Robust Positivity Problems for Linear Recurrence Sequences

The Frontiers of Decidability for Explicitly Given Neighbourhoods

Mihir Vahanwala
Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany

Abstract

Linear Recurrence Sequences (LRS) are a fundamental mathematical primitive for a plethora of applications such as the verification of probabilistic systems, model checking, computational biology, and economics. Positivity (are all terms of the given LRS non-negative?) and Ultimate Positivity (are all but finitely many terms of the given LRS non-negative?) are important open number-theoretic decision problems. Recently, the robust versions of these problems, that ask whether the LRS is (Ultimately) Positive despite small perturbations to its initialisation, have gained attention as a means to model the imprecision that arises in practical settings. However, the state of the art is ill-equipped to reason about imprecision when its extent is explicitly specified. In this paper, we consider Robust Positivity and Ultimate Positivity problems where the neighbourhood of the initialisation, expressed in a natural and general format, is also part of the input. We contribute by proving sharp decidability results: decision procedures at orders our techniques are unable to handle for general LRS would entail significant number-theoretic breakthroughs.

2012 ACM Subject Classification Theory of computation → Logic and verification

Keywords and phrases Dynamical Systems, Verification, Robustness, Linear Recurrence Sequences, Positivity, Ultimate Positivity

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2023.17

Funding The author is partially funded by DFG grant 389792660 as part of TRR 248 (see `https://perspicuous-computing.science`).

Acknowledgements This work is an extension of a recent collaboration with S Akshay, Hugo Bazille, and Blaise Genest. I am especially grateful to Valérie Berthé for her help in number-theoretic aspects of the paper. Our discussions took place during the Bellairs Dynamical Systems workshop organised by my advisor, Joël Ouaknine. I thank Bellairs Research Institute of McGill University, Barbados, for the hospitality and peaceful environment conducive to scientific discussion. Finally, I thank all the anonymous reviewers for their astute observations and constructive suggestions.

1 Introduction

A real Linear Recurrence Sequence (LRS) of order κ is an infinite sequence of real numbers $(u_0, u_1, u_2, \dots)$ having the following property: there exist κ real constants $a_0, \dots, a_{\kappa-1}$, with $a_0 \neq 0$ such that for all $n \geq 0$:

$$u_{n+\kappa} = a_{\kappa-1}u_{n+\kappa-1} + \cdots + a_0 u_n. \tag{1}$$

The constants $a_0, \dots, a_{\kappa-1}$ define the linear recurrence relation $\mathbf{a}$; they are also associated with the characteristic polynomial $X^\kappa - a_{\kappa-1}X^{\kappa-1} - \cdots - a_1 X - a_0$. The initial terms $u_0, \dots, u_{\kappa-1}$ are collectively denoted as the initialisation $\mathbf{c}$. An LRS is uniquely specified by $(\mathbf{a}, \mathbf{c})$. The best-known example is the Fibonacci sequence $(0, 1, 1, 2, 3, 5, 8, \dots)$, satisfying the recurrence relation $u_{n+2} = u_{n+1} + u_n$: it is named after Leonardo of Pisa, who used it to model the population growth of rabbits. LRS have been extensively studied, and found several mathematical and scientific applications since. The monograph of Everest et al. [15] is a comprehensive treatise on the mathematical aspects of Recurrence Sequences.

Important number-theoretic decision problems for Linear Recurrence Sequences include Positivity (is $u_n \geq 0$ for all n?), Ultimate Positivity (is $u_n \geq 0$ for all but finitely many n?) and the closely related Skolem Problem (is $u_n = 0$ for some n?). We remark that a Positive LRS is necessarily Ultimately Positive. As detailed in [24], these problems have applications in software verification, probabilistic model checking, discrete dynamic systems, theoretical biology, and economics. Decidability has been open for decades, with breakthroughs in restricted settings: Mignotte et al. [21] and Vereshchagin [28] independently proved the Skolem Problem to be decidable up to order 4. Further progress has been reliant on spectral restrictions or number-theoretic conjectures, consult, e.g. [19, 1]. Ouaknine and Worrell [24] showed Positivity and Ultimate Positivity are decidable up to order 5 but number-theoretically hard at order 6. For *simple* LRS (those whose characteristic polynomials have no repeated roots), they showed that Positivity is decidable up to order 9 [23] and Ultimate Positivity is decidable at all orders [25]. These results were originally proven for LRS specified by *rational* recurrences and initialisations, but can be generalised to real algebraic input as well. In this paper, we focus on Positivity and Ultimate Positivity for **sequences defined by real algebraic input**.

In contrast, the *uninitialised* variants of these problems are far more tractable. Braverman [9] and Tiwari [27] consider whether *every* possible initialisation keeps the sequence Positive, and decide so in PTIME. More recently, this result has been extended to processes with choices [4]. We argue that practical applications need a middle ground: recurrence relations that arise in practice need to be contextualised by actual instances of sequences; however, considering *precise* initialisations does not account for inherently imprecise real world measurements, and the requirement of safety margins. We thus study robust variants: given a recurrence and an initialisation, do all initialisations in a neighbourhood satisfy (Ultimate) Positivity?

Related Work

In this paper, we focus on the neighbourhood-of-initialisation notion of robustness, which was first introduced in [2], and more comprehensively treated in [3]. Works with a more control-theoretic flavour include [5], which allows for rounding at every step before applying the recurrence; in the same vein, [12, 13] allow for ε-disturbances at every step of the sequence. Our notion of robustness has been considered in [2, 3, 13], however, these works primarily concern themselves with simply deciding whether there *exists* a neighbourhood around the given point that satisfies Positivity, or whether there *exists* a tolerance ε such that the sequence avoids a region despite ε-disturbances at every step. Although they do identify that robust problems are hard when the neighbourhood is given as input, in the absence of decidability results, their hardness results are not sharp.

There are, of course, broader approaches to model and reason about imprecision: [22] considers a model of computation that can take arbitrary real numbers as input, thereby allowing imprecision in both the initialisation and the recurrence. Even in this setting, the focus is on whether the decision is locally constant in *some* neighbourhood of the given instance of the Positivity Problem, as opposed to whether the decision holds for an entire *given* neighbourhood.

Our contribution

We address the gap in the robustness state of the art by exploring the frontiers of decidability when the neighbourhood is given as input. Concretely, our input consists of a linear recurrence relation **a** and a neighbourhood of initialisations centred around **c**. Our problem is to decide whether all initialisations in the **given** neighbourhood result in (Ultimately) Positive sequences.

When neighbourhoods are expressly given as input, their geometry plays a critical role in the decision procedure. The notion of neighbourhoods that we primarily focus on is based on the ℓ^2-norm. We seek to slightly generalise the study of Euclidean ε-ball neighbourhoods undertaken in [2, 3]. More specifically, we use the Mahalanobis distance to define neighbourhoods. Our parameter is the positive definite matrix $\mathbf{S}$, and the neighbourhood of $\mathbf{c}$ it specifies is the set of all points $\mathbf{c}' \in \mathbb{R}^\kappa$ such that $(\mathbf{c}' - \mathbf{c})^T \mathbf{S} (\mathbf{c}' - \mathbf{c}) \leq 1$. The size of neighbourhoods is usually parametrised by an ε: in our case, we can account for it by simply scaling $\mathbf{S}$. In the statistical context, $\mathbf{S}$ is the inverse of a covariance matrix; and thus, our formulation is a rather natural way of capturing noise and measurement errors in the input, whose components may often be correlated. Our **novelty**, to the best of our knowledge, lies in identifying a general and practical way of explicitly specifying neighbourhoods, and establishing the **first decidability results** in such a setting, albeit at low orders or subject to spectral constraints.

As first discussed in [24, Section 5], solving decision problems on Linear Recurrence Sequences in full generality is an endeavour fraught with number-theoretic hardness. Decision procedures for Positivity problems for LRS of higher order would allow number theorists to compute properties of irrational numbers that are considered inaccessible to contemporary techniques. These include the Diophantine approximation type, which intuitively describes the quality of the "best" rational approximation of a given irrational number, and the Lagrange constant, which intuitively describes how well increasingly precise rational approximations of a given irrational number converge. We justify the inability of our techniques to handle LRS of higher orders by **reducing the computation of Diophantine approximation types and Lagrange constants** to robust Positivity problems for LRS of lower orders than ever before. Prior reductions result in LRS of order 6, we demonstrate that hard instances in our setting exist at order 5. Robust problems for Euclidean ε-ball neighbourhoods were also proven hard at order 6 [2, 3]. The setting we consider generalises this specific case, but we concede that the techniques do not address the decidability status for Euclidean ε-ball neighbourhoods at order 5.

Table 1 Main results, summarised. The distinction between uniform and non-uniform refers to whether the threshold index for certifying Ultimate Positivity must be common for the entire neighbourhood.

