{c}) + 2(\text{Re}(x(\mathbf{c}))n + \text{Re}(x'(\mathbf{c}))) \cos(2\pi n\theta) + 2(\text{Im}(x(\mathbf{c}))n + \text{Im}(x'(\mathbf{c}))) \sin(2\pi n\theta)$.

2.2. Skolem and (ultimate) positivity problems.

Definition 2.4. Let $(u_n)_{n \in \mathbb{N}}$ be a rational LRR and $\mathbf{c} \in \mathbb{Q}^\kappa$. The Skolem problem is to determine if there exists $n \in \mathbb{N}$ such that $u_n(\mathbf{c}) = 0$. The positivity (resp. strict positivity) problem is to determine if for all $n \in \mathbb{N}$, $u_n(\mathbf{c}) \geq 0$ (resp. $u_n(\mathbf{c}) > 0$). The ultimate positivity (resp. ultimate strict positivity) problem is to determine if there exists $N \in \mathbb{N}$ such that for all $n \geq N$, $u_n(\mathbf{c}) \geq 0$ (resp. $u_n(\mathbf{c}) > 0$).

In this work, we will be more interested in the complement problem of Skolem: namely, whether $u_n(\mathbf{c}) \neq 0$ for all n . This is of course equivalent in terms of decidability, but this formulation is more meaningful in terms of robustness, where we want to robustly avoid 0.

The famous Skolem-Mahler-Lech theorem states that when \mathbf{c} is algebraic, the set $\{i \mid u_i(\mathbf{c}) = 0\}$ is the union of a finite set F and finitely many arithmetic progressions [Sko34, Mah35, BM76]. These arithmetic progressions can be computed but the hard part lies in deciding if the set F is empty: although we know that there is N such that for all $n > N$, $n \notin F$, we do not have an effective bound on this N in general. The Skolem problem has been shown to be decidable for LRS of order up to 4 [MST84, Ver85] and is still open for LRS of higher order. Also, only an NP hardness bound is known if the order is unrestricted [BP02, ABV17].

For simple LRS, positivity has been shown to be decidable up to order 9 [OW14a]. In [OW14c], it is proved that positivity for simple LRS is hard for $\text{co}\exists\mathbb{R}$, the class of problems whose complements are solvable in the existential theory of the reals. A last result, from [OW14b], shows the difficulty of positivity, linking it to Diophantine approximations: how close one can approximate a transcendental number with a rational number with small denominator. We will follow the reasoning from [OW14b]. We start with two definitions.

- The *Diophantine approximation type* of a real number x is defined as:

$$L(x) = \inf \left\{ c \in \mathbb{R} \mid \left| x - \frac{n}{m} \right| < \frac{c}{m^2}, \text{ for some } n, m \in \mathbb{Z} \right\}.$$

- The *Lagrange constant* of a real number x is defined as:

$$L_\infty(x) = \inf \left\{ c \in \mathbb{R} \mid \left| x - \frac{n}{m} \right| < \frac{c}{m^2}, \text{ for infinitely many } n, m \in \mathbb{Z} \right\}.$$

As mentioned in [OW14b], the Diophantine approximation type and Lagrange constant of most transcendental numbers are unknown. Let $\mathcal{A} = \{p + qi \in \mathbb{C} \mid p, q \in \mathbb{Q} \setminus \{0\}, p^2 + q^2 = 1\}$, i.e., the set of points on the unit circle of \mathbb{C} with rational real and imaginary parts, excluding $1, -1, i$ and $-i$. The set \mathcal{A} consists of algebraic numbers of degree 2, none of which are roots of unity [OW14b]. In particular, writing $p + qi = 2^{i2\pi\theta} = (-1)^{2\theta}$, we have that $\theta \notin \mathbb{Q}$ [OW14b]. We denote:

$$\mathcal{T} = \left\{ \theta \in (-1/2, 1/2] \mid e^{2\pi i\theta} \in \mathcal{A} \right\}.$$

As argued in [OW14b], the set \mathcal{T} is dense in $(-\frac{1}{2}, \frac{1}{2}]$, and consists solely of transcendental numbers. We assume that $\theta \in \mathcal{T}$ is specified by $p = \cos 2\pi i\theta$. In general, we don't have a method to compute $L(\theta)$ or $L_\infty(\theta)$ for $\theta \in \mathcal{T}$, or approximate them with arbitrary precision.

Definition 2.5. We say that a problem is \mathcal{T} -Diophantine hard (resp. \mathcal{T} -Lagrange hard) if its decidability entails that given any $\theta \in \mathcal{T}$ and $\varepsilon > 0$ as input, one can compute a number ℓ such that $|\ell - L(\theta)| < \varepsilon$ (resp. $|\ell - L_\infty(\theta)| < \varepsilon$).

Remarkably, in [OW14b], it is shown that (i) if one can solve the positivity problem in general, then one can also approximate $L(\theta)$ and (ii) if one can solve the ultimate positivity, then one can approximate $L_\infty(\theta)$. That is,

Theorem 2.6 [OW14b]. *Positivity for LRS of order 6 or above is \mathcal{T} -Diophantine hard and ultimate positivity for LRS of order 6 and above is \mathcal{T} -Lagrange hard.*

3. ROBUST SKOLEM, POSITIVITY AND ULTIMATE POSITIVITY

The Skolem and (Ultimate) Positivity problems, as defined in the previous section, consider a single initial configuration \mathbf{c} . In this article, we investigate the notion of robustness, that is, whether the property is true in a neighbourhood of \mathbf{c} , which is important for real systems, where setting \mathbf{c} with an arbitrary precision is not possible. We will consider two variants. The first one fixes the neighbourhood as a ball \mathcal{B} , while the second asks for the existence of a ball \mathcal{B} centred around a given initial configuration \mathbf{c} , such that for every initial configuration in \mathcal{B} , the respective condition is satisfied.

Definition 3.1 (Robustness for Skolem, Positivity, Ultimate Positivity). Let $(u_n)_{n \in \mathbb{N}}$ be the rational linear recurrence relation specified by a rational coefficient vector \mathbf{a} , and an initial algebraic configuration \mathbf{c} . Consider an algebraic ball \mathcal{B} (with algebraic entries for both the centre and the radius). We define the following problems:

- The *robust Skolem* problem is to determine if for all $\mathbf{c}' \in \mathcal{B}$ and all $n \in \mathbb{N}$, we have $u_n(\mathbf{c}') \neq 0$.
- The *robust positivity* problem is to determine if for all $\mathbf{c}' \in \mathcal{B}$ and all $n \in \mathbb{N}$, we have $u_n(\mathbf{c}') \geq 0$.
- The *robust non-uniform ultimate positivity* problem is to determine if for all $\mathbf{c}' \in \mathcal{B}$, there exists $N_{\mathbf{c}'} \in \mathbb{N}$ such that for all $n > N_{\mathbf{c}'}$, we have $u_n(\mathbf{c}') \geq 0$.
- The *robust uniform ultimate positivity* problem is to determine if there exists $N \in \mathbb{N}$ such that for all $\mathbf{c}' \in \mathcal{B}$ and all $n > N$, we have $u_n(\mathbf{c}') \geq 0$.

The \exists -robust variants of each problem asks whether there exists a ball \mathcal{B}' centred around \mathbf{c} such that the above holds over \mathcal{B}' .

Note that for all the variants of \exists -robustness, there exists an open ball of radius $\psi > 0$ for which robust Skolem (resp. positivity, uniform ultimate positivity) holds iff there exists a closed ball of radius $\psi' > 0$ (e.g. $\psi' = \frac{\psi}{2}$) for which it holds. Further, if there exists a ball with real radius, then there exist balls with algebraic and rational radii with the same centre. Thus we do not need to consider open and closed balls separately, nor do we need to explicitly mention the domain of the radius. For the other, i.e., *non* existential, variants of robustness as defined above, the case of closed and open balls can be different, and can also depend on whether the radius is rational or just real algebraic.

Our main results investigate the decidability and complexity of these problems.

Theorem 3.2. *For rational linear recurrence relations and algebraic balls:*

- (1) *for open and closed balls,*
 - (a) *the robust positivity problem is \mathcal{T} -Diophantine hard,*
 - (b) *the robust Skolem problem is \mathcal{T} -Diophantine hard,*
 - (c) *the robust uniform ultimate positivity problem is \mathcal{T} -Lagrange hard*
- (2) *for closed balls, the robust non-uniform ultimate positivity is \mathcal{T} -Lagrange hard.*

These lower bounds hold even for rational linear recurrence relations restricted to order 6 and for balls with rational radius.

We remark that our proof of these lower bounds does not hold for balls whose centres have rational entries.

Our theorem above implies that for uninitialised positivity, one really needs the initial configuration to take a value possibly anywhere in the space rather than in a fixed neighbourhood to obtain decidability via [Tiw04, Bra06]. We remark that Diophantine hardness is known for the non-robust variant of positivity [OW14b], but to the best of our knowledge, it was not known for any variant of the Skolem problem. In fact, in light of the latest results in [BLN⁺22], Diophantine hardness for (exact) Skolem seems unlikely, unless either of Skolem conjecture or p-adic Schanuel conjecture is falsified. We note that the p-adic techniques used therein rely on the input being integral, or rational, or algebraic. Our Diophantine hardness result, on the other hand, has connections to the positivity problem, and is intrinsically related to the common underlying geometry. In a nutshell, the distinction is that the robust variants of the problem implicitly reason about a continuum of initialisations, including those with transcendental coordinates, for which standard results like the Skolem-Mahler-Lech Theorem do not hold.

Surprisingly, we obtain decidability for every linear recurrence relation and every initial configuration when considering a given *open* ball for non-uniform robust ultimate positivity,

or by relaxing the neighbourhood to be as small as desired for any of the variants. This constitutes our second main result:

Theorem 3.3. *The following decidability results hold for rational linear recurrence relations:*

- (1) \exists -robust Skolem, \exists -robust positivity and \exists -robust (non)-uniform ultimate positivity are decidable for a centre \mathbf{c} with algebraic entries. Further:
 - (a) Deciding \exists -robust (non)-uniform ultimate positivity can be done in PSPACE.
 - (b) When the centre \mathbf{c} has rational entries, for any $K \in \mathbb{N}$, deciding \exists -robust Skolem and \exists -robust positivity for LRS of order at most K can be done in PSPACE.
- (2) Robust non-uniform ultimate positivity is decidable in PSPACE for open algebraic balls.

The main difference between our techniques and several past works (except [AGKV16] which is restricted to eigenvalues being roots of unity) is as follows: given an LRR $(u_n)_{n \in \mathbb{N}}$, our intuition and proofs hinge on representing the set P of initial configurations \mathbf{d} from which positivity holds. Formally:

$$P = \{\mathbf{d} \in \mathbb{R}^\kappa \mid u_n(\mathbf{d}) \geq 0 \text{ for all } n \in \mathbb{N}\}.$$

We may note that the set P is convex. To see this, observe that for $d, d' \in P$, for all $\alpha, \beta > 0$ with $\alpha + \beta = 1$, we have $\alpha \mathbf{d} + \beta \mathbf{d}' \in P$ as $u_n(\alpha \mathbf{d} + \beta \mathbf{d}') = \alpha u_n(\mathbf{d}) + \beta u_n(\mathbf{d}') \geq 0$ for all n . We also remark that a definition similar to P is possible for the set S of initial configurations from which 0 is avoided. But it turns out that that set is much harder to represent (e.g., it is not convex in general). Using P surprisingly suffices to deal with robust Skolem as well.


In Section 4, we provide the geometric intuitions behind our ideas as well as set up the notations for the proofs of the above theorems. We exploit the geometric intuitions from Section 4 in Section 5, to prove Theorem 3.2. In Section 6 we prove Theorem 3.3 providing algorithms for the decidable cases.

4. GEOMETRICAL REPRESENTATION OF AN LRR FOR DIOPHANTINE HARDNESS

The number-theoretic hardness of the non-robust variants starts at order 6; in this work, we show corresponding hardness results for the robust variants too. Decidability at lower orders is non-trivial: see [Vah23] for an exposition. Hence, in this section and the next, we will focus on a particular LRR of order $\kappa = 6$, sufficient for the proofs of hardness, i.e. Theorem 3.2. In Section 6, we will generalise some of the constructions explored here to obtain our decidability results stated in Theorem 3.3.

Let $\theta \in \mathcal{T}$, i.e. $e^{i2\pi\theta} = p + qi \in \mathcal{A}$, with both p, q rational and $p^2 + q^2 = 1$. We want to approximate $L(\theta)$ (indeed this is the problem that is ‘‘Diophantine hard’’). For Lagrange hardness, we will adapt the construction and proof in section 5.4, approximating $L_\infty(\theta)$ instead of $L(\theta)$.

Consider the Linear Recurrence Relation of order 6 defined by $\mathbf{a} = (-1, 4p + 2, -(4p^2 + 8p + 3), 8p^2 + 8p + 4, -(4p^2 + 8p + 2), 4p + 2)$. The roots of the characteristic polynomial are $1, e^{i2\pi\theta}, e^{-i2\pi\theta}$, each with multiplicity 2, and all dominant (they have the same modulus 1). Example 3 is a particular case of this \mathbf{a} , with $p = \frac{1}{2} = \cos(\frac{\pi}{3})$. However, notice that $\theta = \frac{1}{3} \notin \mathcal{T}$ as it corresponds to $q = \sin(\frac{\pi}{3}) = \frac{\sqrt{3}}{2} \notin \mathbb{Q}$. Now, since $u_n(\mathbf{c})$ is a real number for any n and real initial configuration \mathbf{c} , we can write the exponential polynomial solution in

Figure 1: Visual representation of the cone $P_{(0,0,0)}$.

the form:

$$u_n(\mathbf{c}) = z_{dom}(\mathbf{c})n - x_{dom}(\mathbf{c})n \cos(2\pi n\theta) - y_{dom}(\mathbf{c})n \sin(2\pi n\theta) \\ + z_{res}(\mathbf{c}) - x_{res}(\mathbf{c}) \cos(2\pi n\theta) - y_{res}(\mathbf{c}) \sin(2\pi n\theta)$$

The coefficients $z_{dom}(\mathbf{c})$, $x_{dom}(\mathbf{c})$, $y_{dom}(\mathbf{c})$ and $z_{res}(\mathbf{c})$, $x_{res}(\mathbf{c})$, $y_{res}(\mathbf{c})$ are associated with the initial configuration \mathbf{c} of the LRS. In the following, we reason in the basis of vectors \vec{z}_{dom} , \vec{x}_{dom} , \vec{y}_{dom} , \vec{z}_{res} , \vec{x}_{res} , \vec{y}_{res} , as the geometrical interpretation is simpler in this basis. We will eventually get back to the original coordinate vector basis at the end of the process. From e.g., [HHHK05, Section 2], we know that we can transform from one basis to the other using an invertible Matrix C with $C \cdot \mathbf{c} = (z_{dom}(\mathbf{c}), x_{dom}(\mathbf{c}), y_{dom}(\mathbf{c}), z_{res}(\mathbf{c}), x_{res}(\mathbf{c}), y_{res}(\mathbf{c}))$.

We study the positivity of u_n by studying the positivity of $v_n = \frac{u_n}{n}$, for all $n \geq 1$. We denote $v_n^{dom}(z_{dom}, x_{dom}, y_{dom}) = z_{dom} - x_{dom} \cos(2\pi n\theta) - y_{dom} \sin(2\pi n\theta)$, which we call the dominant part of v_n , while we denote $v_n^{res}(z_{res}, x_{res}, y_{res}) = \frac{1}{n}(z_{res} - x_{res} \cos(2\pi n\theta) - y_{res} \sin(2\pi n\theta))$, which we call the residual part of v_n . The residual part tends towards 0 when n tends towards infinity because of the coefficient $\frac{1}{n}$.

4.1. High-Level intuition and Geometrical Interpretation. We provide a geometrical interpretation of set P . We cannot characterise it exactly, even in this particular LRR of order $\kappa = 6$ (else we could decide positivity for this case which is known to be Diophantine hard). To describe P , we define its “section” over $(z_{dom}, x_{dom}, y_{dom})$ given $(z_{res}, x_{res}, y_{res})$:

$$P_{(z_{res}, x_{res}, y_{res})} = \{(z_{dom}, x_{dom}, y_{dom}) \mid v_n(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \geq 0 \text{ for all } n\}.$$

It suffices to characterise $P_{(z_{res}, x_{res}, y_{res})}$ for all $(z_{res}, x_{res}, y_{res})$ in order to characterise P , as $P = \{(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \mid (z_{dom}, x_{dom}, y_{dom}) \in P_{(z_{res}, x_{res}, y_{res})}\}$. Among these sets, one is particularly interesting: $P_{(0,0,0)}$, as it is the set of tuples $(z_{dom}, x_{dom}, y_{dom})$ such that $v_n^{dom}(z_{dom}, x_{dom}, y_{dom}) \geq 0$ for all $n \in \mathbb{N}$. Our reason for focussing on this representation of P is three-fold. First, unlike P , the set $P_{(0,0,0)}$ can be characterised exactly, as a cone depicted in Figure 1 (this will be formally shown in Lemma 4.1 below). Second, the set $P_{(z_{res}, x_{res}, y_{res})}$ is in 3 dimensions that we can represent more intuitively than a 6 dimensional set. Last but not least, we can show that $P_{(z_{res}, x_{res}, y_{res})} \subseteq P_{(0,0,0)}$ for all $(z_{res}, x_{res}, y_{res})$ (Lemma 4.3).

On the other hand, we also consider a related set in 6 dimensions:

$$P_{dom} = \{(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \mid \forall n, v_n^{dom}(z_{dom}, x_{dom}, y_{dom}) \geq 0\}. \quad (4.1)$$

We note that $P_{(0,0,0)}$ is the projection of P_{dom} over the 3 dimensions $(z_{dom}, x_{dom}, y_{dom})$. Also, characterising $P_{(0,0,0)}$ is sufficient to characterise P_{dom} as $(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \in$




Figure 2: Sections of $P_{(0,0,0)}$ (in black) and $P_{(z_{res}, x_{res}, y_{res})}$ (in dashed red), carved out by hyperplanes (H_i) (in black) and ($H_i(z_{res}, x_{res}, y_{res})$) (in red) respectively.

P_{dom} iff $(z_{dom}, x_{dom}, y_{dom}) \in P_{(0,0,0)}$. As $P_{(z_{res}, x_{res}, y_{res})} \subseteq P_{(0,0,0)}$ for all $(z_{res}, x_{res}, y_{res})$, we have $P \subseteq P_{dom}$.

We are now ready to represent $P_{(z_{res}, x_{res}, y_{res})}$ given some value $(z_{res}, x_{res}, y_{res})$. We can interpret $P_{(z_{res}, x_{res}, y_{res})}$ in terms of half spaces: $P_{(z_{res}, x_{res}, y_{res})} = \bigcap_{m=1}^{\infty} H_m^+(z_{res}, x_{res}, y_{res})$, with $H_m^+(z_{res}, x_{res}, y_{res}) = \{(z_{dom}, x_{dom}, y_{dom}) \mid v_m(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \geq 0\}$. The half space $H_m^+(z_{res}, x_{res}, y_{res})$ is delimited by the hyperplane

$$H_m(z_{res}, x_{res}, y_{res}) = \{(z_{dom}, x_{dom}, y_{dom}) \mid v_m(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) = 0\}.$$

which is a vector space ($\cos(2\pi m\theta)$ and $\sin(2\pi m\theta)$ are constant when m is fixed).

Consider the case of $(z_{res}, x_{res}, y_{res}) = (0, 0, 0)$. We denote $H_m^+ = H_m^+(0, 0, 0)$ and $H_m = H_m(0, 0, 0)$ for all m . For instance, $H_0 = \{(z_{dom}, x_{dom}, y_{dom}) \mid z_{dom} = x_{dom}\}$, as $v_0^{dom}(z_{dom}, x_{dom}, y_{dom}) = z_{dom} - x_{dom}$.

Define \mathbf{M}_{dom} as the matrix that captures the action of the LRS $(v_n^{dom})_{n \in \mathbb{N}}$ in the subspace of dominant coefficients of the exponential polynomial solution space. We have $H_m = \mathbf{M}_{dom} H_{m-1} = \mathbf{M}_{dom}^m H_0$. We characterise \mathbf{M}_{dom} in Lemma 4.2 as a rotation around $\overrightarrow{z_{dom}}$ of angle $-2\pi\theta$, which allows to characterise H_m as the hyperplane which is the rotation of H_0 of angle $2m\pi\theta$ around $\overrightarrow{z_{dom}}$. That is, the cone shape for $P_{(0,0,0)}$ is obtained by cutting away chunk of the 3D space delimited by hyperplanes (H_m), the rotation $2n\pi\theta$ being dense in $[-\pi, \pi]$.

Coming back to some value $(z_{res}, x_{res}, y_{res}) \neq (0, 0, 0)$, we have that the hyperplane $H_n(z_{res}, x_{res}, y_{res})$ is parallel to the hyperplane H_n (which is tangent to the cone $P_{(0,0,0)}$), because for H_n of the form $uz_{dom} + vx_{dom} + wy_{dom} = 0$, we have $H_n(z_{res}, x_{res}, y_{res})$ is defined by $\{(z_{dom}, x_{dom}, y_{dom}) \mid uz_{dom} + vx_{dom} + wy_{dom} = C\}$, for $C = \frac{z_{res} + x_{res} \cos(2\pi n\theta) + y_{res} \sin(2\pi n\theta)}{n}$ a constant as n is fixed.

Thus, with this idea in mind, we can visualise $P_{(z_{res}, x_{res}, y_{res})}$ as depicted in Figure 2, using $P_{(0,0,0)}$ and the hyperplanes $H_n(z_{res}, x_{res}, y_{res})$ parallel to H_n , with an explicit bound

on the distance from $H_n(z_{res}, x_{res}, y_{res})$ to H_n , which further tends towards 0 as n tends towards infinity. Next, we formalise the above intuition/picture into lemmas.

4.2. Characterisation of $P_{(0,0,0)}$ and representing $P_{(z_{res}, x_{res}, y_{res})}$. We now formalise some of the ideas in the above subsection. First, we start with Lemma 4.1 which shows that $P_{(0,0,0)}$ describes a cone, as displayed on Figure 1.

Lemma 4.1. $P_{(0,0,0)} = \{(z_{dom}, x_{dom}, y_{dom}) \mid z_{dom} \geq \sqrt{x_{dom}^2 + y_{dom}^2}\}$.

Proof. We have $\cos(2\pi n\theta)^2 + \sin(2\pi n\theta)^2 = 1$ and $\cos(2\pi n\theta)$ is dense in $[-1, 1]$ as $\theta \notin \mathbb{Q}$. Denote $X = \cos(2\pi n\theta)$, and study the function $f(X) = x_{dom}X + y_{dom}\sqrt{1 - X^2}$. Its derivative is $f'(X) = x_{dom} - \frac{y_{dom}X}{\sqrt{1-X}\sqrt{1+X}}$. We have $f'(X) = 0$ iff $X = X_0 = \frac{x_{dom}}{\sqrt{x_{dom}^2 + y_{dom}^2}}$. This gives us that $\max f(X) = f(X_0) = \frac{x_{dom}^2 + y_{dom}^2}{\sqrt{x_{dom}^2 + y_{dom}^2}} = \sqrt{x_{dom}^2 + y_{dom}^2}$. Thus, for all $(z_{dom}, x_{dom}, y_{dom})$ with $z_{dom} \geq \sqrt{x_{dom}^2 + y_{dom}^2}$, we have $z_{dom} \geq \max(f(X))$ and $v_n(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \geq z_{dom} - f(X) \geq 0$ for all n . On the other hand, if $z_{dom} < \sqrt{x_{dom}^2 + y_{dom}^2}$, then there exists n such that $f(\cos(2\pi n\theta))$ is arbitrarily close to $\max f(X) > z_{dom}$, and in particular $v_n = z_{dom} - f(\cos(2\pi n\theta)) < 0$. \square

We show now that the linear function \mathbf{M}_{dom} associated with the LRR $(v_n^{dom})_{n \in \mathbb{N}}$ is actually a rotation of angle $-2\pi\theta$.

Lemma 4.2. $\mathbf{M}_{dom}(z_{dom}, x_{dom}, y_{dom}) = (z_{dom}, x_{dom} \cos(2\pi\theta) + y_{dom} \sin(2\pi\theta), y_{dom} \cos(2\pi\theta) - x_{dom} \sin(2\pi\theta))$, that is \mathbf{M}_{dom} is a rotation around axis \vec{z} of angle $-2\pi\theta$.

Proof. We use the formulas $\cos(a + b) = \cos(a)\cos(b) - \sin(a)\sin(b)$ and $\sin(a + b) = \sin(a)\cos(b) + \cos(a)\sin(b)$.

Matrix \mathbf{M}_{dom} transforms $v_n^{dom}(z_{dom}, x_{dom}, y_{dom})$ into $v_{n+1}^{dom}(z_{dom}, x_{dom}, y_{dom})$. Using the formulas above with $a = 2\pi n\theta, b = 2\pi\theta$, we have that for all $n \geq 1$, $v_{n+1}^{dom}(z_{dom}, x_{dom}, y_{dom}) = v_n^{dom}(z_{dom}, x_{dom} \cos(2\pi\theta) + y_{dom} \sin(2\pi\theta), y_{dom} \cos(2\pi\theta) - x_{dom} \sin(2\pi\theta))$ for all n , and thus \mathbf{M}_{dom} transforms $(z_{dom}, x_{dom}, y_{dom})$ into $(z_{dom}, x_{dom} \cos(2\pi\theta) + y_{dom} \sin(2\pi\theta), y_{dom} \cos(2\pi\theta) - x_{dom} \sin(2\pi\theta))$.

To see why the transformed coordinates are obtained by the rotation claimed above, observe the action in the plane of rotation. Indeed, consider a point p in 2D space at Cartesian coordinates (x_{dom}, y_{dom}) . Its polar coordinates are (r, α) , with $r = \sqrt{x_{dom}^2 + y_{dom}^2}$ being the distance between $(0, 0)$ and p . Consider the point at polar coordinates $(r, \alpha - 2\pi\theta)$, i.e. the point obtained by rotating p through an angle of $-2\pi\theta$. Through a straightforward application of trigonometric identities, we can see that it is at Cartesian coordinates $(r \cos(\alpha - 2\pi\theta), r \sin(\alpha - 2\pi\theta)) = (r \cos(\alpha) \cos(2\pi\theta) + r \sin(\alpha) \sin(2\pi\theta), r \sin(\alpha) \cos(2\pi\theta) - r \cos(\alpha) \sin(2\pi\theta)) = (x_{dom} \cos(2\pi\theta) + y_{dom} \sin(2\pi\theta), y_{dom} \cos(2\pi\theta) - x_{dom} \sin(2\pi\theta))$. On comparing with the form obtained in the preceding paragraph, we conclude that the rotation of angle $-2\pi\theta$ transforms (x_{dom}, y_{dom}) into $(x_{dom} \cos(2\pi\theta) + y_{dom} \sin(2\pi\theta), y_{dom} \cos(2\pi\theta) - x_{dom} \sin(2\pi\theta))$. \square

Finally, the following lemma implies that $P \subseteq P_{dom}$.

Lemma 4.3. For all $z_{res}, x_{res}, y_{res}$, we have $P_{(z_{res}, x_{res}, y_{res})} \subseteq P_{(0,0,0)}$.

Proof. We use the following simple but important observation. Let $(u_n)_{n \in \mathbb{N}}$ be an LRS where all roots have modulus 1, i.e., each root is of the form $\gamma = e^{i\theta}$, with distinct values of θ . Let u_j be the j^{th} element of the LRS, with $j \in \mathbb{N}$. Then for all ε, N , there exists $n > N$ with $|u_n - u_j| < \varepsilon$. That is, for each value visited, the LRS will visit arbitrarily close values an infinite number of times. This is the case in particular of v_n^{dom} .

Now, assume for contradiction that there is a configuration $(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}})$ in $P_{(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})} \setminus P_{(0,0,0)}$. Since $(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}) \notin P_{(0,0,0)}$, there exists m with $v_m^{\text{dom}}(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}) < 0$. We let $\varepsilon = \frac{|v_m^{\text{dom}}(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}})|}{3}$ and N such that for all $n > N$, $|v_n^{\text{res}}| < \varepsilon$ (because it converges towards 0 when n tends towards infinity). From the above observation, we obtain an $n > N$ such that $|v_n^{\text{dom}}(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}) - v_m^{\text{dom}}(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}})| < \varepsilon$. Thus:

$$\begin{aligned} v_n(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}, z_{\text{res}}, x_{\text{res}}, y_{\text{res}}) &= v_n^{\text{dom}}(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}) + v_n^{\text{res}}(z_{\text{res}}, x_{\text{res}}, y_{\text{res}}) \\ &< v_m^{\text{dom}}(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}) + \varepsilon + \varepsilon < 0. \end{aligned}$$

A contradiction with $(z_{\text{res}}, x_{\text{dom}}, y_{\text{dom}}) \in P_{(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})}$. □

5. PROOF OF THEOREM 3.2

5.1. Intuition for hardness of (robust) positivity. Consider a vector \mathbf{d} on the surface of P_{dom} , namely $\mathbf{d} = (z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}, z_{\text{res}}, x_{\text{res}}, y_{\text{res}})$, that is, $(z_{\text{dom}}, x_{\text{dom}}, y_{\text{dom}}) \in P_{(0,0,0)}$. Consider the subset of $P_{(0,0,0)}$ which consists of points whose first coordinate z_{dom} is the same as that of \mathbf{d} . For all n , let \mathbf{e}_n be the point of this section where hyperplane H_n is tangent to $P_{(0,0,0)}$. Let τ be the angle made between the centre b of the section, \mathbf{e}_0 and \mathbf{d} . Hence, \mathbf{e}_0 is at angle 0 and \mathbf{e}_n at angle $2\pi n\theta \pmod{2\pi}$. We depict this pictorially in Figure 3.