	Decidability Proof		
Problem: S-Robust	**General**	**Simple**	**Hardness**
Positivity	order ≤ 4	order ≤ 5	Diophantine hard at order 5
Uniform Ultimate Positivity	order ≤ 4	**all orders**	Lagrange hard at order 5
Non-uniform Ultimate Positivity	order ≤ 4	order ≤ 4	[24, 3]: Lagrange hard at order 6

Structure of the paper

The exponential polynomial closed form is an invaluable tool in the study of LRS, and we devote §2 to its exposition. This equips us to introduce our Robust Positivity Problems and intuit their decidability proofs in §3. Linear Recurrences and Diophantine Approximation are intrinsically connected: number-theoretic results form the basis of decision procedures; open problems are a yardstick for hardness reductions. We survey the number theory relevant to us in §4. We then prove our decidability results in the technical §5 and §6, and present our hardness reduction in §7. We provide concluding perspective in §8. We refer the reader to Appendix A for a summary of the standard notation and prerequisites we use.

2 The exponential polynomial closed form

We begin by discussing the exponential polynomial closed form, a perspective that is routinely leveraged to study the behaviour of Linear Recurrence Sequences. Simple LRS (no repeated characteristic roots) have a surjective correspondence with the closed form

$$u_n = \sum_j w_j \rho_j^n + \sum_j (z_j \gamma_j^n + \bar{z_j} \bar{\gamma_j}^n) \tag{2}$$

where each $\rho_j \in \mathbb{R}, \gamma_j, \bar{\gamma_j} \in \mathbb{C}\backslash\mathbb{R}$ are distinct roots of the characteristic polynomial. By straightforward arithmetic on the above expression, we can see that if $(u_n)_{n\in\mathbb{N}}, (v_n)_{n\in\mathbb{N}}$ are simple LRS with sets of characteristic roots U and V respectively, then

- $r_n = u_n + v_n$ is a simple LRS, whose set of roots is $U \cup V$.
- $r_n = u_n \cdot v_n$ is a simple LRS, whose set of roots is $\{\gamma_1\gamma_2 : \gamma_1 \in U, \gamma_2 \in V\}$.

In general, one can encode a linear recurrence $\mathbf{a}$ as a $\kappa \times \kappa$ companion matrix $\mathbf{A}$, and interpret the initialisation $\mathbf{c}$ as a vector. Then, u_n is given by the first coordinate of $\mathbf{A}^n\mathbf{c}$, i.e.

$$\begin{bmatrix} u_n \\ u_{n+1} \\ \vdots \\ u_{n+\kappa-1} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots \\ a_0 & a_1 & a_2 & \dots & a_{\kappa-1} \end{bmatrix}^n \begin{bmatrix} u_0 \\ u_1 \\ \vdots \\ u_{\kappa-1} \end{bmatrix}. \tag{3}$$

Let $\mathbf{e_1}^T$ denote the row vector $\begin{bmatrix} 1 & 0 & \dots & 0 \end{bmatrix}$. We can thus write $u_n = \mathbf{e_1}^T\mathbf{A}^n\mathbf{c}$. It is now a standard fact that LRS have a surjective correspondence with the following **real exponential polynomial** closed form (every sequence given by the expression is an LRS)

$$u_n = \left(\sum_{j=1}^{k_1} \sum_{\ell=0}^{m_j-1} z_{j\ell} \rho_j^n n^\ell \right) + \left(\sum_{j=k_1+1}^{k_2} \sum_{\ell=0}^{m_j-1} (x_{j\ell} \cos n\theta_j + y_{j\ell} \sin n\theta_j) \rho_j^n n^\ell \right) \tag{4}$$

where ρ_j (alternately, $\rho_j e^{i\theta_j}$) are roots of the characteristic polynomial defined by $\mathbf{a}$, each with multiplicity m_j. There are k_1 distinct real roots, and k_2 distinct pairs of complex conjugates among the roots. Given the recurrence relation $\mathbf{a}$, the coefficients $z_{j\ell}, x_{j\ell}, y_{j\ell}$ are each linear functions of $\mathbf{c}$. Indeed, the exponential polynomial solution expresses the sequence as a linear combination of closed-form functions that satisfy the recurrence. For a choice of basis functions, let $\mathbf{q_n}$ denote the vector whose entries are these functions, evaluated at n. In the above example, the entries are of the form $\rho_j n^\ell, \rho_j n^\ell \cos n\theta_j, \rho_j n^\ell \sin n\theta_j$, and so on. The choice of $\{\mathbf{q_n}\}_{n\in\mathbb{N}}$ can differ in "phase": one can replace $\cos n\theta, \sin n\theta$ by $\cos(n\theta - \varphi), \sin(n\theta - \varphi)$ for some choice of φ. Such a choice of basis defines a generalised $\kappa \times \kappa$ Vandermonde matrix $\mathbf{V}$ [16], whose i^{th} row is $\mathbf{q_{i-1}}^T$. By construction, $\mathbf{e_1}^T\mathbf{A}^n\mathbf{V} = \mathbf{q_n}^T$. By defining $\mathbf{p} = \mathbf{V}^{-1}\mathbf{c}$, u_n can thus be equivalently expressed as the inner (dot) product $\langle \mathbf{p}, \mathbf{q_n} \rangle$. The coefficients $z_{j\ell}, x_{j\ell}, y_{j\ell}$ are entries of this very $\mathbf{p}$.

Roots such that $|\rho_j|$ is the largest are called **dominant**. The growth rate of a term in the above expression is governed by $\rho_j^n n^\ell$. Terms with the fastest growth are called **dominant terms**, and they drive the asymptotic behaviour of the LRS. A standard, intuitive prerequisite for Ultimate Positivity is that the leading terms in the exponential polynomial expression must include one that is real and strictly positive, otherwise their dominant contribution oscillates between positive and negative. It is formalised by applying [9, Lemma 4] to the dominant terms in expression 4 and arguing that the contribution from the remaining terms vanishes asymptotically.

▶ **Proposition 1.** *If the characteristic polynomial has no real dominant root of maximum multiplicity, then in any full-dimensional neighbourhood of initialisations, there exists an initialisation, such that the sequence has infinitely many positive terms, and infinitely many negative terms.*

Henceforth, we assume that the characteristic polynomial has a real positive dominant root of maximum multiplicity, for otherwise the answer to Ultimate Positivity is trivially NO. In this paper, we shall always work with algebraic numbers[1], and thus factorising the characteristic polynomial and checking this condition is a standard procedure.

We define $\langle \mathbf{p}, \mathbf{q_n} \rangle_{dom}$ to be the normalised contribution of the dominant terms in the exponential polynomial solution. That is, if the dominant growth rate is $\rho^n n^\ell$, we pick terms with that growth rate, and divide their contribution by $\rho^n n^\ell$. For example, if $u_n = p_1 2^n + p_2 2^n \cos(n\theta - \varphi) + p_3 2^n \sin(n\theta - \varphi) + p_4$ then

$$\langle \mathbf{p}, \mathbf{q_n} \rangle_{dom} = p_1 + p_2 \cos(n\theta - \varphi) + p_3 \sin(n\theta - \varphi).$$

We define $\mu(\mathbf{c}) = \liminf_{n \in \mathbb{N}} \langle \mathbf{p}, \mathbf{q_n} \rangle_{dom}$. Note that μ is an intrinsic property of the initialisation $\mathbf{c}$ and the sequence it generates, and hence is invariant under the choice of "phase shift" φ while defining $\{\mathbf{q_n}\}_{n \in \mathbb{N}}$. In our example, assuming θ is not a rational multiple of π, it is $p_1 - \sqrt{p_2^2 + p_3^2}$

3 Robust Positivity Problems

In this paper, we shall focus on defining and tackling Robust Positivity problems. Our input consists of a linear recurrence relation $\mathbf{a}$, an initialisation $\mathbf{c}$, and a positive definite matrix $\mathbf{S}$ that is used to define a neighbourhood around $\mathbf{c}$. **All input is real algebraic.**

▶ **Problem 1** (**S**-Robust Positivity)**.** *Decide whether for all* $\mathbf{c}'$ *such that* $(\mathbf{c}' - \mathbf{c})^T \mathbf{S} (\mathbf{c}' - \mathbf{c}) \leq 1$*, the LRS* $(\mathbf{a}, \mathbf{c}')$ *is positive.*

▶ **Problem 2** (**S**-Robust Uniform Ultimate Positivity)**.** *Decide whether there exists an* N *such that for all* $\mathbf{c}'$ *with* $(\mathbf{c}' - \mathbf{c})^T \mathbf{S} (\mathbf{c}' - \mathbf{c}) \leq 1$*, the LRS* $(\mathbf{a}, \mathbf{c}')$ *is positive from the* N^{th} *term onwards.*

We can switch the order in which N and $\mathbf{c}'$ are quantified, and query a weaker notion of Robust Ultimate Positivity:

▶ **Problem 3** (**S**-Robust Non-uniform Ultimate Positivity)**.** *Decide whether for all* $\mathbf{c}'$ *with* $(\mathbf{c}' - \mathbf{c})^T \mathbf{S} (\mathbf{c}' - \mathbf{c}) \leq 1$ *, there exists an* N *such that the LRS* $(\mathbf{a}, \mathbf{c}')$ *is positive from the* N^{th} *term onwards.*

The attentive reader might have already noticed that we depart from convention and specify neighbourhoods as *closed* balls. Although [3] does not solve the problems we consider in this paper, it makes crucial observations about the geometry: for Problems 1 and 2, there is no difference between open and closed balls. On the other hand, Problem 3 becomes considerably easier with open balls, and its decidability in this case is tackled in [3] itself.