We have that $u_n(\mathbf{d}) \geq 0$ for all n iff \mathbf{d} is in the intersection of all half spaces defined by $H_i(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})$. As $2\pi n\theta \pmod{2\pi}$ is dense in $[0, 2\pi)$, for all $\beta > 0$, there is a n such that \mathbf{e}_n is at angle $\alpha_n \in [\tau - \beta, \tau + \beta]$, hence H_n will be ε -close to \mathbf{d} . To know whether \mathbf{d} is in the half space defined by $H_n(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})$, we need to compare the distance ε between H_n and \mathbf{d} , with the value of n . If the value of n is too large, then the distance between $H_n(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})$ and H_n is smaller than ε , and \mathbf{d} is in the half space $H_n^+(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})$.

In other words, for $(u_n(\mathbf{d}))_{n \in \mathbb{N}}$ not to be positive, n needs to be both small enough and such that $2\pi n\theta \pmod{2\pi}$ is close to τ . This is similar to $L(\theta)$ being small, as shown in Lemma 5.2.

Now, for robust positivity (Theorem 3.2.1.a), we consider a ball \mathcal{B} entirely in P_{dom} , tangent to the surface of P_{dom} only on point \mathbf{d} . The ball will be positive iff the curvature of the ball is steeper than the curvature from hyperplanes $H_n(z_{\text{res}}, x_{\text{res}}, y_{\text{res}})_{n \in \mathbb{N}}$ around \mathbf{d} , as shown in Lemma 5.3. This will correspond again to computing $L(\theta)$, thus showing hardness.

5.2. Formalising the proof for closed balls and robust positivity. In this section, we formalise the intuition given above, in the case of a closed ball and for robust positivity, i.e., Theorem 3.2.1.a. We will extend this to the other cases of Theorem 3.2 in the next subsection.

We recall the following:




Figure 3: Representation of a section of $P_{(0,0,0)}$, with hyperplanes H_0, H_{10} being represented.

- The *Diophantine approximation type* of a real number x is defined as:

$$L(x) = \inf \left\{ c \in \mathbb{R} \mid \left| x - \frac{n}{m} \right| < \frac{c}{m^2}, \text{ for some } n, m \in \mathbb{Z} \right\}.$$

- The *Lagrange constant* of a real number x is defined as:

$$L_\infty(x) = \inf \left\{ c \in \mathbb{R} \mid \left| x - \frac{n}{m} \right| < \frac{c}{m^2}, \text{ for infinitely many } n, m \in \mathbb{Z} \right\}.$$

Now, for all $x \in \mathbb{R}$, we denote by $[x]$ the quantity $\min_{n \in \mathbb{Z}} |x - 2\pi n|$. Notice that $[x] \in [0, \pi]$; we follow the equivalent convention of [LS97] to state that $L(\theta) = \frac{1}{2\pi} \inf_{n > 0} n[2\pi n\theta]$, and that $L_\infty(\theta) = \frac{1}{2\pi} \liminf_{n > 0} n[2\pi n\theta]$.

We show how to ε -approximate $L(\theta)$ in the following, using an oracle for robust positivity, following ideas in [OW14b]. To compute some ℓ that is ε -close to $L(\theta)$ for a given $\varepsilon > 0$, we perform a binary search on ℓ . By definition, it is clear that $L_\infty(\theta) \geq L(\theta) \geq 0$. An old observation of Dirichlet shows that every real number has Lagrange constant at most 1. Due to the work of Khintchine [Khi35], it is further known that these constants lie between 0 and $1/\sqrt{5}$. So, for the binary search, we start with a lower bound $\ell_{min} = 0$ and an upper bound $\ell_{max} = 1/\sqrt{5}$. For $\ell := \frac{\ell_{min} + \ell_{max}}{2}$, we want to know if $\ell \geq L(\theta) - \varepsilon$ (and then we set $\ell_{min} := \ell$) or whether $\ell \leq L(\theta) + \varepsilon$ (and then we set $\ell_{max} := \ell$).

Let $\varepsilon > 0$ and ℓ be a guess to check against $L(\theta)$. We define a *closed ball* \mathcal{B}_ψ of radius $\sqrt{2}\psi$, centred at $\mathbf{c} = (2 + \psi, 2 - \psi, 0, 0, 2\pi\ell)$, with $\psi < \frac{1}{3}$ and $\psi < \pi\ell$. We now show that the ball \mathcal{B}_ψ is entirely in P_{dom} , only touching the surface of P_{dom} at the point $\mathbf{d} = (2, 2, 0, 0, 2\pi\ell)$:

Lemma 5.1. *We have that \mathbf{d} is on the surfaces of both B_ψ and P_{dom} .*

Further, every $\mathbf{d}' = (z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \in \mathcal{B}_\psi \setminus \{\mathbf{d}\}$ satisfies

- $x_{dom} < z_{dom}$, and
- \mathbf{d}' is in P_{dom} but not at the surface of P_{dom} .

Proof. First, $\mathbf{d} = (2, 2, 0, 0, 0, 2\pi\ell)$ satisfies $\|\mathbf{c} - \mathbf{d}\| = \sqrt{2}\psi$, so $\mathbf{d} \in B_\psi$, on the surface of B_ψ . Further, \mathbf{d} satisfies the equation of Lemma 4.1: $2^2 = 4 = (2+0)^2$, so $\mathbf{d} \in P_{dom}$ at its surface.

Let $\mathbf{d}' = (z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res}) \in \mathcal{B}_\psi \setminus \{\mathbf{d}\}$.

Thus, $(z_{dom} - (2 + \psi))^2 + (x_{dom} - (2 - \psi))^2 + y_{dom}^2 + z_{res}^2 + x_{res}^2 + (y_{res} - 2\pi\ell)^2 \leq 2\psi^2$. In particular, we have $(z_{dom} - (2 + \psi))^2 + (x_{dom} - (2 - \psi))^2 + y_{dom}^2 \leq 2\psi^2$, i.e. $(z_{dom} - 2)^2 + (x_{dom} - 2)^2 + y_{dom}^2 \leq 2\psi(z_{dom} - x_{dom})$. That is, $x_{dom} \leq z_{dom}$ as $(z_{dom} - 2)^2 + (x_{dom} - 2)^2 + y_{dom}^2 \geq 0$. Further, equality $z_{dom} = x_{dom}$ would imply that $z_{dom} = x_{dom} = 2$ and $y_{dom} = 0$; Since the ball is centred at $(2 + \psi, 2 - \psi, 0, 0, 0, 2\pi\ell)$ and has radius $\sqrt{2}\psi$, this implies that $\mathbf{d}' = \mathbf{d}$, which is not possible. Hence $x_{dom} < z_{dom}$.

For the last result, by Lemma 4.1, it suffices to prove that $x_{dom}^2 + y_{dom}^2 < z_{dom}^2$. Assume by contradiction that $z_{dom}^2 - x_{dom}^2 - y_{dom}^2 \leq 0$. As $\mathbf{d}' \in \mathcal{B}_\psi$, we have in particular $(z_{dom} - 2 - \psi)^2 + (x_{dom} - 2 + \psi)^2 + y_{dom}^2 \leq 2\psi^2$. Adding the two inequalities, we obtain:

$$(z_{dom} - 2)^2 + z_{dom}^2 - 2\psi(z_{dom} - 2) + 2x_{dom}(\psi - 2) + 4 - 4\psi \leq 0.$$

That is $(z_{dom} - 2)^2 + (z_{dom} - 2)^2 + 4z_{dom} - 2\psi z_{dom} + 2x_{dom}(\psi - 2) \leq 0$.

Hence $2(z_{dom} - 2)^2 + 2(z_{dom} - x_{dom})(2 - \psi) \leq 0$, a contradiction: indeed, the sum of a positive term and a strictly positive term ($x_{dom} < z_{dom}$ by the above statement) cannot be negative or zero. \square

In other words, \mathbf{d} is the only point where the ball \mathcal{B}_ψ intersects the surface of P_{dom} .

We now explain the relationship between the positivity of $(u_n(\mathbf{d}))_{n \in \mathbb{N}}$ and $L(\theta)$, which is the crux of the proof of Theorem 2.6 by [OW14b].

Lemma 5.2. *For every $\ell, \varepsilon > 0$, we can compute an $n_1 > 0$ such that for all $n \geq n_1$, we have $u_n(\mathbf{d}) \geq 0$ and $2\pi n\theta \bmod 2\pi \in [0, \pi]$ together imply $n[2\pi n\theta] > 2\pi\ell - \varepsilon$.*

Proof. Assume that $u_n(\mathbf{d}) \geq 0$ and $2\pi n\theta \bmod 2\pi \in [0, \pi]$. First, notice that as $2\pi n\theta \bmod 2\pi \in [0, \pi]$, we have

$$u_n(\mathbf{d}) = 2n - 2n \cos([2\pi n\theta]) - 2\pi\ell \sin([2\pi n\theta]).$$

Let $\alpha_n = [2\pi n\theta] \geq 0$. Considering the Taylor development for α_n close to 0 of $(1 - \cos(\alpha_n))$ and $\sin(\alpha_n)$, we get $u_n(\mathbf{d})/n = \frac{2}{2}\alpha_n^2 - \frac{2\pi\ell\alpha_n}{n} + f(\alpha_n)$, with $f(\alpha_n) = O(\alpha_n^3)$. We have $u_n(\mathbf{d}) \geq 0$, thus $\frac{2\pi\ell\alpha_n}{n}$ is smaller than $\alpha_n^2(1 + \frac{f(\alpha_n)}{\alpha_n^2})$, that is $\alpha_n \geq \frac{2\pi\ell}{n(1 + \frac{f(\alpha_n)}{\alpha_n^2})}$. There exists a computable value $\alpha_0 > 0$ such that $\alpha_n < \alpha_0$ implies $1 - \frac{\varepsilon}{2\pi\ell} \leq \frac{1}{(1 + \frac{f(\alpha_n)}{\alpha_n^2})} \leq 1 + \frac{\varepsilon}{2\pi\ell}$.

That is, if $\alpha_n \leq \alpha_0$, then $n[2\pi n\theta] > 2\pi\ell - \varepsilon$ and we are done.

We define $n_1 = \lfloor \frac{2\pi\ell - \varepsilon}{\alpha_0} \rfloor + 1$, and thus $n_1\alpha_0 > 2\pi\ell - \varepsilon$. Hence if $n > n_1$ and $\alpha_n > \alpha_0$, we also have $2\pi\ell - \varepsilon < n\alpha_n$. \square

We reason about the case $2\pi n\theta \bmod 2\pi \in [\pi, 2\pi)$ symmetrically with respect to the case $2\pi n\theta \bmod 2\pi \in [0, \pi)$: it suffices to consider the configuration $\mathbf{d}^- = (2, 2, 0, 0, 0, -2\pi\ell)$. If both $u_n(\mathbf{d}) \geq 0$ and $u_n(\mathbf{d}^-) \geq 0$, then $n[2\pi n\theta] > 2\pi\ell - \varepsilon$. The configuration \mathbf{d}^- can be handled by considering the ball \mathcal{B}_ψ^- of radius $\sqrt{2}\psi$ centred in $\mathbf{c}^- = (2 + \psi, 2 - \psi, 0, 0, 0, -2\pi\ell)$.

We turn now to the positivity of the balls $\mathcal{B}_\psi \cup \mathcal{B}_\psi^-$. We show that ψ can be chosen small enough such that the following is true:

Lemma 5.3. *For any ℓ, ε , let \mathbf{d}, \mathbf{d}' be defined as above. Then for any n_1 , we can choose ψ small enough and $n_2 > n_1$ such that for all $n > n_2$, if $n[2\pi n\theta] \geq 2\pi\ell + \varepsilon$, then $u_n(\mathbf{d}') \geq 0$ for all $\mathbf{d}' \in \mathcal{B}_\psi \cup \mathcal{B}_\psi^-$.*

Proof. We first consider the initial configuration $\mathbf{d}' \in \mathcal{B}_\psi \cup \mathcal{B}_\psi^-$ which minimises $u_n(\mathbf{d}')$:

Claim 5.4.

$$\min_{\mathbf{d}' \in \mathcal{B}_\psi \cup \mathcal{B}_\psi^-} u_n(\mathbf{d}') = n(2 - \psi)(1 - \cos(2\pi n\theta)) - 2\pi\ell |\sin(2\pi n\theta)| - 2\psi(\sqrt{n^2 + 1} - n).$$

Proof. Any point $\mathbf{d}' \in \mathcal{B}_\psi$ can be expressed as $\mathbf{c} + \mathbf{r}$, where \mathbf{c} is the centre of one of the two spheres, and \mathbf{r} is an arbitrary vector whose length does not exceed $\sqrt{2}\psi$. We know that $u_n(\mathbf{d}') = \mathbf{x}_n^T(\mathbf{c} + \mathbf{r})$, where $\mathbf{x}_n = [n, n \cos(2\pi n\theta), n \sin(2\pi n\theta), 1, \cos(2\pi n\theta), \sin(2\pi n\theta)]^T$ is fixed. As discussed in the previous lemma, the minimum contribution from $\mathbf{x}_n^T \mathbf{c}$ is (note that the choice is over the centres of the two spheres)

$$n(2 + \psi) - n(2 - \psi) \cos(2\pi n\theta) - 2\pi\ell |\sin(2\pi n\theta)|$$

which can also be rearranged as

$$2n\psi + n(2 - \psi)(1 - \cos(2\pi n\theta)) - 2\pi\ell |\sin(2\pi n\theta)|.$$

It now remains to independently optimise over \mathbf{r} . For this, we note $\mathbf{x}_n^T \mathbf{r}$ is minimised when \mathbf{r} has longest possible length, and is oriented opposite to \mathbf{x}_n . In this case, $\mathbf{x}_n^T \mathbf{r}$ will be the negative product of the lengths of \mathbf{x}_n and \mathbf{r} . This is $-\sqrt{2n^2 + 2} \cdot \sqrt{2}\psi$, which simplifies to

$$-2\psi\sqrt{n^2 + 1}.$$

Adding the two contributions gives the result. \square

Thus, we have $\frac{u_n(\mathbf{d}')}{n} = (2 - \psi)(1 - \cos(2\pi n\theta)) - \frac{2\pi\ell}{n} |\sin(2\pi n\theta)| - 2\psi \left(\sqrt{1 + \frac{1}{n^2}} - 1 \right)$.