[1] The field of algebraic numbers $\overline{\mathbb{Q}}$ is the algebraic closure of the rationals $\mathbb{Q}$. Arithmetic and polynomial factoring over $\overline{\mathbb{Q}}$ can be performed with exact precision. We use $\mathbb{A}$ to denote the field of real algebraic numbers, and refer the reader to Appendix A for an initiation to these number fields.

3.1 Uniform Variants: The foundation

In general, an arbitrary point $\mathbf{c}'$ is expressed as $\mathbf{c} + \mathbf{d}$, where $\mathbf{d} \in \mathcal{P}$, a full-dimensional neighbourhood symmetric about the origin. Observe equation 3. The n^{th} term of the LRS is non-negative throughout the neighbourhood if and only if for all $d \in \mathcal{P}$, $\mathbf{e_1}^T\mathbf{A}^n(\mathbf{c}+\mathbf{d}) \geq 0$. We can use the symmetry of $\mathcal{P}$ about the origin to rewrite the above as

$$\mathbf{e_1}^T\mathbf{A}^n\mathbf{c} \geq \max_{\mathbf{d}\in\mathcal{P}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d} \geq 0. \tag{5}$$

As a simple illustration, assume that the neighbourhood is defined by a polytope rather than a positive definite matrix. This situation arises, for instance, when the metric is based on the ℓ^1- or ℓ^∞-norm, as opposed to the ℓ^2-norm. In this simple example, $\mathcal{P}$ is a polytope, hence $\mathbf{e_1}^T\mathbf{A}^n\mathbf{d}$ is maximised at one of the finitely many corners $\{\mathbf{d_1}, \ldots, \mathbf{d_k}\}$. Thus, Robust (Uniform Ultimate) Positivity is decided by using the state of the art [24] to check the (Ultimate) Positivity of each of the LRS $(\mathbf{a}, \mathbf{c}+\mathbf{d_i})$. The geometry of our setting is not simple enough to allow such a straightforward approach. **The overview of our approach to Problems 1 and 2 is as follows.**

1. Decide (constructively for Problem 1) whether there exists an N_1 such that $\mathbf{e_1}^T\mathbf{A}^n\mathbf{c} \geq 0$ for all $n > N_1$. If N_1 is explicitly required, the state of the art is able to tackle LRS of order ≤ 5 [24] and simple LRS satisfying spectral conditions that always hold at orders up to 9 [23]. In the non-constructive case, it can further handle all simple LRS [25].
2. Use linear-algebraic arguments to define a real algebraic LRS $(v_n)_{n=0}^{\infty}$, such that $v_n \geq 0$ if and only if $|\mathbf{e_1}^T\mathbf{A}^n\mathbf{c}| \geq \max_{\mathbf{d}\in\mathcal{P}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d}$.
3. Decide (constructively for Problem 1) whether there exists N_2 such that $v_n \geq 0$ for all $n > N_2$. Positivity throughout the neighbourhood is thus guaranteed beyond step $N = \max(N_1, N_2)$. If either N_1 or N_2 does not exist, then Robust Ultimate Positivity, and hence Robust Positivity, does not hold.
4. **Only for Problem 1:** Explicitly check inequality 5 for $n \leq N$.

Our novelty lies in Step 2 and identifying when Step 3 can be implemented. We now discuss how we perform Steps 2 and 4 when $\mathcal{P} = \mathcal{B}_\mathbf{S}$, a neighbourhood of vectors $\mathbf{d}$ such that $\mathbf{d}^T\mathbf{S}\mathbf{d} \leq 1$. The defining parameter $\mathbf{S}$ is a real algebraic positive definite matrix. We note that since $\mathbf{S}$ is positive definite, it can be factored as $\mathbf{G}^T\mathbf{G}$, where $\mathbf{G}$ is a real algebraic invertible matrix. We denote $\mathbf{G}\mathbf{d} = \mathbf{f}$. We argue that $\mathbf{G}^{-1}$ bijectively maps the Euclidean unit ball $\mathcal{B}$ to $\mathcal{B}_\mathbf{S}$. The bijection is clear from the invertibility of the matrix. Suppose $\mathbf{d} = \mathbf{G}^{-1}\mathbf{f}$, where $\mathbf{f} \in \mathcal{B}$, i.e. $\mathbf{f}^T\mathbf{f} \leq 1$. Then $\mathbf{d}^T\mathbf{S}\mathbf{d} = \mathbf{d}^T\mathbf{G}^T\mathbf{G}\mathbf{d} = \mathbf{f}^T\mathbf{f} \leq 1$. Hence,

$$\max_{\mathbf{d}\in\mathcal{B}_\mathbf{S}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d} = \max_{\mathbf{f}\in\mathcal{B}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{G}^{-1}\mathbf{f}. \tag{6}$$

$\mathcal{B}$ is a convex set; thus a linear function will necessarily be maximised at its boundary, i.e. when $||\mathbf{f}|| = 1$. Given $\mathbf{h}$, the linear function mapping $\mathbf{f}$ to $\mathbf{h}^T\mathbf{f}$ is maximised over the unit Euclidean ball when $\mathbf{f}$ is aligned along $\mathbf{h}$; the maximum value is $||\mathbf{h}||$. We can thus perform Step 4 because

$$\max_{\mathbf{d}\in\mathcal{B}_\mathbf{S}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d} = \left|\left|\left(\mathbf{e_1}^T\mathbf{A}^n\mathbf{G}^{-1}\right)^T\right|\right|. \tag{7}$$

For Step 2, we need a necessary and sufficient condition for $|\mathbf{e_1}^T\mathbf{A}^n\mathbf{c}| \geq \max_{\mathbf{d}\in\mathcal{B}_\mathbf{S}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d}$, in terms of the positivity of an LRS at iterate n. We consider the first part of inequality 5, substitute equation 7 in, square both sides, and transfer all terms to the left:

$$(\mathbf{e_1}^T\mathbf{A}^n\mathbf{c})^2 - (\mathbf{e_1}^T\mathbf{A}^n\mathbf{g_1})^2 - \cdots - (\mathbf{e_1}^T\mathbf{A}^n\mathbf{g}_\kappa)^2 \geq 0. \tag{8}$$

Crucially, $\mathbf{g_1}, \ldots, \mathbf{g}_\kappa$ are the linearly independent columns of the invertible $\mathbf{G}^{-1}$. Only Step 3 remains to be addressed: we must (constructively) decide whether there exists N_2 such that the previous inequality holds for all $n > N_2$. In §5, we give the technical details, thus proving our first main decidability result.

▶ **Theorem 2** (First Main Decidability Result). *Problem 2 is decidable for simple LRS. Problem 1 is decidable for simple LRS up to order 5. Problems 1 and 2 are decidable for general LRS up to order 4.*

3.2 The non-uniform variant: An overview

As discussed at length in [2, 3], $\mu(\mathbf{c}') = \liminf_{n\in\mathbb{N}} \langle \mathbf{p}', \mathbf{q_n} \rangle_{dom} \geq 0$ is necessary for the Ultimate Positivity of $\mathbf{c}'$; $\mu(\mathbf{c}') > 0$ is sufficient.

Our strategy for Problem 3 is as follows.

1. Use the First Order Theory of the Reals to check that $\mu(\mathbf{c}') \geq 0$ for all $\mathbf{c}'$ in the given neighbourhood, and detect the critical boundary cases when $\mu(\mathbf{c}') = 0$.
2. Exploit the low dimensionality to decide the critical boundary cases when $\mu(\mathbf{c}') = 0$.

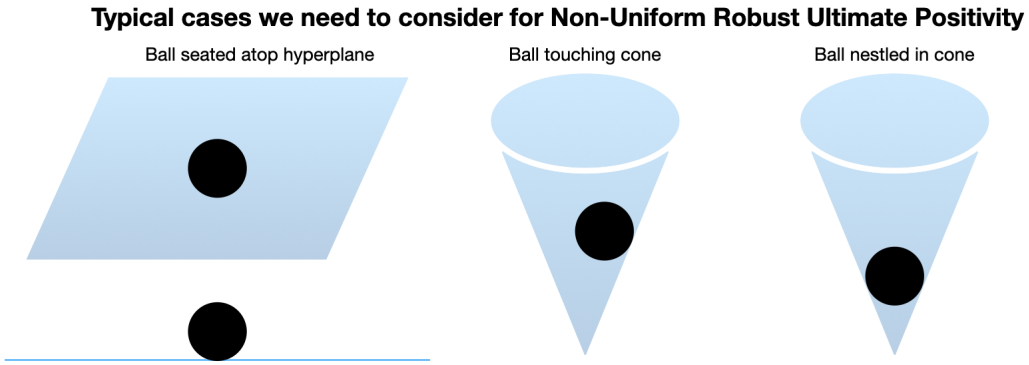

Figure 1 Visual intuition. The region $\mu \geq 0$ is defined by the intersection of halfspaces. The orientation of the neighbourhood relative to this region is deduced with the First Order Theory of the Reals. When there are finitely many halfspaces, the critical case is marked by the ball being tangent to the separating hyperplane(s) at finitely many discrete points. In low dimensions, Ultimate Positivity can be decided for these boundary cases using existing techniques. When there are infinitely many halfspaces, they carve out a region that resembles a cone. The neighbourhood can either touch the cone as before, or be nestled in it, having a continuous, connected region of tangency. In the latter case, Robust Ultimate Positivity can be handled with number-theoretic arguments in the low-dimensional setting.

We adopt this strategy (see Figure 1) and prove our second decidability result in §6.

▶ **Theorem 3** (Second Decidability Result). *Problem 3 is decidable up to order 4.*

4 Diophantine Approximation

We justify the inability of our techniques to generalise to LRS of higher order by establishing a connection to a number-theoretic hurdle: that of Diophantine Approximation. Diophantine Approximation is a vast and active number-theoretic field of research, one of whose concerns is the approximation of reals by rational numbers. A key tool in this regard is the continued fraction expansion $[a_0; a_1, a_2, \ldots]$ of an irrational t:

$$t = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\cdots}}}$$

where $a_0, a_1, a_2, \cdots \in \mathbb{N}$. Truncating this expansion at progressively greater depths yields a series of increasingly accurate approximations. The quality of the rational approximation depends not only on its accuracy but also on the size of the denominator. As discussed in the Introduction, evaluating the quality of the approximation, or that of the convergence, seems inaccessible to contemporary number theory.

The above intuition about the quality of the approximation is captured in the following definition of $L(t)$, the (homogenous) Diophantine approximation type:

$$L(t) = \inf \left\{ c \in \mathbb{R} : \left| t - \frac{p}{q} \right| < \frac{c}{q^2} \text{ for some } p, q \in \mathbb{Z} \right\}. \tag{9}$$

Similarly, the quality of the convergence is formalised by defining $L_\infty(t)$, the (homogenous) Lagrange constant:

$$L_\infty(t) = \inf \left\{ c \in \mathbb{R} : \left| t - \frac{p}{q} \right| < \frac{c}{q^2} \text{ for infinitely many } p, q \in \mathbb{Z} \right\}. \tag{10}$$

For technical purposes, we use an equivalent definition that relates to the continued fraction perspective, and allows for a slight generalisation. We follow Lagarias and Shallit's terminology [18] and use $[x]$ to denote the shortest distance from x to an integer; for $b \in \mathbb{R}$, $[x]_b$ denotes the shortest distance from x to an integer multiple of b. It is easy to observe the property $[x]_b = b[x/b]$.

▶ **Definition 4** (Diophantine Approximation Type). *The homogenous Diophantine approximation type $L(t)$ is defined to be $\inf_{n \in \mathbb{N}_{>0}} n[nt]$. The inhomogeneous Diophantine approximation type $L(t, s)$ is defined to be $\inf_{n \in \mathbb{N}_{>0}} n[nt - s]$, $s \notin \mathbb{Z} + t\mathbb{Z}$.*

▶ **Definition 5** (Lagrange constant). *The homogenous Lagrange constant $L_\infty(t)$ is defined to be $\liminf_{n \in \mathbb{N}} n[nt]$. The inhomogeneous Lagrange constant $L_\infty(t, s)$ is defined to be $\liminf_{n \in \mathbb{N}} n[nt - s]$, $s \notin \mathbb{Z} + t\mathbb{Z}$.*

From the definitions, it is clear that $0 \leq L(t) \leq L_\infty(t)$, and also $0 \leq L(t, s) \leq L_\infty(t, s)$. Due to the work of Khintchine [17], it is known that these constants lie between 0 and $1/\sqrt{5}$. In our setting, the irrational t comes from the argument θ of the characteristic root $\rho e^{i\theta}$ of the LRS: $t = \theta/2\pi$. The following observation that arises from the above fact and properties of the cosine function is pivotal to our novel low-dimensional decidability result for Uniform Robustness.

▶ **Lemma 6.** *For all θ that are not rational multiples of 2π, and all φ, there exist infinitely many $n \in \mathbb{N}$ such that $1 - \cos(n\theta - \varphi) \leq \frac{1}{2} [n\theta - \varphi]_{2\pi}^2 = 2\pi^2 [nt - s]^2 \leq \frac{2\pi^2}{5n^2}$.*

The properties of the LRS are driven by whether the characteristic root $\rho e^{i\theta}$ is a scaled root of unity, i.e. θ is a rational multiple of 2π. If yes, decision procedures are often much simpler; if not, we appeal to the number theory discussed in this section. One can detect whether an algebraic characteristic root is a root of unity by brute enumeration.

▶ **Lemma 7.** *Let α be an algebraic number of degree d. Then if α is a k^{th} root of unity, $k \leq 2d^2$.*

Proof. The degree of the k^{th} root of unity is precisely $\Phi(k)$, where Φ denotes the Euler totient function. The desired inequality follows from the property that $\Phi(k) \geq \sqrt{k/2}$. ◀

We record a number-theoretic fact which describes the density of the integer multiples of an irrational x modulo 1 in the unit interval: its proof relies on continued fraction expansions and the Ostrowski numeration system [8, 7], and is deferred to Appendix B. This result is decisive when considering Non-Uniform Robustness.

▶ **Lemma 8.** *For every irrational number x, strictly decreasing real positive function ψ, and interval $\mathcal{I} = [a, b] \subset [0, 1]$, $a \neq b$, there exists $y_0 \in \mathcal{I}$ such that $[nx - y_0] < \psi(n)$ for infinitely many even n, and $y_1 \in \mathcal{I}$ such that $[nx - y_1] < \psi(n)$ for infinitely many odd n.*

The familiar density theorem is an immediate corollary of the above powerful result. Indeed, we can consider an interval of length $\varepsilon/2$, and take $\psi(n) = \varepsilon/2$.

▶ **Lemma 9.** *Let x be irrational, and $y \in [0, 1)$. For every $\varepsilon > 0$, there exist infinitely many even n, and infinitely many odd n such that $[nx - y] < \varepsilon$.*

The following application of the density theorem is central to the computation of the discrete $\mu(\mathbf{c}) = \liminf_{n\in\mathbb{N}} \langle \mathbf{p}, \mathbf{q_n} \rangle_{dom}$.

▶ **Lemma 10.** *Suppose θ is not a rational multiple of 2π. Let h_1, h_2 be real functions such that $h_1(t)$ is continuous with period 2π. Then*
$\liminf_n \left(h_1(n\theta) + h_2(-1^n)\right) = \min_{t\in[0,2\pi], b\in\{-1,1\}} \left(h_1(t) + h_2(b)\right)$.