We let $\alpha_n = [2\pi n\theta] \geq 0$. Note that $|\sin \alpha_n| = \sin[\alpha_n]$

We first show that if $n\alpha_n > 36\pi$ then $\mathcal{B}_\psi \cup \mathcal{B}_\psi^-$ is positive, i.e. \mathbf{d}' is:

Claim 5.5. If $\alpha_n n > 36\pi$, then $u_n(\mathbf{d}') > 0$ for all $\mathbf{d}' \in \mathcal{B}_\psi \cup \mathcal{B}_\psi^-$.

Proof. Assume that $\alpha_n n > 36\pi$. We use $\psi \leq 1$, $\sqrt{1 + y} - 1 \leq \frac{y}{2}$, $\sin y \leq y$ and $1 - \cos y \geq \frac{y^2}{2} - \frac{y^4}{24}$ to obtain:

$$\frac{u_n(\mathbf{d}')}{n} \geq \frac{\alpha_n^2}{2} - \frac{\alpha_n^4}{24} - \frac{2\pi\ell\alpha_n}{n} - \frac{\psi}{n^2}. \quad (5.1)$$

As $\alpha_n \leq \pi$, we have $\alpha_n^2 < 10$ and $\frac{\alpha_n^2}{2} - \frac{\alpha_n^4}{24} \geq \frac{\alpha_n^2}{12}$. That is:

$$\frac{u_n(\mathbf{d}')}{n} \geq \alpha_n \left(\frac{\alpha_n}{12} - \frac{2\pi\ell}{n} \right) - \frac{\psi}{n^2}.$$

As $\ell \leq 1$ and $\alpha_n \geq \frac{36\pi}{n}$, we have:

$$\begin{aligned} \frac{u_n(\mathbf{d}')}{n} &\geq \alpha_n \left(\frac{3\pi}{n} - \frac{2\pi}{n} \right) - \frac{\psi}{n^2} \\ &\geq \frac{36\pi^2}{n^2} - \frac{\psi}{n^2}. \end{aligned}$$

Hence, as $\psi \leq 1 < 36\pi^2$, we have $u_n(\mathbf{d}') > 0$. \square

Hence, we can assume without loss of generality that $\alpha_n n \leq 36\pi$.

Assume now that $2\pi\ell + \varepsilon \leq \alpha_n n \leq 36\pi$. We want to show that $u_n(\mathbf{d}') \geq 0$. Assume by contradiction that $\frac{u_n(\mathbf{d}')}{n} < 0$. Similarly as (5.1), we have:

$$(2 - \psi)\left(\frac{\alpha_n^2}{2} - \frac{\alpha_n^4}{24}\right) - \frac{2\pi\ell\alpha_n}{n} - \frac{\psi}{n^2} \leq \frac{u_n(\mathbf{d}')}{n} < 0.$$

We choose $\tau_1 = \frac{2\pi\ell}{2\pi\ell + \frac{\varepsilon}{3}} < 1$ and $0 < \psi \leq 2 - 2\tau_1$. We get:

$$\tau_1\alpha_n^2 - \tau_1\frac{\alpha_n^4}{12} < \frac{2\pi\ell}{n}\alpha_n + \frac{\psi}{n^2}.$$

We multiply by $\frac{n}{\tau_1\alpha_n} > 0$ and obtain:

$$\alpha_n n < \frac{2\pi\ell}{\tau_1} + \frac{\alpha_n^3 n}{12} + \frac{\psi}{\tau_1\alpha_n n}.$$

Now, $\frac{1}{\alpha_n n} \leq \frac{1}{2\pi\ell}$, $\alpha_n^3 n = O(\frac{1}{n^2})$ as $\alpha_n \leq \frac{36\pi}{n}$, and choosing n_2 with $\left(\frac{36\pi}{n_2}\right)^3 n_2 \leq \frac{12\varepsilon}{3}$ and ψ with $\psi < \frac{2\tau_1\pi\ell\varepsilon}{3}$, we obtain the contradiction:

$$\alpha_n n < (2\pi\ell + \frac{\varepsilon}{3}) + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = 2\pi\ell + \varepsilon. \quad \square$$

Now, assuming we have an oracle for ψ -robust positivity, we shall show how to use Lemma 5.2 (along with its symmetric analog for \mathcal{B}_ψ^-) and Lemma 5.3 to deduce either $L(\theta) < \ell + \varepsilon$ or $L(\theta) > \ell - \varepsilon$. These lemmata directly allow us to also reason about $\liminf n[2\pi n\theta]$ just the way we do about $\inf n[2\pi n\theta]$, i.e. the reduction from computing $L_\infty(\theta)$ to uniform ψ -robust ultimate positivity (Theorem 3.2.1.c) is a straightforward adaptation of the proof we give of Theorem 3.2.1.a.

Proof of Theorem 3.2 for robust positivity and closed balls. Let $\varepsilon > 0$. Assume that an ℓ has been fixed, such that we want to know either $L(\theta) < \ell + \varepsilon$ or $L(\theta) > \ell - \varepsilon$. First, we fix n_2 and $\psi < \frac{1}{3}$ and $\psi < \pi\ell$ as per Lemma 5.3. The high level plan is to use the robust positivity oracle to query whether neighbourhoods that are subsets of \mathcal{B}_ψ (symmetrically \mathcal{B}_ψ^-) containing \mathbf{d} (symmetrically \mathbf{d}^-) are positive from iterate n_2 onwards. If yes, it means that \mathbf{d} and \mathbf{d}^- in particular are positive, and we use Lemma 5.2 and its symmetric statement to argue $L_{\geq n_2}(\theta) > \ell - \varepsilon$. If not, it means that the larger \mathcal{B}_ψ and \mathcal{B}_ψ^- themselves are not positive, and we use the contrapositive of Lemma 5.3 to argue that $L_{\geq n_2}(\theta) < \ell + \varepsilon$.

To complete the reduction, we need to decide whether $L_{\leq n_2}(\theta) = \frac{1}{2\pi} \inf_{0 < n \leq n_2} n[2\pi n\theta] < \ell + \varepsilon$. This is straightforward as it only involves computing $n[2\pi n\theta]$ to sufficient precision for a bounded number of indices n . If $L_{\leq n_2}(\theta) < \ell + \varepsilon$, then we know $L(\theta) \leq L_{\leq n_2}(\theta) < \ell + \varepsilon$.

We remark that \mathcal{B}_ψ (symmetrically \mathcal{B}_ψ^-) corresponds to a ball in the coordinates $(z_{dom}, x_{dom}, y_{dom}, z_{res}, x_{res}, y_{res})$ of the coefficient space. In fact, as outlined in the plan above, our reduction needs to supply the initialisation $(v_{n_2}, \dots, v_{n_2+5})$ as input in the original coordinates $(v_0, v_1, v_2, v_3, v_4, v_5)$. The ball \mathcal{B}_ψ is mapped to an hyper-ellipsoid \mathcal{O} in the original coordinates. We can explicitly define a smaller ball $\mathcal{B}' \subseteq \mathcal{O}$ in the original coordinates, containing (the image of) $d \in \mathcal{B}'$. Hence \mathcal{B} is positive implies that \mathcal{B}' is.

We symmetrically do the same thing for \mathcal{B}_ψ^- , and query both neighbourhoods, thus implementing the reduction as outlined above.

Finally, we show the statement that we can restrict the balls to having rational radius and centre with real algebraic entries. The radius restriction is simple, as we can choose \mathcal{B}' arbitrarily small. Hence in particular we can choose the radius to be rational.

For the centre, notice that the initial configuration $\mathbf{d} = (2, 2, 0, 0, 0, 2\pi\ell)$ is not a priori algebraic. We can however restrict ourselves to choosing ℓ of the form $\frac{q}{\pi}$, with $q \in \mathbb{Q}$. This does not impede our search for lower and upper bounds on $L_{\leq n_2}(\theta)$: we can define $r, r' \in \mathbb{Q}$ such that $r < \frac{q}{\pi} - 0.45\varepsilon < L_{\leq n_2}(\theta) < \frac{q}{\pi} + 0.45\varepsilon < r'$, with $r' - r < \varepsilon$.

Now, this choice of $\ell = \frac{q}{\pi}$ makes $\mathbf{d} = (2, 2, 0, 0, 0, 2\pi\ell)$ rational. As the linear operator $M^n H$ that transforms the coefficient space to the input space is algebraic, it means that $M^n H\mathbf{d}$ is algebraic as well. The normal to \mathcal{B} at d is rational, thus the normal to $\mathcal{O} = M^n H\mathcal{B}$ at the algebraic $M^n H\mathbf{d}$ is algebraic, and we obtain a centre of $\mathcal{B}' \subset \mathcal{O}$ that is algebraic (since the radius is rational). \square

This completes the proof of Theorem 3.2.1.a for closed balls.

5.3. Case of Open Balls and robust Skolem. In this subsection, we extend the proof of Theorem 3.2 to show that considering open or closed balls does not make a difference for the Diophantine hardness. Further, there is also no difference whether we consider the robust Skolem problem (0 is avoided), the robust positivity problem (negative numbers are avoided), or the robust strict positivity problem (negative and 0 are avoided). Thus, this establishes Theorem 3.2.1.b; and 1.a for open balls.

Let \mathcal{B}_ψ be an open ball and $cl(\mathcal{B}_\psi)$ its topological closure, which is the closed ball consisting of \mathcal{B}_ψ and its surface. With the next lemma, we argue that open and closed balls can very often be reasoned about interchangeably. Consider the following statements:

- (1) Robust positivity holds for the closed ball $cl(\mathcal{B}_\psi)$.
- (2) Robust positivity holds for the open ball \mathcal{B}_ψ .
- (3) Robust strict positivity holds for the open ball \mathcal{B}_ψ .
- (4) Robust Skolem holds for the open ball \mathcal{B}_ψ .
- (5) Robust strict positivity holds for the closed ball $cl(\mathcal{B}_\psi)$.
- (6) Robust Skolem holds for the closed ball $cl(\mathcal{B}_\psi)$.

We show that equivalence results between these statements. This allows us to conclude that having open or closed balls does not make a difference for \mathcal{T} -Diophantine hardness of Skolem and (strict) positivity. Formally, we have the following.

Lemma 5.6. *(1), (2) and (3) are equivalent. Further, for balls \mathcal{B}_ψ containing at least one initial configuration \mathbf{d}_0 in its interior that is strictly positive, i.e. $u_n(\mathbf{d}_0) > 0$ for all n , both (3) and (4) are equivalent and (5) and (6) are equivalent.*

Proof. (1) implying (2) is trivial. (2) implies (1): we show the contrapositive. Suppose there exists an initial configuration \mathbf{d} on the surface of the ball \mathcal{B}_ψ and an integer n such that $u_n(\mathbf{d}) = y < 0$. Recall that M is the companion matrix, and $u_n(\mathbf{d})$ is the first component of $(M^n \cdot \mathbf{d})$, so $u_n(x)$ is a continuous function. Thus, there exists a neighbourhood of \mathbf{d} , such that for all \mathbf{d}' in the neighbourhood, $u_n(\mathbf{d}') < y/2 < 0$. This neighbourhood intersects the open ball \mathcal{B}_ψ enclosed by the surface, and picking d' in this intersection shows that Robust Positivity does not hold in the open ball.

(3) implying (2) is trivial. (2) implies (3): Assume for the sake of contradiction that there is an initial configuration \mathbf{c}' in the open ball \mathcal{B}_ψ such that $u_n(\mathbf{c}') = 0$. Consider any open O around \mathbf{c}' entirely in the open ball \mathcal{B}_ψ . We have that \mathbf{c}' is on hyperplane H_n by

definition. That is, there are initial configurations in O on both sides of H_n . In particular, there is an initial configuration \mathbf{c}'' in O , hence in \mathcal{B}_ψ , with $\mathbf{c}'' \notin H_n^+$, i.e. $u_n(\mathbf{c}'') < 0$, a contradiction with \mathcal{B}_ψ being robustly positive.

(3) implies (4) is trivial. (4) implies (3): We consider the contrapositive: if we have an initial configuration \mathbf{d}_1 of \mathcal{B}_ψ which is not strictly positive, then $u_n(\mathbf{d}_1) \leq 0$ for some n , and there is a barycenter \mathbf{d}_2 between $\mathbf{d}_0, \mathbf{d}_1$ which satisfies $u_n(\mathbf{d}_2) = 0$, i.e. negation of (4). To be more precise, we can choose $\mathbf{d}_2 = \frac{-u_n(\mathbf{d}_1)}{u_n(\mathbf{d}_1) - u_n(\mathbf{c})} \mathbf{d}_0 + \frac{u_n(\mathbf{d}_1)}{u_n(\mathbf{d}_1) - u_n(\mathbf{d}_0)} \mathbf{d}_1$.

Now, (5) and (6) are equivalent for balls containing at least one initial configuration \mathbf{d}_0 that is strictly positive in its interior (same proof as for the equivalence between (3) and (4) above). However, notice that (5,6) are not equivalent with (1,2,3,4) in general. \square

We are now ready to prove Theorem 3.2 for open balls \mathcal{B}_ψ . It suffices to remark that the centre \mathbf{c} of \mathcal{B}_ψ is strictly in the interior of P_{dom} , and thus it will be eventually strictly positive by Lemma 4.1, that is there exists $n_2 > n_1$ such that $u_n(\mathbf{c}) > 0$ for all $n > n_2$, and we can choose $\mathbf{d}_0 = \mathbf{M}^{n_2} \mathbf{c}$. Hence by Lemma 5.6, robustness (for $n > \max(n_1, n_2)$) of positivity, strict positivity and Skolem are equivalent on \mathcal{B} , and these are equivalent with robust positivity of $cl(\mathcal{B})$ which was proved \mathcal{T} -Diophantine hard in the previous section.

It remains to prove Theorem 3.2 for robust Skolem for closed balls \mathcal{B} . For that, it suffices to easily adapt Lemma 5.2, replacing $(u_n(\mathbf{d}))_{n > n_2}$ positive by strictly positive, and obtain the \mathcal{T} -Diophantine hardness for robust strict positivity of closed balls. We again apply Lemma 5.6 ((5) and (6) are equivalent) to obtain hardness for robust Skolem of closed balls.

5.4. Case of Robust Ultimate Positivity. We now turn to approximating $L_\infty(\theta)$ using robust ultimate positivity, i.e., to show Theorem 3.2.1.c. The idea is similar to approximating $L(\theta)$ that we developed in previous sections.

Let $\ell > 0$ and $\varepsilon > 0$. We use n_2 and ψ as defined in Lemma 5.3.

Proposition 5.7. *\mathcal{B}_ψ is (uniformly) ultimately positive implies that $L_\infty(\theta) \geq \ell - \varepsilon$, and \mathcal{B}_ψ is not (uniformly) ultimately positive implies that $L_\infty(\theta) \leq \ell + \varepsilon$.*

Proof. Recall $L_\infty = \frac{1}{2\pi} \liminf_{n > 0} n[2\pi n\theta]$.