We note that despite the observations and results mentioned in the preceding discussion, the Diophantine approximation type and Lagrange constant of most transcendental numbers are unknown. For instance, computing $L_\infty(\pi)$ is a longstanding and mathematically interesting open problem. We refer the reader to [24, Section 5] for a cursory survey of the history of relevant developments in the field of Diophantine approximation. This source reduces the computation of the constants $L(t)$ and $L_\infty(t)$ for transcendental numbers such as $t = \arcsin(3/5)/2\pi$ to the non-robust variants of Positivity problems for LRS of order 6. In §7, we prove analogous hardness results for robust Positivity problems of order 5. To that end, we define a similar classes of transcendental numbers relevant to our reduction. Let

$$\mathcal{A}_{alg} = \{p + iq \in \mathbb{C} \mid p, q \in \mathbb{A},\ p^2 + q^2 = 1, \forall n.\ (p + iq)^n \neq 1\}, \tag{11}$$

$$\mathcal{A}_{rat} = \{p + iq \in \mathbb{C} \mid p, q \in \mathbb{Q},\ p^2 + q^2 = 1, \forall n.\ (p + iq)^n \neq 1\}. \tag{12}$$

i.e., the set $\mathcal{A}_{alg}$ (resp. $\mathcal{A}_{rat}$) consists of algebraic (resp. rational) numbers on the unit circle in $\mathbb{C}$, none of which are roots of unity. In particular, writing $p + qi = e^{i2\pi\theta}$, we have that $\theta \notin \mathbb{Q}$. We denote:

$$\mathcal{T}_{alg} = \left\{t \in (-1/2, 1/2] \mid e^{2\pi it} \in \mathcal{A}_{alg}\right\};\ \mathcal{T}_{rat} = \left\{t \in (-1/2, 1/2] \mid e^{2\pi it} \in \mathcal{A}_{rat}\right\}. \tag{13}$$

The sets $\mathcal{T}_{alg}, \mathcal{T}_{rat}$ are dense in $(-\frac{1}{2}, \frac{1}{2}]$. We represent $t \in \mathcal{T}_{alg}$ (resp. $\mathcal{T}_{rat}$) by the real algebraic (resp. rational) numbers $(\cos 2\pi t, \sin 2\pi t)$. In general, we don't have a method to compute or approximate to arbitrary precision $L(t), L(t, s), L_\infty(t), L_\infty(t, s)$ given $s, t \in \mathcal{T}_{alg}$, or even $s, t \in \mathcal{T}_{rat}$.

▶ **Definition 11** (Number-theoretic hardness). *Let $\mathcal{T}$ denote one of $\mathcal{T}_{alg}, \mathcal{T}_{rat}$ as defined above. A decision problem is said to be $\mathcal{T}$-Diophantine hard (resp. $\mathcal{T}$-Lagrange hard), if its decidability entails that given any $s,t \in \mathcal{T}$ and rational $\varepsilon > 0$, one can compute rational ℓ_1, ℓ_2 such that: (a) $|\ell_1 - L(t)| < \varepsilon$ (resp. $|\ell_1 - L_\infty(t)| < \varepsilon$); (b) $|\ell_2 - L(t,s)| < \varepsilon$ (resp. $|\ell_2 - L_\infty(t,s)| < \varepsilon$).*

▶ **Theorem 12** (Main Hardness Result). *Problem 1 (resp. Problem 2) is $\mathcal{T}_{alg}$-Diophantine hard (resp. $\mathcal{T}_{alg}$-Lagrange hard) at order 5. For rational input, Problem 1 (resp. Problem 2) is $\mathcal{T}_{rat}$-Diophantine hard (resp. $\mathcal{T}_{rat}$-Lagrange hard) at order 5.*

As noted in [3], in view of the Lagrange hardness (Definition 11) of Ultimate Positivity at order 6 [24], Problem 3, which asks whether the given neighbourhood consists entirely of initialisations that produce an Ultimately Positive sequence, is also Lagrange hard at order 6. The idea is to use the existing reduction from the computation of Lagrange constants to Ultimate Positivity, and extend it to the robust variant: one simply constructs a neighbourhood of initialisations that has the hard instance of Ultimate Positivity on its surface, but otherwise lies entirely in the region where Ultimate Positivity is guaranteed.

▶ **Theorem 13.** *Problem 3 is $\mathcal{T}_{alg}$-Lagrange hard at order 6.*

5 Decidability of Uniform Robustness

In this section, we prove Theorem 2 by showing that we can implement Step 3 of the overview in §3.1: (constructively) decide whether there exists N_2 such that for all $n > N_2$,

$$v_n := (\mathbf{e_1}^T \mathbf{A}^n \mathbf{c})^2 - (\mathbf{e_1}^T \mathbf{A}^n \mathbf{g_1})^2 - \cdots - (\mathbf{e_1}^T \mathbf{A}^n \mathbf{g_4})^2 \geq 0. \tag{14}$$

▶ **Theorem 14** (First Main Decidability Result, restated). *Problem 2 (Robust Uniform Ultimate Positivity) is decidable for simple LRS. Problem 1 (Robust Positivity) is decidable for simple LRS up to order 5. Problems 1 and 2 are decidable for general LRS up to order 4.*

5.1 Simple LRS

We begin by treating simple LRS. The goal is to show that the current state of the art is equipped to handle instances relevant to this setting. Recall the discussion on the point-wise sums of products of simple LRS, surrounding equation 2. If the original LRS is simple, then inequality 14 is also an instance of Ultimate Positivity for simple LRS; indeed, its input can be seen to be real algebraic. In case we are only interested in Robust Ultimate Positivity, the non-constructive decision procedure [25] suffices, because it completely solves Ultimate Positivity for simple LRS.

As a corollary of the proof of the decidability of Positivity of simple LRS up to order 9 [23], Ultimate Positivity for simple LRS is *constructively* decidable if one of the following holds: **(a)** all characteristic roots have the same modulus; **(b)** there are at most three pairs of complex conjugates among the dominant (maximal modulus) characteristic roots.

We argue that for the original simple LRS $(u_n)_n$, 5 is the highest order that guarantees that at least one of the conditions holds for the resulting simple LRS $(v_n)_n$ in inequality 14. For this, we recall the property discussed after equation 2: if U is the set of characteristic roots of $(u_n)_n$, then the set of characteristic roots of $v_n = u_n^2$ is $V = \{\lambda_1\lambda_2 : \lambda_1, \lambda_2 \in U\}$. By Proposition 1, U contains a real positive dominant root ρ. It is clear that the dominant roots of V result from, and only from multiplying together pairs of dominant roots from U. If U does not have non-real dominant roots, neither does V. If $\lambda, \bar{\lambda} \in U$ are dominant non-real roots, then $\lambda\bar{\lambda} = \rho^2 \in \mathbb{R}$. If $U = \{\rho, \lambda_1, \lambda_2, \bar{\lambda_1}, \bar{\lambda_2}\}$, all dominant, then all roots of V are

dominant, and condition **(a)** is met. The only remaining case is that U has one pair of complex conjugates among its dominant roots: the scenario that results in most dominant roots in V is $U_{dom} = \{\rho, -\rho, \lambda, \bar{\lambda}\}$. Then, the dominant roots in V are $\{\rho^2, -\rho^2, \lambda^2, \bar{\lambda}^2, \pm\rho\lambda, \pm\rho\bar{\lambda}\}$: three conjugate pairs, and condition **(b)** is met. Finally, we record that order 5 is maximal: consider $U = \{\rho, \lambda_1, \lambda_2, \bar{\lambda_1}, \bar{\lambda_2}, \alpha\}$ with α non-dominant. Then V has five pairs of complex conjugates among its dominant roots, along with the presence of non-dominant roots.

5.2 Non-simple LRS

We treat order 4 LRS: our techniques naturally apply to lower orders too. We make extensive use of the real exponential polynomial closed form 4 and the surrounding discussion. The key lies in observing that in critical inequality 14, each squared term satisfies the same recurrence, and hence we can express it as

$$\langle \mathbf{p}, \mathbf{q_n} \rangle^2 - \langle \mathbf{b_1}, \mathbf{q_n} \rangle^2 - \cdots - \langle \mathbf{b_4}, \mathbf{q_n} \rangle^2 \geq 0 \quad \Leftrightarrow \quad \langle \mathbf{x}, \mathbf{r_n} \rangle \geq 0. \tag{15}$$

We choose $\{\mathbf{q_n}\}_{n \in \mathbb{N}}$ judiciously, expand the squares, use trigonometric identities, and collect the terms in $\mathbf{r_n}$ to obtain a new LRS. If all the characteristic roots of the original LRS are real, then $\mathbf{q_n}$ is free of trigonometric terms, and hence so is $\mathbf{r_n}$. Thus $\langle \mathbf{x}, \mathbf{r_n} \rangle$ is also an LRS with all real characteristic roots, and constructively deciding the existence of N_2 is easily done through elementary growth arguments. We shall thus assume the presence of a pair of complex conjugates among the characteristic roots. As discussed through Proposition 1, any decision regarding Ultimate Positivity is NO in the absence of a real positive dominant root. At order 4, this means that there is **exactly one pair of complex conjugates** among the roots. We further assume, without loss of generality, that **the real positive dominant root is unity**. We shall also assume **non-degeneracy**, i.e. the ratio of any pair of distinct roots of the characteristic polynomial is not a root of unity. This can be detected, courtesy Lemma 7. In our restricted setting, degeneracy can arise because: (a) -1 is a characteristic root; (b) a characteristic root is of the form $\rho e^{2\pi i \cdot \frac{\ell}{k}}$, i.e. a scaled k^{th} root of unity. In this case, any LRS $\langle \mathbf{v}, \mathbf{q_n} \rangle$ with roots $\{1, \alpha, \rho e^{\pm 2\pi i \cdot \frac{\ell}{k}}\}$ can be decomposed as the interleaving of $2k$ real LRS $\{u_n^{(i)}\}_{i=0}^{2k-1}$ where $u_n^{(i)} = \langle \mathbf{v}, \mathbf{q_{2nk+i}} \rangle$. Each of these LRS has characteristic roots $\{1, \rho^{2k}\} \cup \{\alpha^{2k}\}$.

The only possibility, therefore, is that the characteristic roots are $1, 1, \gamma, \bar{\gamma}$. Let $0 < |\gamma| = \rho \leq 1$, where $\gamma = \rho e^{i\theta}$ is not a scaled root of unity. We take inequality 15 as the starting point for our computations. Let $\mathbf{q_n} = \begin{bmatrix} n & 1 & \rho^n \cos(n\theta - \varphi) & \rho^n \sin(n\theta - \varphi) \end{bmatrix}^T$. Let $\mathbf{u_1}^T, \ldots, \mathbf{u_4}^T$ be the rows of the **invertible** matrix $\begin{bmatrix} \mathbf{b_1} & \ldots & \mathbf{b_4} \end{bmatrix}$. The table below shows the terms and coefficients on simplifying inequality 15.

Term of $\mathbf{r_n}$	**Coefficient** in $\mathbf{x}$	**Explicitly**
n^2	z_2	$p_1^2 - \langle \mathbf{u_1}, \mathbf{u_1} \rangle$
n	z_1	$2p_1p_2 - 2\langle \mathbf{u_1}, \mathbf{u_2} \rangle$
1	z_0	$p_2^2 - \langle \mathbf{u_2}, \mathbf{u_2} \rangle$
$n\rho^n \cos(n\theta - \varphi)$	x_2	$2p_1p_3 - 2\langle \mathbf{u_1}, \mathbf{u_3} \rangle$
$n\rho^n \sin(n\theta - \varphi)$	y_2	$2p_1p_4 - 2\langle \mathbf{u_1}, \mathbf{u_4} \rangle$
$\rho^n \cos(n\theta - \varphi)$	x_1	$2p_2p_3 - 2\langle \mathbf{u_2}, \mathbf{u_3} \rangle$
$\rho^n \sin(n\theta - \varphi)$	y_1	$2p_2p_4 - 2\langle \mathbf{u_2}, \mathbf{u_4} \rangle$
ρ^{2n}	w	$\frac{1}{2}(p_3^2 + p_4^2) - \frac{1}{2}(\langle \mathbf{u_3}, \mathbf{u_3} \rangle + \langle \mathbf{u_4}, \mathbf{u_4} \rangle)$
$\rho^{2n} \cos(2n\theta - 2\varphi)$	x_0	$\frac{1}{2}(p_3^2 - p_4^2) - \frac{1}{2}(\langle \mathbf{u_3}, \mathbf{u_3} \rangle - \langle \mathbf{u_4}, \mathbf{u_4} \rangle)$
$\rho^{2n} \sin(2n\theta - 2\varphi)$	y_0	$p_3p_4 - \langle \mathbf{u_3}, \mathbf{u_4} \rangle$

If $\rho < 1$, then the dominant growth rate for the problem to be non-trivial is $n^2, n, 1$, or ρ^{2n}: indeed, if the growth rate is ρ^n or $n\rho^n$, the leading terms would all be trigonometric, and the sign of the sequence oscillates. The former cases can be solved with straightforward growth arguments, while the last case results in an order 3 LRS that can easily be dealt with [23, 24]. We thus assume $\rho = 1$. Again, if $z_2 \neq 0$, then decidability is trivial because the dominant growth rate of n^2 is dictated by a single term; hence we assume $z_2 = 0$. In this case, there are two groups of terms, based on growth rate: one with n, the other with 1. To study these groups, we define

$$f(t) = z_1 + x_2 \cos(t - \varphi) + y_2 \sin(t - \varphi), \tag{16}$$

$$g(t) = z_0 + w + x_1 \cos(t - \varphi) + y_1 \sin(t - \varphi) + x_0 \cos(2t - 2\varphi) + y_0 \sin(2t - 2\varphi). \tag{17}$$

Since θ is not a rational multiple of 2π, $\{n\theta \bmod 2\pi\}$ is dense in $[0, 2\pi]$, and we invoke Lemma 10 to deduce

$$\liminf_{n\in\mathbb{N}} f(n\theta) = \min_{t\in[0,2\pi]} f(t) = z_1 - \sqrt{x_2^2 + y_2^2} = \mu. \tag{18}$$

If $\mu < 0$, then the critical inequality $nf(n\theta) + g(n\theta) \geq 0$ will be violated infinitely often. If $\mu > 0$, we can compute an N_2 beyond which it is guaranteed to be satisfied. We thus concern ourselves with the case where $\mu = 0$. Recall the discussion around μ when its concept was first defined after Proposition 1: it is an intrinsic property of the problem itself, and invariant under the "phase" φ chosen in the basis of solutions. We thus assume that φ is chosen in such a way that the minimum is attained at φ, i.e. $f(\varphi) = 0$. This choice can be made by applying the trigonometric identity $\cos(a - b) = \cos a \cos b + \sin a \sin b$ to $f(t)$. This means that $y_2 = 0$, and we choose $-z_1 = x_2 < 0$.

Now, if $g(\varphi) > 0$, we compute a positive lower bound on $f(t)$ for t such that $g(t) < 0$. This then results in an N_2 beyond which $nf(n\theta) + g(n\theta) \geq 0$ is guaranteed. If $g(\varphi) < 0$, then the inequality has infinitely many violations. This is due to Lemma 6, which asserts that there are infinitely many n for which $f(n\theta) \leq 2\pi^2 z_1/5n^2$. These n are necessarily such that $n\theta$ is close to φ, and the negativity of $g(n\theta)$ is thus decisive.

The final case that remains is $g(\varphi) = 0$. We argue that remarkably, it does not arise at all!

▶ **Lemma 15.** *If $z_2 = \mu = 0$, it cannot be the case that $f(\varphi) = g(\varphi) = 0$.*

Proof. Suppose, for the sake of contradiction, the scenario actually occurs. This means that $z_2 = z_1 + x_2 = z_0 + w + x_1 + x_0 = 0$. From the table, these respectively imply

$$\begin{aligned} p_1^2 &= \langle \mathbf{u_1}, \mathbf{u_1} \rangle, \\ p_1(p_2 + p_3) &= \langle \mathbf{u_1}, \mathbf{u_2} + \mathbf{u_3} \rangle, \\ (p_2 + p_3)^2 &= \langle \mathbf{u_2} + \mathbf{u_3}, \mathbf{u_2} + \mathbf{u_3} \rangle. \end{aligned}$$

This implies that $|\langle \mathbf{u_1}, \mathbf{u_2} + \mathbf{u_3} \rangle| = ||\mathbf{u_1}|| \cdot ||\mathbf{u_2} + \mathbf{u_3}||$, i.e. $\mathbf{u_1}$ is a scaled multiple of $\mathbf{u_2} + \mathbf{u_3}$. This contradicts the fact that the rows of the invertible $\begin{bmatrix} \mathbf{b_1} & \dots & \mathbf{b_4} \end{bmatrix}$ are linearly independent, and we are done. ◀

6 Non-uniform Robustness: Decidability at order four

In this section, we prove Theorem 3. The techniques naturally apply to lower orders, and we omit their explicit treatment. Recall the critical condition from our overview in §3.2:

$$\mu(\mathbf{c}') = \liminf_{n\in\mathbb{N}} \langle \mathbf{p}', \mathbf{q_n} \rangle_{dom} \geq 0 \tag{19}$$

for all $\mathbf{c}'$ in the neighbourhood is necessary for the decision to be YES; the inequality holding strictly is sufficient. Critical cases arise when the surface of the neighbourhood touches the region where $\mu = 0$, and the non-dominant terms, if any, can potentially have a negative contribution. We demonstrate that these can be detected and dealt with.

Since Proposition 1 guarantees the existence of a real positive dominant term, $\langle \mathbf{p}, \mathbf{q_n} \rangle_{dom}$ can only be of one of the following forms: **(I)** z; **(II)** $z + w(-1)^n$; **(III)** $z + x\cos n\theta + y \sin n\theta$; **(IV)** $z + x\cos n\theta + y\sin n\theta + w(-1)^n$, where x, y, z, w are linear in the initialisation $\mathbf{c}$. Cases (I), (II), and (III), (IV) where θ is a rational multiple of 2π (detected with Lemma 7) are the easiest. The region $\mu \geq 0$ is carved out by *finitely* many halfspaces, defined by separating hyperplanes of the form $z + bw + c_0 x + s_0 y = 0$. By elementary linear algebra and co-ordinate geometry (e.g. by working in a basis where the neighbourhood is a perfect hypersphere), one can determine whether $\mu > 0$ for the entire neighbourhood, or whether $\mu < 0$ for some points in the neighbourhood, or whether the neighbourhood touches a hyperplane. Each hyperplane has at most one point of tangency, whose algebraic coordinates can be solved for. These critical points are low-dimensional instances of Ultimate Positivity, and can be decided with the state of the art [25].