Assume that \mathcal{B}_ψ is (uniformly) ultimately positive. Then in particular \mathbf{d} is positive. Applying Lemma 5.2, we obtain that $n[2\pi n\theta] < 2\pi\ell - \varepsilon$ for only finitely many n , and hence the limit infimum $L_\infty(\theta) \geq \ell - \varepsilon$.

Now, assume that \mathcal{B}_ψ is not (uniformly) ultimately positive. Then by Lemma 5.3, which we can apply as ψ is sufficiently small, for infinitely many $n > n_2$, we have that $n[2\pi n\theta] < 2\pi\ell + \varepsilon$. Hence the limit infimum $L_\infty(\theta) \leq \ell + \varepsilon$. \square

This settles the case of closed balls. Now, when we have a uniform constant on ultimate positivity, the case of open balls is equivalent to the case of closed balls, as shown in the following Lemma:

Lemma 5.8. *Let \mathcal{B}_ψ be an open ball. Then $cl(\mathcal{B}_\psi)$ is robustly uniformly ultimately positive iff \mathcal{B}_ψ is robustly uniformly ultimately positive.*

Proof. One implication is obvious as $\mathcal{B}_\psi \subseteq cl(\mathcal{B}_\psi)$. For the other, assume that \mathcal{B}_ψ is robustly uniformly ultimately positive. Let N be the uniform constant for ultimate positivity over \mathcal{B}_ψ . So for all $n > N$, we know that for all $\mathbf{d}' \in \mathcal{B}_\psi$, $u_n(\mathbf{d}') \geq 0$. Fix a $n > N$. By continuity, we obtain that for all $\mathbf{d}' \in cl(\mathcal{B}_\psi)$, we also have $u_n(\mathbf{d}') \geq 0$. Hence $cl(\mathcal{B}_\psi)$ is robustly uniformly ultimately positive with the same uniform ultimate positive constant. \square

This terminates the proof of the last cases of Theorem 3.2. Notice that this last Lemma is not true when the constant on ultimate positivity is not uniform over the ball. And indeed, we will show a major complexity difference in the next section.

6. PROOF OF THEOREM 3.3

We now turn to the proof of Theorem 3.3, generalising elements from Section 4.

6.1. Intuitions for the proof of Theorem 3.3. Let $(u_n)_{n \in \mathbb{N}}$ be a recurrence relation defined by coefficients $\mathbf{a} \in \mathbb{Q}^\kappa$. As before, we will consider $(v_n)_{n \in \mathbb{N}} = (\frac{u_n}{f_n})_{n \in \mathbb{N}}$, for f_n such that the dominant coefficients of $(v_n)_{n \in \mathbb{N}}$ are of the form $\alpha e^{in\theta}$. We will then decompose the exponential solution of $(v_n)_{n \in \mathbb{N}}$ into two: a *dominant term* $(v_n^{dom})_{n \in \mathbb{N}}$ made of coefficients $\alpha e^{in\theta}$, and a *residue* $(v_n^{res})_{n \in \mathbb{N}}$ with $(v_n^{res})_{n \in \mathbb{N}} \xrightarrow{n \rightarrow +\infty} 0$. Recall that, denoting by \mathbf{c}^{dom} the projection of an initial configuration \mathbf{c} on dominant space, we defined $P_{dom} = \{\mathbf{c} \mid \forall n, v_n^{dom}(\mathbf{c}^{dom}) \geq 0\}$.

The non-negativity of $v_n^{dom}(\mathbf{c})$ for all n , and thus membership in P_{dom} , is necessary for the Ultimate Positivity of the LRS initialised by \mathbf{c} . However, the decidability of this prerequisite is unclear from the above formulation of P_{dom} . We define a function $\text{dominant}(\mathbf{c}, \mathbf{t})$, whose second argument \mathbf{t} comes from a continuous domain, namely the Masser Torus T [Mas88]. This torus T is compact, and Lemma 6.4 establishes an important closure property: P_{dom} can equivalently be defined as the set of points \mathbf{c} for which $\mu(\mathbf{c}) = \min_{\mathbf{t} \in T} \text{dominant}(\mathbf{c}, \mathbf{t})$ is non-negative. This definition is more accessible: Renegar's theorem [Ren92] states that $\mu(\mathbf{c})$ can be computed explicitly, which we will use in some cases but manage to avoid in others, as detailed below.

Before giving the formal proofs, we provide our intuitions for all 4 decidable cases. We start with the simplest case \exists -robust ultimate positivity, then move to robust non-uniform ultimate positivity, then \exists -robust positivity and finally \exists -robust Skolem.

(1) For \exists -robust uniform ultimate positivity, we consider different possibilities for $\mu = \mu(\mathbf{c}_0)$, and show the following in Proposition 6.5:

- If $\mu < 0$, then \mathbf{c}_0 is not in P_{dom} , that is $u_n(\mathbf{c}_0) < 0$ for an infinite number of indices $n \in \mathbb{N}$, so \mathbf{c}_0 cannot be ultimately positive and we are done.
- If $\mu = 0$, then \mathbf{c}_0 is on the surface of P_{dom} , so it is arbitrarily close to point not in P_{dom} , so \mathbf{c}_0 cannot be *robustly* ultimately positive.
- The last case $\mu > 0$, means that \mathbf{c}_0 is in the interior of P_{dom} , and so in particular there exists a closed ball \mathcal{B} around \mathbf{c}_0 entirely in the interior of P_{dom} . In particular, for some N large enough, \mathcal{B} is in the intersection of all the half-space $\mathcal{H}_n, n > N$, and thus \mathbf{c}_0 is robustly uniformly ultimately positive for the uniform bound N .

Thus, we have reduced deciding \exists -robust uniform ultimate positivity to checking the sign of $\mu(\mathbf{c}_0)$. As mentioned earlier we can explicitly compute $\mu(\mathbf{c}_0)$ using Renegar's result [Ren92] and see whether it is positive or not. In this case, we actually do not need to compute $\mu(\mathbf{c}_0)$ explicitly, only test whether it is strictly positive (since that is the only case where we can be \exists -robust ultimately positive), which can be formulated as an FO formula over Reals.

(2) The next case is robust *non-uniform* ultimate positivity for an *open* ball \mathcal{B} of algebraic radius ψ , centred in \mathbf{c}_0 . We show in Proposition 6.8 that \mathcal{B} is robustly *non-uniformly* ultimately positive iff $\mathcal{B} \subseteq P_{dom}$, i.e. iff $\mu(\mathbf{c}) > 0$ (the inequality is strict as \mathcal{B} is open

and P_{dom} is closed) for all $\mathbf{c} \in \mathcal{B}$, which can be effectively tested in PSPACE using a FO formula over reals.

(3) For \exists -robust positivity, we again do a case analysis according to the signs of $\mu = \mu(\mathbf{c}_0)$, by observing that:

- if $\mu < 0$, that is $\mathbf{c}_0 \notin P_{dom}$, then there exists a n such that the LRS initialised by \mathbf{c}_0 is negative, so $(u_n(\mathbf{c}_0))_{n \in \mathbb{N}}$ is not robustly positive;
- if $\mu = 0$, that is \mathbf{c}_0 is at the surface of P_{dom} , then there exists a configuration arbitrarily close to \mathbf{c} such that the LRS from that configuration is negative, i.e. $(u_n(\mathbf{c}_0))_{n \in \mathbb{N}}$ is not robustly positive either;
- if $\mu > 0$, that is $\mathbf{c}_0 \in P_{dom}$ and not on the surface, then as the residue has negligible contribution to $(v_n)_{n \in \mathbb{N}}$ for large n , we show that the LRS will ultimately avoid negative numbers beyond a threshold index n_{thr} depending in the exact value of μ , which can be computed using Renegar's result [Ren92].

Having assured ourselves of the long run behaviour, it suffices to check the value of the LRS up to n_{thr} , where the residue can have significant contribution, to see whether the LRS is strictly positive, in which case it satisfies robust positivity.

(4) The last case is for \exists -robust Skolem. Analogous to robust positivity problem that sought to avoid negative numbers and hence required v_n^{dom} to be always positive, here we seek to avoid 0 and thus require v_n^{dom} to be non-zero, i.e. have strictly positive absolute value. We thus concern ourselves with $\nu = \min_{\mathbf{t} \in T} |\text{dominant}(\mathbf{c}, \mathbf{t})|$, where T is the continuous, compact torus. We split our analysis based on the possible values of ν , which we show in Proposition 6.11 can be computed effectively:

- $\nu > 0$. Then as the residue has negligible contribution to $(v_n)_{n \in \mathbb{N}}$ for large n , we show that the LRS will ultimately avoid zero beyond a threshold index n_{thr} . Having assured ourselves of the long run behaviour, it suffices to check the value of the LRS up to n_{thr} , where the residue can have significant contribution, to see whether the LRS satisfies robust Skolem.
- $\nu = 0$. Then we show in Proposition 6.12 that the LRS does not satisfy robust Skolem: no matter how small we pick a neighbourhood around \mathbf{c}_0 , there will always exist a \mathbf{c} in that neighbourhood that hits zero at some iteration.

The rest of this section formalises the above intuitions to prove Theorem 3.3.

6.2. Masser's Torus and relation with P_{dom} . We first define the normalised exponential polynomial solution $(v_n)_{n \in \mathbb{N}}$:

Definition 6.1. Let $(u_n)_{n \in \mathbb{N}}$ be an LRS of general term $u_n(\mathbf{c}) = \sum_{i=1}^r \sum_{j=0}^{m_i-1} p_{ij} n^j \gamma_i^n$, with ρ being the modulus of the dominant roots and $m+1$ the maximal multiplicity of a dominant root. Define $v_n(\mathbf{c}) = \frac{u_n(\mathbf{c})}{n^m \rho^n}$ for $n > 0$, and $v_0(\mathbf{c}) = u_0(\mathbf{c})$.

We call every term of v_n which converges towards 0 as n tends towards infinity residual, while the other terms, of the form $\alpha e^{i\theta}$ are dominant. We denote $\{\theta_j \mid j = 1, \dots, k\}$ the set of θ in dominant terms, and $\alpha_j(\mathbf{c})$ the associated coefficient. We denote the total contribution from the dominant terms of the form $\alpha_j(\mathbf{c}) e^{in\theta_j}$ by $v_n^{dom}(\mathbf{c})$. We have that:

$$v_n^{dom}(\mathbf{c}) = \sum \alpha_j(\mathbf{c}) e^{in\theta_j} \quad \text{and} \quad v_n^{res}(\mathbf{c}) = v_n(\mathbf{c}) - v_n^{dom}(\mathbf{c}) = \mathcal{O}\left(\frac{1}{n}\right) \rightarrow_{n \rightarrow \infty} 0.$$

As we explained in Section 6.1, knowing whether $(v_n^{dom})_{n \in \mathbb{N}} > 0$ for all $n \in \mathbb{N}$ is crucial in order to solve \exists -robust Skolem, positivity, and (uniform or not) ultimate positivity. In

Section 4 and 5, we dealt with hardness via an example which had 3 dominant roots, and it was rather simple to determine the min/max value (computed in the proof of Lemma 4.1). The general case is not so easy however. We follow the method defined in [OW14a, Theorem 4]. Towards this goal, define $\mathbf{s}^n = (e^{in\theta_1}, \dots, e^{in\theta_k})$ for all $n \in \mathbb{N}$. Computing the range of \mathbf{s}^n over \mathbb{N} is not simple. One idea is to perform a (continuous) relaxation, as provided by (a weak version of) Kronecker's Theorem:

Theorem 6.2 (Kronecker [Bou66]). *Let $\theta_1, \dots, \theta_k, \varphi_1, \dots, \varphi_k \in [0, 2\pi)$ such that for every tuple $(\lambda_1, \dots, \lambda_k)$ of integers with $e^{i(\lambda_1\theta_1 + \dots + \lambda_k\theta_k)} = 1$, we have $e^{i(\lambda_1\varphi_1 + \dots + \lambda_k\varphi_k)} = 1$.*

Then for any $\epsilon > 0$, there exists an arbitrarily large n such that for all $1 \leq j \leq k$ we have $|e^{in\theta_j} - e^{i\varphi_j}| \leq \epsilon$.

Intuitively, one can replace $\{\mathbf{s}^n = (e^{in\theta_1}, \dots, e^{in\theta_k}) \mid n \in \mathbb{N}\}$ by a set T as long as

- (1) for all n , $\mathbf{s}^n = (e^{in\theta_1}, \dots, e^{in\theta_k}) \in T$ and,
- (2) for all $(e^{i\varphi_1}, \dots, e^{i\varphi_k}) \in T$, we have $e^{i(\lambda_1\varphi_1 + \dots + \lambda_k\varphi_k)} = 1$ for every tuple $(\lambda_1, \dots, \lambda_k)$ of integers such that $e^{i(\lambda_1\theta_1 + \dots + \lambda_k\theta_k)} = 1$.

The first item ensures that T is a relaxation of $\{\mathbf{s}^n \mid n \in \mathbb{N}\}$, while the second ensures that $\{\mathbf{s}^n \mid n \in \mathbb{N}\}$ is dense in T by Kronecker's Theorem 6.2. To define such a set T , we have to extract the integer multiplicative relations between the roots $e^{i\theta_1}, \dots, e^{i\theta_k}$, to fulfill the second requirement. A tedious but now rather classical way to obtain these relations is by invoking a deep result from Masser [Mas88]:

Theorem 6.3 (Masser [Mas88]). *Let $K \in \mathbb{N}$ be fixed. Let $e^{i\theta_1}, \dots, e^{i\theta_k}$ be complex algebraic numbers of unit modulus, given as input. Consider the free abelian group L defined by $L = \{(\lambda_1, \dots, \lambda_k) \in \mathbb{Z}^k : e^{i(\lambda_1\theta_1 + \dots + \lambda_k\theta_k)} = 1\}$. The group L has a finite generator set $\{\mathbf{l}_1, \dots, \mathbf{l}_p\} \subset \mathbb{Z}^k$ with $p \leq k$. Each entry in the generator set is polynomially bounded in size($e^{i\theta_1}$), \dots , size($e^{i\theta_k}$). The generator set can be computed in polynomial space. Further, if the decision problem guarantees that k is at most K , the generator set can be computed in polynomial time.*

Using the generator set of L provided by Theorem 6.3, we can define a set T as a Torus as follows:

$$T = \{\mathbf{t} = (e^{i\varphi_1}, \dots, e^{i\varphi_k}) \mid \forall \mathbf{l} = (\ell_1, \dots, \ell_k) \in L, e^{i(\ell_1\varphi_1 + \dots + \ell_k\varphi_k)} = 1\}.$$

By definition of T , we have $\mathbf{s}^n \in T$ for all $n \in \mathbb{N}$, so the first requirement is fulfilled. For the second requirement, let $(e^{i\varphi_1}, \dots, e^{i\varphi_k}) \in T$ and $(\lambda_1, \dots, \lambda_k) \in \mathbb{Z}^k$ such that $e^{i(\lambda_1\theta_1 + \dots + \lambda_k\theta_k)} = 1$. By the definition of a generator set, there exist $\alpha_1, \dots, \alpha_p \in \mathbb{Z}^k$ with $(\lambda_1, \dots, \lambda_k) = \sum_1^p \alpha_i \mathbf{l}_i$. In particular, $e^{i(\lambda_1\varphi_1 + \dots + \lambda_k\varphi_k)} = \prod_{\mathbf{l}=(\ell_1, \dots, \ell_k) \in L} [e^{i(\ell_1\varphi_1 + \dots + \ell_k\varphi_k)}]^{\alpha_i} = \prod_{\mathbf{l}=(\ell_1, \dots, \ell_k) \in L} 1^{\alpha_i} = 1$, and the second requirement on T from above is also fulfilled.