We therefore assume that θ is not a rational multiple of 2π, and we are in Case (III) or (IV). We apply Lemma 10, we get that

$$\mu(\mathbf{c}) = \liminf_{n\in\mathbb{N}} \langle \mathbf{p}, \mathbf{q_n} \rangle_{dom} = \min_{t\in\mathbb{R}, b\in\{\pm 1\}} z + x\cos t + y \sin t + wb = z - \sqrt{x^2 + y^2} - |w|. \quad (20)$$

If we are in Case (IV), there are no non-dominant roots, and $\mu \geq 0$ throughout the neighbourhood is necessary as well as sufficient for the decision to be YES. This is an algebraic condition, and can be checked using the First Order Theory of the Reals.[2]

Case (III) remains. $\langle \mathbf{q_n}, \mathbf{p} \rangle = z + x\cos n\theta + y \sin n\theta + w\alpha^n$, where $0 < |\alpha| < 1$. As discussed, we can use the First Order Theory of the Reals to check the sufficient $\mu > 0$, and the necessary $\mu \geq 0$ throughout the neighbourhood. We consider the scenario where the necessity check succeeds, but the sufficiency check fails. The decision can be NO only if there are points on the surface of the neighbourhood where $\mu = 0$, and the non-dominant $w\alpha^n$ can make a negative contribution. We describe how these points are found and analysed. First, we observe that the region $\mu \geq 0$ is given by the cone $z - \sqrt{x^2 + y^2} \geq 0$. It can be intuited as being carved out by a continuum of hyperplanes $z + x\cos\phi + y\sin\phi = 0$. We encode the above discussion to find the critical points with the following first order formula with free variable c, which stands for $\cos\phi$

$$\chi_1(c) := \exists s \exists \mathbf{c}'.\ (\mathbf{c}' - \mathbf{c})^T \mathbf{S} (\mathbf{c}' - \mathbf{c}) = 1 \wedge z' + cx' + sy' = 0 \wedge c^2 + s^2 = 1 \wedge w' \sim 0. \quad (21)$$

In the above $\sim$ is $\neq$ if the non-dominant root $\alpha < 0$, and is $<$ if $\alpha > 0$. We can use Theorem 17 to get an equivalent quantifier free formula: this comprises purely of polynomial (in-)equalities in the free variable c. The set of c, and hence $\cos\phi$, satisfying these, consists of finitely many intervals. Of course, Ultimate Positivity is guaranteed when this set is empty: it means there are no points threatening to violate Ultimate Positivity.

[2] Given input $\mathbf{a}, \mathbf{c}, \mathbf{S}$, we can choose $\mathbf{V}$ (cf equation 4). Our formula is

$$\forall \mathbf{c}'.\ (\mathbf{c}' - \mathbf{c})^T \mathbf{S} (\mathbf{c}' - \mathbf{c}) \leq 1 \Rightarrow \Bigg(\exists z', x', y', w', r_1, r_2.\ \left(\begin{bmatrix} z' & x' & y' & w' \end{bmatrix}^T = \mathbf{V}^{-1} \mathbf{c}' \right) \wedge (r_1, r_2 \geq 0)$$
$$\wedge (z' - r_1 - r_2 \geq 0) \wedge (r_1^2 = w'^2) \wedge (r_2^2 = x'^2 + y'^2) \Bigg).$$

We first dispose of the case where all intervals consist of single points. Consider an interval $\{c_0\}$ consisting of a single point. This is illustrated by the case of the ball touching the cone in Figure 1. Due to its origins and discrete occurrence, c_0 must be a root of a polynomial obtained by quantifier elimination on χ_1, and is hence algebraic. The corresponding critical point is the point of tangency of the neighbourhood with a hyperplane with a real algebraic equation. Thus, it generates a real algebraic instance of Ultimate Positivity, which can be decided with the techniques of [25].

If, however, the set of c satisfying χ_1 consists of intervals that have more than one point, then the techniques of [25] to decide Ultimate Positivity for a single point with algebraic coordinates are no longer accessible, in fact, the decision will always be NO in our setting. This situation is illustrated by the case of the ball nestled in cone in Figure 1. Let $[\phi_1, \phi_2]$ be an interval of ϕ such that: (i) all values of c between $\cos\phi_1$ and $\cos\phi_2$ satisfy χ_1, (ii) The corresponding witnesses z' are at most z_0, and (iii) The corresponding witnesses w' have magnitude at least some fixed w_0. In order to ensure robust Ultimate Positivity, we must have for each ϕ (and corresponding point on the cone with coordinates $z'(c), x' = -cz', y' = -z'\sin\phi, w'$) in this interval, the following inequality is violated only finitely often:

$$\langle \mathbf{p}', \mathbf{q_n} \rangle = z' - z'\cos n\theta\cos\phi - z'\sin n\theta\sin\phi + w\alpha^n = z' - z'\cos(n\theta - \phi) + w\alpha^n \geq 0. \quad (22)$$

We consider an even weaker inequality (courtesy Lemma 6), which, in this context, we argue is bound to be violated infinitely often:

$$z_0[n\theta - \phi]_{2\pi}^2 \geq 2w_0\alpha^n. \quad (23)$$

The argument hinges on Lemma 8, which we restate:

▶ **Lemma 16.** *For every irrational number x, strictly decreasing real positive function ψ, and interval $\mathcal{I} = [a, b] \subset [0, 1]$, $a \neq b$, there exists $y_0 \in \mathcal{I}$ such that $[nx - y_0] < \psi(n)$ for infinitely many even n, and $y_1 \in \mathcal{I}$ such that $[nx - y_1] < \psi(n)$ for infinitely many odd n.*

Now, if $\alpha < 0$, we use Lemma 16 on the irrational $\theta/2\pi$, and the decreasing $\sqrt{\frac{w_0|\alpha|^n}{2\pi^2 z_0}}$ to argue that there exists a ϕ in the desired interval, such that the weaker inequality will be violated for infinitely many n of the appropriate parity. The case $\alpha > 0$ is even simpler, as the parity does not matter. Thus, we can return NO if we are in the case where the set of c satisfying χ_1 (equation 21) consists of intervals that contain more than a single point.

7 Uniform Robustness: Hardness at order five

We shall prove Theorem 12 in this section. That is, given $s, t \in \mathcal{T}$ as defined in equation 13, we shall give $\mathbf{a}, \mathbf{c}, \mathbf{S}$ (real algebraic or rational as appropriate) such that varying $\mathbf{S} = \mathbf{G}^T\mathbf{G}$ while invoking $\mathbf{S}$-Robust Positivity decision procedures will enable us to approximate $L(t), L(t, s)$ and $L_\infty, L_\infty(t, s)$ to arbitrary precision.

We recall $t \in \mathcal{T}$ is specified by $(\cos 2\pi t, \sin 2\pi t)$. This tuple is real algebraic for $t \in \mathcal{T}_{alg}$, and rational for $t \in \mathcal{T}_{rat}$. If we wish to approximate the homogenous constants, we specify $s = 0$ with the representation $(1, 0)$. Our linear recurrence relation $\mathbf{a}$ is such that the roots of the characteristic polynomial are $1, 1, 1, e^{2\pi it}, e^{-2\pi it}$, i.e. the characteristic polynomial is $(X-1)^3(X^2 - 2X\cos 2\pi t + 1)$.