Notice that the Torus T is *independent of the initial configuration \mathbf{c}* .

Denoting, for every initial configuration \mathbf{c} and element of the torus $\mathbf{t} \in T$, $\text{dominant}(\mathbf{c}, \mathbf{t}) = \sum_j \alpha_j(\mathbf{c}) t_j$, we obtain that $v_n^{\text{dom}}(\mathbf{c}) = \text{dominant}(\mathbf{c}, \mathbf{s}^n)$ for all n and all \mathbf{c} . We have that $v_n^{\text{dom}}(\mathbf{c}) \in \{\text{dominant}(\mathbf{c}, \mathbf{t}) \mid \mathbf{t} \in T\}$. The second statement translates into:

$$\forall \mathbf{t} \in T, \epsilon > 0, \exists n, \text{ with } \forall \mathbf{c}, |v_n^{\text{dom}}(\mathbf{c}) - \text{dominant}(\mathbf{c}, \mathbf{t})| \leq \epsilon. \quad (6.1)$$

Using the above, we can relate T with P_{dom} :

Lemma 6.4.

$$P_{\text{dom}} = \{\mathbf{c} \mid \forall \mathbf{t} \in T, \text{dominant}(\mathbf{c}, \mathbf{t}) \geq 0\}.$$

Proof. Let us denote $P'_{dom} = \{\mathbf{c} \mid \forall t \in T, \text{dominant}(\mathbf{c}, \mathbf{t}) \geq 0\}$. Remember that the definition is $P_{dom} = \{\mathbf{c} \mid \forall n \in \mathbb{N}, v_n^{dom}(\mathbf{c}) \geq 0\}$. We now prove that $P'_{dom} = P_{dom}$. As for all n and all \mathbf{c} , we have $v_n^{dom}(\mathbf{c}) = \text{dominant}(\mathbf{c}, \mathbf{s}^n)$, we have that $P'_{dom} \subseteq P_{dom}$ directly.

We now prove that $P_{dom} \subseteq P'_{dom}$. Let $\mathbf{c} \in P_{dom}$, i.e. $\forall n \in \mathbb{N}, v_n^{dom}(\mathbf{c}) \geq 0$. Assume by contradiction that $\mathbf{c} \notin P'_{dom}$. Hence there exists $\mathbf{t} \in T$ with $\text{dominant}(\mathbf{t}, \mathbf{c}) < 0$. Denote $\varepsilon = -\frac{\text{dominant}(\mathbf{t}, \mathbf{c})}{2} > 0$. Using (6.1), we obtain an n with $v_n^{dom}(\mathbf{c}) < -\frac{\varepsilon}{2} < 0$, a contradiction with $\mathbf{c} \in P_{dom}$. \square

6.3. Relation between sign of $\mu(\mathbf{c})$ and (robust) positivity. Let \mathbf{c} be an initial configuration and let T be Masser's Torus as described above. Then, we define

$$\mu(\mathbf{c}) = \min_{\mathbf{t} \in T} \text{dominant}(\mathbf{c}, \mathbf{t}). \quad (6.2)$$

We now state the crucial proposition that will allow us to relate the sign of $\mu(\mathbf{c})$ with the position of \mathbf{c} with respect to P_{dom} , which will be used to characterise \exists -robust positivity and ultimate positivity around \mathbf{c} .

Proposition 6.5.

- (1) If $\mu(\mathbf{c}) < 0$, then (i) $\mathbf{c} \notin P_{dom}$ and (ii) there exist an infinite number of n such that $v_n(\mathbf{c}) < 0$.
- (2) If $\mu(\mathbf{c}) = 0$, then (i) \mathbf{c} is on the surface of P_{dom} and (ii) $\forall \varepsilon > 0, \exists \mathbf{c}_\varepsilon$ such that $|\mathbf{c} - \mathbf{c}_\varepsilon| \leq \varepsilon$ and an infinite number of n with $v_n(\mathbf{c}_\varepsilon) < 0$.
- (3) If $\mu(\mathbf{c}) > 0$, then (i) \mathbf{c} is in the interior of P_{dom} and for all $\varepsilon > 0$, (ii) there exists N , such that for all \mathbf{c}_ε with $|\mathbf{c} - \mathbf{c}_\varepsilon| \leq \varepsilon$ and all $n > N$, we have $v_n(\mathbf{c}_\varepsilon) > 0$.

Proof. First, we will prove a useful claim, relating $v_n(\mathbf{c})$ and the distance from \mathbf{c} to hyperplane H_n : For every n , let $\text{distance}(\mathbf{c}, H_n)$ be the distance between an initial configuration \mathbf{c} and the hyperplane $H_n = \{\mathbf{c}' \mid v_n(\mathbf{c}') = 0\} = \{\mathbf{c}' \mid u_n(\mathbf{c}') = 0\}$.

Claim 6.6. There exists C such that for all n , $\text{distance}(\mathbf{c}, H_n) \leq C \cdot |v_n(\mathbf{c})|$.

Proof. Let $n \in \mathbb{N}$. We have $\text{distance}(\mathbf{c}, H_n) = \frac{|u_n(\mathbf{c})|}{\|\mathbf{y}\|}$ for \mathbf{y} the first row of \mathbf{M}^n by basic geometry. Let H be the transformation matrix between the basis of initial configurations and the basis of the exponential polynomial solution of $(u_n)_{n \in \mathbb{N}}$. Let $\mathbf{x} = (x_1, \dots, x_\kappa)$ with $x_i = n^k \rho_j^n$ so that to cover every root ρ_j and multiplicities $k = 1, \dots, m_j$. We have $u_n(\mathbf{c}) = \mathbf{y} \cdot \mathbf{c} = \mathbf{x} \cdot (H \cdot \mathbf{c})$ for all initial configurations \mathbf{c} , i.e., $\mathbf{y} = \mathbf{x} \cdot H$. That is, there exists a constant $D > 0$ depending upon H with $\|\mathbf{y}\| \geq Dn^m \rho^n$ for ρ the modulus of a dominant root and $m + 1$ the highest multiplicity of a root of modulus ρ . We obtain $\text{distance}(\mathbf{c}, H_n) \leq \frac{|u_n(\mathbf{c})|}{Dn^m \rho^n} = \frac{|v_n(\mathbf{c})|}{D}$. \square

Now we prove the relationship between the sign of $\mu(\mathbf{c})$ and the position wrt P_{dom} .

- If $\mu(\mathbf{c}) < 0$, then let $\varepsilon = -\mu(\mathbf{c}) > 0$. Using Equation (6.1), taking $t \in T$ realising the minimum value $\mu(\mathbf{c})$, we have an n such that $v_n^{dom}(\mathbf{c}) - \text{dominant}(\mathbf{c}, \mathbf{t}) \leq \frac{\varepsilon}{2}$. In particular, $v_n^{dom}(\mathbf{c}) \leq -\frac{\varepsilon}{2} < 0$ and $\mathbf{c} \notin P_{dom}$.
- If $\mu(\mathbf{c}) > 0$, then in particular for all n , $v_n^{dom}(\mathbf{c}) = \text{dominant}(\mathbf{c}, \mathbf{s}^n) \geq \mu(\mathbf{c}) > 0$, meaning that $\mathbf{c} \in P_{dom}$, in its interior.

- The last case is $\mu(\mathbf{c}) = 0$. Then by the previous reasoning, $\mathbf{c} \in P_{dom}$, but not necessarily in the interior. To see that it is on the surface of P_{dom} , it suffices to use Equation (6.1) again. By contradiction, assume that it is not the surface of P_{dom} . Hence there exists a distance $\varepsilon > 0$ to the surface, hence at least ε away from all hyperplanes H_n . Taking $t \in T$ realising the minimum value $\mu(\mathbf{c})$, we have an n such that $|v_n^{dom}(\mathbf{c}) - \text{dominant}(\mathbf{c}, \mathbf{t})| \leq \frac{\varepsilon}{2C}$. In particular, $v_n^{dom}(\mathbf{c}) \leq \frac{\varepsilon}{2C}$, a contradiction with $\text{distance}(\mathbf{c}, H_n) > \varepsilon$ using Claim 6.6.

This completes the proof of 1(i), 2(i) and 3(i).

We now turn to the existence of counterexamples of positivity in the neighbourhood of \mathbf{c} or of a neighbourhood of \mathbf{c} entirely positive, i.e., proof of 1(ii), 2(ii) and 3(ii). For this, we use Claim 6.6 in the following way: if for all $\alpha > 0$, there exist an infinite number of n_α with $|v_{n_\alpha}(\mathbf{c})| < \alpha$, then there exists an infinite number of n with $\text{distance}(\mathbf{c}, H_n) < \varepsilon$ (choose $n = n_\alpha$ for $\alpha = \frac{\varepsilon}{2C}$). This means that for any neighbourhood we pick, there will be a violation of positivity in it for infinitely many n .

- Let $\mu = 0$. We prove that for all $\alpha > 0$, we have a n_α such that $|v_{n_\alpha}(\mathbf{c})| \leq \alpha$, which suffices by Claim 6.6. Let $\alpha > 0$ arbitrarily small, and let N such that for all $n > N$, $|v_n^{res}(\mathbf{c})| \leq \frac{\alpha}{2}$. This N exists as $v_n^{res}(\mathbf{c}) \rightarrow_{n \rightarrow \infty} 0$. By Equation (6.1), as $\mu = 0$, there exists an (infinite number as n can be chosen arbitrarily large of) $n_\alpha > N$ with $|v_{n_\alpha}^{dom}(\mathbf{c})| < \frac{\alpha}{2}$. Thus $|v_{n_\alpha}(\mathbf{c})| \leq |v_{n_\alpha}^{dom}(\mathbf{c})| + |v_{n_\alpha}^{res}(\mathbf{c})| \leq \alpha$, thus proving 2(ii).
- Let $\mu < 0$. We still have $|v_n^{res}(\mathbf{c})| \leq \frac{\mu}{2}$ for all $n > N$. Again by Equation 6.1, we have an infinite number of $n > N$ such that $v_n^{dom}(\mathbf{c}) < \frac{\mu}{2}$, and we get $v_n(\mathbf{c}) = v_n^{dom}(\mathbf{c}) + v_n^{res}(\mathbf{c}) < \mu$, and thus 1(ii) follows.
- The last statement is with $\mu(\mathbf{c}) > 0$. We have that \mathbf{c} is guaranteed to be ε away from all hyperplanes H_n for sufficiently large n (i.e. $\geq N$). Hence considering the ball \mathcal{B} of radius $\frac{\varepsilon}{2}$ and of centre \mathbf{c} , all the initial configurations $\mathbf{c}_\varepsilon \in \mathcal{B}$ are at distance at least $\frac{\varepsilon}{2}$ to all hyperplanes H_n , and in particular we also have $v_n(\mathbf{c}_\varepsilon) > 0$. \square

6.4. Deciding \exists -robust ultimate positivity and

robust non-uniform ultimate positivity for open balls. We are now ready to provide the algorithms and proof for the decidable cases of robust ultimate positivity. First, we consider \exists -robust ultimate positivity. For this case, there is no difference on uniformity of open/closed ball.

Proposition 6.7. *Let \mathbf{c}_0 be an initial configuration. \mathbf{c}_0 is \exists -robustly ultimately positive iff $\mu(\mathbf{c}_0) > 0$. Further, this condition can be checked in PSPACE.*

Proof. We adapt a similar analysis as the one that appeared in [OW14a]. In one direction, if $\mu(\mathbf{c}_0) > 0$, then by Proposition 6.5 3., we infer that \mathbf{c}_0 is \exists -robust ultimately positive. In the other direction, suppose $\mu(\mathbf{c}_0) \leq 0$, then by Proposition 6.5 1. or 2. we obtain that for all ε -neighbourhoods of \mathbf{c}_0 there are an infinite number of n such that $v_n(\mathbf{c}_\varepsilon) < 0$, i.e., \mathbf{c}_0 is not \exists -robust ultimately positive.

For the complexity, we observe that to check whether $\mu(\mathbf{c}_0) > 0$, its complement can be checked by a call to an oracle for existential/universal first order theory of reals (which is decidable in PSPACE). That is, \mathbf{c}_0 is \exists -robust ultimately positive if the following first order (universal) theory of reals sentence holds:

$$\forall \mathbf{t} \in T, \text{dominant}(\mathbf{c}_0, \mathbf{t}) > 0. \quad (6.3)$$

Algorithm 1: to check robust non-uniform ultimate positivity for an open ball \mathcal{B} and to check \exists -robust ultimate positivity of \mathbf{c}_0 .

Input: Companion matrix $\mathbf{M} \in \mathbb{Q}^{\kappa \times \kappa}$ of $(u_n)_{n \in \mathbb{N}}$, initial configuration $\mathbf{c}_0 \in \mathbb{Q}^\kappa$, open ball \mathcal{B} with centre \mathbf{c}'_0 and radius ψ

- 1 Compute $\{\gamma_j\}_j \leftarrow$ eigenvalues of \mathbf{M} , $\rho \leftarrow \max_j |\gamma_j|$, $k \leftarrow |\{\gamma_i/\rho \mid |\gamma_i| = \rho\}|$, $\{e^{i\theta_j}\}_{j=1}^k \leftarrow \{\gamma_i/\rho \mid |\gamma_i| = \rho\}$
- 2 Determine T Torus obtained by applying Masser's result (Theorem 6.3) to $\{\theta_j\}_{j=1}^k$
- 3 **if** $\forall \mathbf{t} \in T$, $\text{dominant}(\mathbf{c}_0, \mathbf{t}) > 0$ (check using Equation 6.3) **then**
- 4 | **return** \mathbf{c}_0 is \exists -robust ultimate positive
- 5 **end**
- 6 **if** $\forall \mathbf{d} \in \mathbb{R}^\kappa, \forall \mathbf{t} \in T, \|\mathbf{d}\| \leq \psi \rightarrow \text{dominant}(\mathbf{c}'_0 + \mathbf{d}, \mathbf{t}) \geq 0$ (check using Equation 6.4) **then**
- 7 | **return** \mathcal{B} is robust non-uniform ultimate positive
- 8 **end**

In the above, we briefly remark that T is compact, and hence the minimum value of a dominant is well defined, should one need to compute it. The only difficulty is to express $\forall \mathbf{t} \in T$ as a universally quantified formula of the first order theory of reals. For that, it suffices to remark that $(a_1 + ib_1, \dots, a_k + ib_k) \in T$ iff $\|a_i + ib_i\|^2 = 1$ for all $i \leq k$ and $\forall \mathbf{l} = (\ell_1, \dots, \ell_k)$ in the generator set, where the complex part is $\text{Im}((a_1 + ib_1)^{\ell_1} \dots (a_k + ib_k)^{\ell_k}) = 0$ and the real part is $\text{Re}((a_1 + ib_1)^{\ell_1} \dots (a_k + ib_k)^{\ell_k}) = 1$. Notice that $a_i, b_i \in \mathbb{R}$ are real variables, and $(\ell_i)_{i \leq k}$ are fixed integers completely determined by the input, with magnitude polynomial in the size of the input. We can assume without loss of generality that $\ell_i \in \mathbb{N}$ for all i , as $(a_i + ib_i)^{-\ell} = (a_i - ib_i)^\ell$ since the modulus $\|a_i + ib_i\| = 1$. This makes the size of the obtained universal sentence polynomial in the size of the input, and we can indeed determine its truth in PSPACE. \square

Next, we turn to the decidability of robust *non-uniformly* ultimate positivity for open balls, which is slightly more complex.

Proposition 6.8. *Let \mathcal{B} an open ball with centre \mathbf{c}_0 and radius a real algebraic number ψ . \mathcal{B} is robustly non-uniformly ultimately positive iff $\mathcal{B} \subseteq P_{\text{dom}}$. Further, this condition can be checked in PSPACE.*

Proof. In one direction, if $\mathcal{B} \subseteq P_{\text{dom}}$, then for all $\mathbf{c} \in \mathcal{B}$, we have $\mathbf{c} \in P_{\text{dom}}$. In fact, we obtain that \mathbf{c} is in the relative interior of the set P_{dom} since \mathcal{B} is open. By Proposition 6.5, this implies that for any such $\mathbf{c} \in \mathcal{B}$, $\mu(\mathbf{c}) > 0$. This follows since if $\mu(\mathbf{c}) \leq 0$, that would imply \mathbf{c} is either on the surface of P_{dom} or not in P_{dom} at all, which would contradict the assumption. Now, let $N_{\mathbf{c}}$ such that $v_n^{\text{res}}(\mathbf{c})$ is negligible wrt $\mu(\mathbf{c})$. Then $u_n(\mathbf{c}) > 0$ for all $n > N_{\mathbf{c}}$. That is, \mathcal{B} is robustly (non uniformly) ultimately positive.

In the other direction, if $\mathcal{B} \not\subseteq P_{\text{dom}}$, then consider a configuration $\mathbf{c} \in \mathcal{B} \setminus P_{\text{dom}}$. That is, $\mu(\mathbf{c}) < 0$ (again by contradiction as $\mu(\mathbf{c}) \geq 0$ would mean \mathbf{c} is in P_{dom}). Applying Proposition 6.5, there exist an infinite number of n such that $u_n(\mathbf{c}) < 0$, that is, \mathcal{B} is not robustly ultimately positive.

Finally, we show how to effectively test whether $\mathcal{B} \subseteq P_{\text{dom}}$. Recall that \mathbf{c}_0 is the centre of \mathcal{B} and ψ its radius. Then, it suffices to solve the following first order theory of reals

Algorithm 2: to check \exists -robust positivity of initial configuration \mathbf{c}_0

Input: Companion matrix $\mathbf{M} \in \mathbb{Q}^{\kappa \times \kappa}$ of $(u_n)_{n \in \mathbb{N}}$ and initial configuration $\mathbf{c}_0 \in \mathbb{Q}^\kappa$

- 1 Compute $\{\gamma_j\}_j \leftarrow$ eigenvalues of \mathbf{M} , $\rho \leftarrow \max_j |\gamma_j|$, $k \leftarrow |\{\gamma_i/\rho \mid |\gamma_i| = \rho\}|$,
 $\{e^{i\theta_j}\}_{j=1}^k \leftarrow \{\gamma_i/\rho \mid |\gamma_i| = \rho\}$
- 2 Determine T Torus obtained by applying Masser's result (Theorem 6.3) to $\{\theta_j\}_{j=1}^k$
- 3 $\mu \leftarrow \min_{\mathbf{t} \in T} \text{dominant}(\mathbf{c}_0, \mathbf{t})$ ([Ren92])
- 4 **if** $\mu \leq 0$ **then**
- 5 | **return** NO (Proposition 6.5)
- 6 **else**
- 7 | Compute N_{thrpos} such that $|v_n^{res}(\mathbf{c}_0)| < \frac{\mu}{2}$ for all $n > N_{thrpos}$
- 8 | **foreach** $n \in \{0, 1, \dots, N_{thrpos}\}$ **do**
- 9 | | **if** $M^n \mathbf{c}_0 \leq 0$ **then**
- 10 | | | **return** NO
- 11 | | **end**
- 12 | **end**
- 13 | **return** YES
- 14 **end**

sentence, using the effective first order definability of Torus T as before.

$$\forall \mathbf{d} \in \mathbb{R}^\kappa, \forall \mathbf{t} \in T, \|\mathbf{d}\| \leq \psi \rightarrow \text{dominant}(\mathbf{c}_0 + \mathbf{d}, \mathbf{t}) \geq 0. \quad (6.4)$$

We remark in particular that this sentence is part of the universal fragment, and its truth can thus be decided in PSPACE, as required. \square

6.5. Decidability for \exists -robust positivity and \exists -robust Skolem. Finally, we consider decidability for \exists -robust Skolem and positivity as stated in Theorem 3.3, using Proposition 6.5. Unlike for ultimate positivity, where we do not need to resort to [Ren92], here we will have to compute $\mu(\mathbf{c}_0)$ explicitly, which can be done using Renegar's result. More precisely, for \exists -robust positivity and Skolem, we need to check a number of first steps. This number can be obtained effectively given $\min_{\mathbf{t} \in T} \text{dominant}(\mathbf{c}_0, \mathbf{t})$, which is precisely what is being computed by Renegar's result [Ren92].

The procedure for deciding \exists -robust positivity is given in Algorithm 2. The rationale for its correctness is as follows: first, we compute $\mu \leftarrow \min_{\mathbf{t} \in T} |\text{dominant}(\mathbf{c}_0, \mathbf{t})|$ using [Ren92], for \mathbf{c}_0 the initial configuration around which we are looking for a neighbourhood. If $\mu \leq 0$, then we declare \exists -robust positivity does not hold. Otherwise, we compute N_{thrpos} such that $|v_n^{res}(\mathbf{c}_0)| < \frac{\mu}{2}$ for all $n > N_{thrpos}$. Such a N_{thrpos} exists as $|v_n^{res}(\mathbf{c}_0)|$ tends to 0 for n tending towards ∞ .

Then, we check if $v_n(\mathbf{c}_0) \leq 0$ for some $n \leq N_{thrpos}$. If yes, then \exists -robust positivity does not hold as \mathbf{c}_0 will be arbitrarily close to initial configuration \mathbf{c} with $v_n(\mathbf{c}) < 0$. Otherwise, \exists -robust positivity holds, because we know that after N_{thrpos} steps, $|v_n^{res}(\mathbf{c}_0)| < \frac{\mu}{2}$, and thus $v_n(\mathbf{c}) = v_n^{dom}(\mathbf{c}_0) + v_n^{res}(\mathbf{c}_0) > \mu - \frac{\mu}{2} \geq \frac{\mu}{2}$. This lower bound ensures that there is a neighbourhood around \mathbf{c}_0 which remains positive.

Decidability for arbitrary real algebraic input is clear because each step of the algorithm is decidable (resp. computable). We now argue about the complexity when the input is

rational, and the a priori bound on the order of the LRS is hard-wired into the decision problem. The assumption on the bounded order ensures both μ and $1/\mu = 2^{s^{\mathcal{O}(1)}}$ are bounded by [Ren92]. We have $N_{thrpos} = 2^{s^{\mathcal{O}(1)}}$ because $v_n^{res}(\mathbf{c}_0) = \mathcal{O}(\frac{1}{n})$. This is the number of iterates we have to explicitly check, which gives the PSPACE complexity for rational input. This is because our strategy involves guessing the index of violation in coNP, constructing an integer straight line program that computes the term at the guessed step via iterated squaring, and then using a PosSLP oracle to check whether the centre of the neighbourhood is strictly positive at the guessed iterate. PosSLP was famously shown by Allender et. al. to be in PSPACE [ABKPM06]. A rational LRS can easily be scaled up to an integer LRS¹, that preserves the signs of each term, and enables it to be cast as a division-free integer straight line program.

We now turn to deciding Robust \exists -Skolem. While for \exists -Robust positivity we considered μ , for Skolem we will need its absolute value counterpart. Namely, we define:

$$\nu = \min_{\mathbf{t} \in T} |\text{dominant}(\mathbf{c}_0, \mathbf{t})|. \quad (6.5)$$

We now show how to adapt Renegar's result [Ren92, Theorems 1.1 and 1.2] to effectively compute ν when the order of LRS is bounded a priori.

Theorem 6.9 (Tarski [Tar51]). *For every formula $\tau(\mathbf{y})$ in the First Order Theory of the Reals, we can compute an equivalent quantifier-free formula $\chi(\mathbf{y})$.*

In specific scenarios, this computation can be considered efficient.

Theorem 6.10 (Renegar). *Let $M \in \mathbb{N}$ be fixed. Let $\tau(\mathbf{y})$ be a First Order formula with free variables \mathbf{y} , interpreted over the Theory of the Reals. Assume that the total number of free and bound variables in $\tau(\mathbf{y})$ is bounded by M . Denote the maximum degree of the polynomials in $\tau(\mathbf{y})$ by d and the number of atomic predicates in $\tau(\mathbf{y})$ by n . Then there is a procedure which computes an equivalent quantifier-free formula*

$$\chi(\mathbf{y}) = \bigvee_{i=1}^I \bigwedge_{j=1}^{J_i} h_{i,j}(\mathbf{y}) \sim_{i,j} 0$$

in disjunctive normal form, where each $\sim_{i,j}$ is either $>$ or $=$, with the following properties:

- (1) Each of the i and J_i (for $1 \leq i \leq I$) is bounded by a polynomial in (nd) .
- (2) The degree of $\chi(\mathbf{y})$ is bounded by a polynomial in (nd) .
- (3) The height of $\chi(\mathbf{y})$, i.e. the largest coefficient in the polynomials in χ is bounded by $2^{\text{size}(\tau(\mathbf{y}))(nd)^{\mathcal{O}(1)}}$

Moreover, the procedure runs in time polynomial in the size of the input formula $\tau(\mathbf{y})$.

Proposition 6.11. μ, ν are algebraic and computable (in finite time). Further, if K is a bound on the order of the LRS: (a) this computation runs in polynomial time, and furthermore, (b) if μ, ν are nonzero, we have $|\mu|, |\nu| < 2^{s^{\mathcal{O}(1)}}$ and $\frac{1}{|\mu|}, \frac{1}{|\nu|} < 2^{s^{\mathcal{O}(1)}}$

¹To scale the characteristic roots by ρ , replace the characteristic polynomial $X^\kappa - a_{\kappa-1}X^{\kappa-1} \dots - a_0$ by $X^\kappa - \rho a_{\kappa-1}X^{\kappa-1} - \rho^\kappa a_0$. Similarly, replace the initial terms by $u_0, \rho u_1, \dots, \rho^{\kappa-1} u_{\kappa-1}$. This results in replacing the sequence $(u_n)_n$ by the integer $(\rho^n u_n)_n$, for an appropriate choice of ρ , e.g. least common multiple of denominators.

Proof. We obtain the statement directly from Renegar [Ren92], i.e. μ, ν are roots of polynomials obtained via quantifier elimination on the following formulae that can be expressed in the First Order Theory of Reals. μ is the unique satisfying assignment to

$$(\forall \mathbf{t}. \mathbf{t} \in T \Rightarrow \text{dominant}(\mathbf{c}_0, \mathbf{t}) \geq \mu) \wedge (\exists \mathbf{t}_0 \in T. \text{dominant}(\mathbf{c}_0, \mathbf{t}_0) = \mu) \quad (6.6)$$

and is hence algebraic, with degree and height constrained by the formula.

Note that $|x| \geq b$ can be expressed as $(x \geq b \vee x \leq -b)$. Similarly, $|x| = b$ can be expressed as $b \geq 0 \wedge (x = b \vee x = -b)$. ν is the unique satisfying assignment to

$$(\forall \mathbf{t}. \mathbf{t} \in T \Rightarrow |\text{dominant}(\mathbf{c}_0, \mathbf{t})| \geq \nu) \wedge (\exists \mathbf{t}_0 \in T. |\text{dominant}(\mathbf{c}_0, \mathbf{t}_0)| = \nu). \quad (6.7)$$

Computability follows by performing quantifier elimination on the above formulae, which, as discussed in the proof of Proposition 6.7, will have size polynomial in that of the input. This yields polynomial equations, of which μ (resp. ν) are roots, and hence algebraic.

Now, if the order of the LRS is bounded by K , then this ensures that formulae (6.6) and (6.7) use boundedly many variables (depending on K). This is seen by generalising the arguments made in [OW14a, Section 3.1]. Essentially, from the bound on order one can derive bounds on the dimension of the torus \mathbb{T} and the number of constraints that define it, and thus the number of variables. Hence, we can apply Theorem 6.10. From this the complexity bounds on μ, ν follow as quantifier elimination runs in polynomial time under these assumptions. The equivalent quantifier free formula, and hence μ, ν themselves have degree polynomial and height single exponential in the size of the original formula (and thus the input). \square

For \exists -robust Skolem we adapt the above argument but with ν instead of μ . We start with the following useful property about ν .

Proposition 6.12. *If $\nu = 0$, then $\forall \varepsilon > 0$, $\exists \mathbf{c}_\varepsilon, n$ with $|\mathbf{c}_0 - \mathbf{c}_\varepsilon| \leq \varepsilon$ such that $v_n(\mathbf{c}_\varepsilon) = 0$.*

Proof. Assume $\nu = 0$, and let $\varepsilon > 0$. We again prove that for all $\alpha > 0$, we have a n_α such that $|v_{n_\alpha}(\mathbf{c}_0)| \leq \alpha$, which suffices by Lemma 6.6. Let $\alpha > 0$ arbitrarily small, and let N such that for all $n > N$, $|v_n^{res}(\mathbf{c}_0)| \leq \frac{\alpha}{2}$. This N exists as $v_n^{res}(\mathbf{c}_0) \rightarrow_{n \rightarrow \infty} 0$. By Kronecker, as $\nu = 0$, there exists $n_\alpha > N$ with $|v_{n_\alpha}^{dom}(\mathbf{c}_0)| < \frac{\alpha}{2}$. Thus $|v_{n_\alpha}(\mathbf{c}_0)| \leq |v_{n_\alpha}^{dom}(\mathbf{c}_0)| + |v_{n_\alpha}^{res}(\mathbf{c}_0)| \leq \alpha$. \square

The procedure to decide \exists -robust Skolem is given in Algorithm 3. Basically, we compute $\nu \leftarrow \min_{\mathbf{t} \in T} |\text{dominant}(\mathbf{c}, \mathbf{t})|$ using Proposition 6.11, for \mathbf{c}_0 the initial configuration around which we are looking for a neighbourhood. If $\nu = 0$, we declare \exists -robust Skolem does not hold, according to Proposition 6.12.

Otherwise, we compute N_{thrSk} such that $|v_n^{res}(\mathbf{c}_0)| < \frac{\nu}{2}$ for all $n > N_{thrSk}$. Then we check if $v_n(\mathbf{c}_0) = 0$ for some $n \leq N_{thrSk}$. If yes, then \exists -robust Skolem does not hold. Otherwise we know that \exists -robust Skolem holds.

Indeed, if $\nu > 0$, for all $n > N_{thrSk}$, $|v_n(\mathbf{c}_0)| = |v_n^{dom}(\mathbf{c}_0) + v_n^{res}(\mathbf{c}_0)| > \nu - \frac{\nu}{2} \geq \frac{\nu}{2} > 0$, and this remains > 0 in a neighbourhood of \mathbf{c}_0 . Similarly, by hypothesis, $|v_n(\mathbf{c}_0)| > 0$ for the finite number of $n < N_{thrSk}$, and in particular we have a lower bound > 0 that ensure it stays strictly positive in a neighbourhood of \mathbf{c}_0 .

The complexity follows since both ν and $1/\nu = 2^{s^{\mathcal{O}(1)}}$ are bounded (Proposition 6.11) and thus have $N_{thrSk} = 2^{s^{\mathcal{O}(1)}}$ since $v_n^{res}(\mathbf{c}_0) = \mathcal{O}(\frac{1}{n})$.

This finally completes the proof of Theorem 3.3.

Algorithm 3: to check \exists -robust Skolem of initial configuration \mathbf{c}_0

Data: Companion matrix $\mathbf{M} \in \mathbb{Q}^{\kappa \times \kappa}$ of $(u_n)_{n \in \mathbb{N}}$ and initial configuration $\mathbf{c}_0 \in \mathbb{Q}^{\kappa}$

- 1 $\{\gamma_j\}_j \leftarrow$ eigenvalues of \mathbf{M} , $\rho \leftarrow \max_j |\gamma_j|$, $\{e^{i\theta_j}\}_{j=1}^{\kappa} \leftarrow \{\gamma_i/\rho \mid |\gamma_i| = \rho\}$
- 2 Determine T Torus obtained by applying Masser's result (Theorem 6.3) to $\{\theta_j\}_{j=1}^{\kappa}$
- 3 $\nu \leftarrow \min_{\mathbf{t} \in T} |\text{dominant}(\mathbf{c}_0, \mathbf{t})|$ (Proposition 6.11)
- 4 **if** $\nu = 0$ **then**
- 5 | **return** NO (Proposition 6.12)
- 6 **else**
- 7 | Compute N_{thrSk} such that $|v_n^{res}(\mathbf{c}_0)| < \frac{\nu}{2}$ for all $n > N_{thrSk}$
- 8 | **foreach** $n \in \{0, 1, \dots, N_{thrSk}\}$ **do**
- 9 | | **if** $\mathbf{M}^n \mathbf{c}_0 = 0$ **then**
- 10 | | | **return** NO
- 11 | | **end**
- 12 | **end**
- 13 | **return** YES
- 14 **end**

7. CONCLUSION

We have formulated a natural notion of robustness for the Skolem and (ultimate) positivity problems and shown several results: for a given neighbourhood around an initial configuration \mathbf{c}_0 , we show Diophantine hardness for the problems. Interestingly, this is the first Diophantine hardness result for a variant of Skolem as far as we know. This implies that for uninitialised positivity, the fact that the initial configuration \mathbf{c}_0 is arbitrary is crucial to obtain decidability [Tiw04, Bra06], as having a fixed ball around \mathbf{c}_0 is not sufficient. This is the case for all considered problems except the non-uniform variant of ultimate positivity for open balls, where we obtain decidability. Our results are for rational LRRs with the initial configuration having real-algebraic entries. We leave open hardness for the case of rational LRRs where the initial configuration has only rational entries.

We proved decidability of \exists -robust Skolem and (ultimate) positivity problems around an initial configuration in full generality with real-algebraic entries. These questions are also arguably more practical as in a real system, it is often impossible to determine the initial configuration with absolute accuracy. Our results can provide a precision with which it is sufficient to set the initial configuration. Beyond these technical results, we provided geometrical reinterpretations of Skolem/positivity, shedding a new light on this hard open problem.

REFERENCES

- [AAGT15] Manindra Agrawal, S. Akshay, Blaise Genest, and P. S. Thiagarajan. Approximate verification of the symbolic dynamics of Markov chains. *J. ACM*, 62(1):2:1–2:34, 2015.
- [ABGV22] S. Akshay, Hugo Bazille, Blaise Genest, and Mihir Vahanwala. On Robustness for the Skolem and Positivity Problems. In Petra Berenbrink and Benjamin Monmege, editors, *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. URL: <https://drops.dagstuhl.de/opus/volltexte/2022/15815>, doi:10.4230/LIPIcs.STACS.2022.5.

- [ABKPM06] Eric Allender, Peter Bürgisser, Johan Kjeldgaard-Pedersen, and Peter Miltersen. On the complexity of numerical analysis. *SIAM Journal on Computing*, 38, 01 2006. doi:10.1109/CCC.2006.30.
- [ABM⁺20] S. Akshay, Nikhil Balaji, Aniket Murhekar, Rohith Varma, and Nikhil Vyas. Near-Optimal Complexity Bounds for Fragments of the Skolem Problem. In *37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020)*, volume 154, pages 37:1–37:18, 2020.
- [ABV17] S. Akshay, Nikhil Balaji, and Nikhil Vyas. Complexity of restricted variants of Skolem and related problems. In Kim G. Larsen, Hans L. Bodlaender, and Jean-François Raskin, editors, *42nd International Symposium on Mathematical Foundations of Computer Science, MFCS 2017, August 21-25, 2017 - Aalborg, Denmark*, volume 83 of *LIPICs*, pages 78:1–78:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
- [AGKV16] S. Akshay, Blaise Genest, Bruno Karelövic, and Nikhil Vyas. On regularity of unary probabilistic automata. In *STACS'16*, pages 8:1–8:14. *LIPICs*, 2016.
- [AGV18] S. Akshay, Blaise Genest, and Nikhil Vyas. Distribution-based objectives for Markov Decision Processes. In *33rd Symposium on Logic in Computer Science (LICS 2018)*, volume IEEE, pages 36–45, 2018.
- [BFJ⁺20] Christel Baier, Florian Funke, Simon Jantsch, Toghrol Karimov, Engel Lefauchaux, Joël Ouaknine, Amaury Pouly, David Purser, and Markus A. Whiteland. Reachability in dynamical systems with rounding. In Nitin Saxena and Sunil Simon, editors, *40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2020, December 14-18, 2020, BITS Pilani, K K Birla Goa Campus, Goa, India (Virtual Conference)*, volume 182 of *LIPICs*, pages 36:1–36:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- [BHM14] M. Biscaia, D. Henriques, and P. Mateus. Decidability of approximate Skolem problem and applications to logical verification of dynamical properties of Markov chains. *ACM Trans. Comput. Logic*, 16(1), December 2014.
- [BLN⁺22] Yuri Bilu, Florian Luca, Joris Nieuwveld, Joël Ouaknine, David Purser, and James Worrell. Skolem meets Schanuel. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22-26, 2022, Vienna, Austria*, volume 241 of *LIPICs*, pages 20:1–20:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [BM76] Jean Berstel and Maurice Mignotte. Deux propriétés décidables des suites récurrentes linéaires. *Bulletin de la Société Mathématique de France*, 104:175–184, 1976.
- [Bou66] N. Bourbaki. *Elements of Mathematics: General Topology (Part 2)*. Addison-Wesley, 1966.
- [BP02] Vincent D Blondel and Natacha Portier. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. *Linear algebra and its Applications*, 351:91–98, 2002.
- [BPR06] S. Basu, R. Pollack, and M. F. Roy. *Algorithms in Real Algebraic Geometry*. Springer, 2nd edition, 2006.
- [Bra06] Mark Braverman. Termination of integer linear programs. In *International Conference on Computer Aided Verification*, pages 372–385. Springer, 2006.
- [Coh93] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1993.
- [DKM⁺21] Julian D’Costa, Toghrol Karimov, Rupak Majumdar, Joël Ouaknine, Mahmoud Salamati, Sadegh Soudjani, and James Worrell. The pseudo-Skolem problem is decidable. In Filippo Bonchi and Simon J. Puglisi, editors, *46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23-27, 2021, Tallinn, Estonia*, volume 202 of *LIPICs*, pages 34:1–34:21. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- [DKM⁺22] Julian D’Costa, Toghrol Karimov, Rupak Majumdar, Joël Ouaknine, Mahmoud Salamati, and James Worrell. The pseudo-reachability problem for diagonalisable linear dynamical systems. In Stefan Szeider, Robert Ganian, and Alexandra Silva, editors, *47th International Symposium on Mathematical Foundations of Computer Science, MFCS 2022, August 22-26, 2022, Vienna, Austria*, volume 241 of *LIPICs*, pages 40:1–40:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.
- [EvdPSW03] Graham Everest, Alfred J. van der Poorten, Igor E. Shparlinski, and Thomas Ward. Recurrence Sequences. In *Mathematical surveys and monographs*, 2003.

- [HHHK05] V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. Skolem’s problem - on the border between decidability and undecidability. Technical Report 683, Turku Centre for Computer Science, 2005.
- [Khi35] A Khintchine. Neuer beweis und verallgemeinerung eines hurwitzschen satzes. *Mathematische Annalen*, 111:631–637, 1935.
- [LS97] J. C. Lagarias and J. O. Shallit. Linear fractional transformations of continued fractions with bounded partial quotients. *Journal de théorie des nombres de Bordeaux*, 9:267–279, 1997.
- [Mah35] Kurt Mahler. *Eine arithmetische Eigenschaft der Taylor-koeffizienten rationaler Funktionen*. Noord-Hollandsche Uitgevers Mij, 1935.
- [Mas88] D. W. Masser. Linear relations on algebraic groups. In *New Advances in Transcendence Theory*. Cambridge University Press, 1988.
- [Mig82] M. Mignotte. Some useful bounds. In *Computer Algebra*, 1982.
- [MST84] Maurice Mignotte, Tarlok Nath Shorey, and Robert Tijdeman. The distance between terms of an algebraic recurrence sequence. *Journal für die reine und angewandte Mathematik*, 349:63–76, 1984.
- [Neu21] Eike Neumann. Decision problems for linear recurrences involving arbitrary real numbers. *Logical Methods in Computer Science*, 17(3), 2021.
- [OW14a] Joël Ouaknine and James Worrell. On the positivity problem for simple linear recurrence sequences. In *International Colloquium on Automata, Languages, and Programming*, pages 318–329. Springer, 2014.
- [OW14b] 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. SIAM, 2014.
- [OW14c] Joël Ouaknine and James Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In *International Colloquium on Automata, Languages, and Programming*, pages 330–341. Springer, 2014.
- [Ren92] James Renegar. On the computational complexity and geometry of the first-order theory of the reals, part i: Introduction. preliminaries. the geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. *J. Symb. Comput.*, 13:255–300, 1992.
- [Sko34] Th Skolem. Ein verfahren zur behandlung gewisser exponentialer gleichungen, 8de skand. mat. Kongr. Forh. Stockholm, 1934.
- [Tar51] A. Tarski. *A Decision Method for Elementary Algebra and Geometry*. University of California Press, 1951.
- [Tiw04] Ashish Tiwari. Termination of linear programs. In *Computer-Aided Verification, CAV*, volume 3114 of *LNCS*, pages 70–82. Springer, July 2004.
- [Vah23] Mihir Vahanwala. Robust positivity problems for linear recurrence sequences. *Corr arXiv:2305.04870*, 2023. URL: <https://arxiv.org/abs/2305.04870>, arXiv:2305.04870, doi: <https://doi.org/10.48550/arXiv.2305.04870>.
- [Ver85] NK Vereshchagin. The problem of appearance of a zero in a linear recurrence sequence. *Mat. Zametki*, 38(2):609–615, 1985.

APPENDIX A.

In this Appendix, we make a few remarks regarding the representation of numbers. A complex number α is said to be algebraic if it is a root of a polynomial with integer coefficients. For an algebraic number α , its defining polynomial p_α is the unique polynomial of least degree of $\mathbb{Z}[X]$ such that the GCD of its coefficients is 1 and α is one of its roots. Given a polynomial $p \in \mathbb{Z}[X]$, we denote the length of its representation $\text{size}(p)$, its height $H(p)$ the maximum absolute value of the coefficients of p and $d(p)$ the degree of p . When the context is clear, we will only use H and d .

A separation bound provided in [Mig82] has established that for distinct roots α and β of a polynomial $p \in \mathbb{Z}[X]$, $|\alpha - \beta| > \frac{\sqrt{6}}{d^{(d+1)/2} H^{d-1}}$. This bound allows one to represent an

algebraic number α as a 4-tuple $(p, a, b, r) \in \mathbb{Z}[X] \times \mathbb{Q}^3$ where α is the only root of p at distance $< r$ from $a + ib$, and we denote $\text{size}(\alpha)$ the size of this representation, i.e., number of bits needed to write down this 4-tuple.

Further, we note that two distinct algebraic numbers α and β , are always roots of $p_\alpha p_\beta$, and we have that

$$\frac{1}{|\alpha - \beta|} = 2^{(\|\alpha\| + \|\beta\|)^{O(1)}} \quad (\text{A.1})$$

Given a polynomial $p \in \mathbb{Z}[X]$, one can compute its roots in polynomial time wrt $\text{size}(p)$ [BPR06]. Since algebraic numbers form a field, given α, β two algebraic numbers, one can always compute the representations of $\alpha + \beta, \alpha\beta, \frac{1}{\alpha}, \text{Re}(\alpha), \text{Im}(\alpha)$ in polynomial time wrt $\text{size}(\alpha) + \text{size}(\beta)$ [BPR06, Coh93].