Recalling the discussion surrounding equation 4, $u_n = \mathbf{e_1}^T\mathbf{A}^n\mathbf{c} = \langle \mathbf{p}, \mathbf{q_n} \rangle$. For the problem instance we create in our reduction, we choose rational $\mathbf{p} = \begin{bmatrix} r & 0 & 1 + \frac{r}{2} & -1 & 0 \end{bmatrix}^T$, where r is a parameter we use to tune our guess for $L(t, s)$ and $L_\infty(t, s)$; we choose our basis of closed

form solutions such that $\mathbf{q_n}^T = \mathbf{e_1}^T\mathbf{A}^n\mathbf{V} = \begin{bmatrix} n^2 & n & 1 & \cos(2\pi(nt-s)) & \sin(2\pi(nt-s)) \end{bmatrix}$. We choose $\mathbf{S} = \mathbf{G}^T\mathbf{G}/r^2$, where $\mathbf{G} = \mathbf{V}^{-1}$. Observe that depending on whether s, t are in $\mathcal{T}_{alg}$ or $\mathcal{T}_{rat}$, the input $\mathbf{a}, \mathbf{c}, \mathbf{S}$ is constructed to be real algebraic or rational, as desired. Our critical inequality is

$$\langle \mathbf{p}, \mathbf{q_n} \rangle = \mathbf{e_1}^T\mathbf{A}^n\mathbf{c} \geq \max_{d \in \mathcal{B}_\mathbf{S}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d} = ||r(\mathbf{e_1}^T\mathbf{A}^n\mathbf{G}^{-1})^T|| = r||\mathbf{q_n}||. \tag{24}$$

Let $\Psi(n, r)$ denote the proposition $\langle \mathbf{p}, \mathbf{q_n} \rangle \geq r||\mathbf{q_n}||$. Our reduction works by proving that for any guess $r > 0$, given $\varepsilon > 0$, we can compute an N such that for all $n \geq N$

$$\Psi(n, r) \Rightarrow n[nt - s] > \frac{(1-\varepsilon)\sqrt{7r}}{4\pi}. \tag{25}$$

$$\neg\Psi(n, r) \Rightarrow n[nt - s] < \frac{\sqrt{7r}}{(1-\varepsilon)4\pi}. \tag{26}$$

To compute $L_\infty(t, s) = \liminf_{n \in \mathbb{N}} n[nt - s]$ by increasingly precise approximations, we query Robust Uniform Ultimate Positivity: does $\Psi(n, r)$ hold for all but finitely many n? If the decision is YES, then we use property 25 to argue that for any ε, $n[nt - s]$ exceeds $\frac{(1-\varepsilon)\sqrt{7r}}{4\pi}$ for all but finitely many n, hence $L_\infty(t, s)$ must be at least $\frac{\sqrt{7r}}{4\pi}$. Conversely, if the decision is NO, we use property 26 to deduce that for any ε, $n[nt - s]$ falls short of $\frac{\sqrt{7r}}{(1-\varepsilon)4\pi}$ for infinitely many n, hence $L_\infty(t, s)$ must be at most $\frac{\sqrt{7r}}{4\pi}$.

By definition, $L(t, s) = \inf_{n \in \mathbb{N}_{>0}} n[nt - s]$. Given the guess r, precision ε, the corresponding N, and oracle access to whether $\Psi(n, r)$ holds for all $n \geq N$, it follows from properties 25 and 26 that we can detect the truth of one of the cases among $\inf_{n \geq N} n[nt - s] \geq \frac{(1-\varepsilon)\sqrt{7r}}{4\pi}$ and $\inf_{n \geq N} n[nt - s] \leq \frac{\sqrt{7r}}{(1-\varepsilon)4\pi}$: indeed, the implications allow us to identify one from the set of valid cases, which is guaranteed to be non-empty. By obtaining sufficiently precise numerical approximations of $n[nt - s]$ for each $n < N$ in the finite prefix, one has a procedure for approximating $L(t, s)$.

In order to establish the number-theoretic hardness of Robust Positivity, we now explain how we use it as an oracle to decide whether $\Psi(n, r)$ holds for all $n \geq N$. Note that as it is, our query specifies a recurrence $\mathbf{a}$, an initialisation $\mathbf{c}$, a neighbourhood defined by $\mathbf{S}$ asks for the Robust Positivity of a *suffix* of the sequence, as opposed to the entire sequence. We create a new instance with updated $\mathbf{c}'$ and $\mathbf{S}'$ to implement the shift:

$$\forall n \geq N.\ \mathbf{e_1}^T\mathbf{A}^n\mathbf{c} \geq \max_{\mathbf{d} \in \mathcal{B}_\mathbf{S}} \mathbf{e_1}^T\mathbf{A}^n\mathbf{d} \ \Leftrightarrow\ \forall n.\ \mathbf{e_1}^T\mathbf{A}^n(\mathbf{A}^N\mathbf{c}) \geq \max_{\mathbf{d} \in \mathcal{B}_\mathbf{S}} \mathbf{e_1}^T\mathbf{A}^n(\mathbf{A}^N\mathbf{d}) \tag{27}$$

$$\Leftrightarrow\ \forall n.\ \mathbf{e_1}^T\mathbf{A}^n(\mathbf{c}') \geq \max_{\mathbf{d}' \in \mathcal{B}_{\mathbf{S}'}} \mathbf{e_1}^T\mathbf{A}^n(\mathbf{d}'). \tag{28}$$

It is clear that $\mathbf{c}' = \mathbf{A}^N\mathbf{c}$ and that $\mathbf{d}' = \mathbf{A}^N\mathbf{d}$. Using the same reasoning as we did in the derivation of equation 6, we argue $\mathbf{S}' = (\mathbf{A}^{-N})^T\mathbf{S}\mathbf{A}^{-N}$. The reduction is thus complete, but for the proof of properties 25 and 26, which we defer to Appendix C.

8 Extensions and Perspective

It is interesting to investigate what kinds of decidability and hardness results hold for neighbourhoods specified using norms other than those based on the standard matrix inner product. For instance, our techniques for $\mathbf{S}$-Robust Non-uniform Ultimate Positivity hinged on the First Order Theory of the Reals and were rather agnostic to the exact shape of the neighbourhood: we can easily extend the same techniques to arbitrary *semi-algebraic*

neighbourhoods. Perhaps, other results could also be universal across a wider class of norms, and there could be a profound underlying linear-algebraic reason whose discovery would be mathematically significant.

While contributing towards a sharp and comprehensive picture of what is *decidable* about Robust Positivity for LRS, we found it remarkable that number-theoretic analyses involving Diophantine approximation, which usually show up in the context of hardness, also play a significant role in our *decidability* proofs! However, a rather conspicuous gap in our picture is the status of **S**-Robust Non-uniform Ultimate Positivity at order 5: this seems to require even more delicate analysis.

An obvious, but possibly tedious future direction would be to tie up the book-keeping loose ends, and meticulously account for the complexity of our techniques. We chose to work with algebraic numbers; in settings involving rational numbers where scaling to integers and accessing an PosSLP oracle is viable, the complexity usually lies in PSPACE. However, this might blow up significantly in the absence of efficient positivity testing for a different class of arithmetic circuit.

In the grand Formal Methods scheme, the study of Hyperproperties [10] is an exciting natural way robustness problems for Linear Dynamical Systems could fit in. Hyperlogics reason about sets of traces of an infinite time system, rather than a single trace. They gained importance as a means to verify security in view of attacks like Meltdown and Spectre. A quintessential hyperproperty, for instance, would specify a reasonable notion of *indistinguishability* of traces. In that regard, our notions of **S**-Robust Positivity and **S**-Robust Uniform Ultimate Positivity bear striking resemblance. Exploring deeper connections is a fascinating future research avenue.

References

1 S. Akshay, Nikhil Balaji, Aniket Murhekar, Rohith Varma, and Nikhil Vyas. Near-Optimal Complexity Bounds for Fragments of the Skolem Problem. In Christophe Paul and Markus Bläser, editors, *37th International Symposium on Theoretical Aspects of Computer Science (STACS 2020)*, volume 154 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 37:1–37:18, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.STACS.2020.37.

2 S. Akshay, Hugo Bazille, Blaise Genest, and Mihir Vahanwala. On robustness for the skolem and positivity problems. In *STACS 2022*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.STACS.2022.5.

3 S. Akshay, Hugo Bazille, Blaise Genest, and Mihir Vahanwala. On robustness for the skolem, positivity and ultimate positivity problems, 2022. arXiv:2211.02365.

4 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.

5 Christel Baier, Florian Funke, Simon Jantsch, Toghrul Karimov, Engel Lefaucheux, 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*, volume 182 of *LIPIcs*, pages 36:1–36:17, 2020.

6 S. Basu, R. Pollack, and M. F. Roy. *Algorithms in Real Algebraic Geometry*. Springer, 2nd edition, 2006.

7 Valérie Berthé and Jungwon Lee. Dynamics of ostrowski skew-product: I. limit laws and hausdorff dimensions, 2022. arXiv:2108.06780.

8 Avraham Bourla. The ostrowski expansions revealed, 2016. arXiv:1605.07992.

9 Mark Braverman. Termination of integer linear programs. In *International Conference on Computer Aided Verification*, pages 372–385. Springer, 2006.

10 Michael R. Clarkson and Fred B. Schneider. Hyperproperties. In *2008 21st IEEE Computer Security Foundations Symposium*, pages 51–65, 2008. doi:10.1109/CSF.2008.7.

11 H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1993.

12 Julian D'Costa, Toghrul 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*, volume 202 of *LIPIcs*, pages 34:1–34:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

13 Julian D'Costa, Toghrul 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)*, volume 241 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 40:1–40:13, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.MFCS.2022.40.

14 Manuel Eberl and René Thiemann. Factorization of polynomials with algebraic coefficients. *Archive of Formal Proofs*, November 2021. Formal proof development. URL: https://isa-afp.org/entries/Factor_Algebraic_Polynomial.html.

15 Graham Everest, Alfred J. van der Poorten, Igor E. Shparlinski, and Thomas Ward. Recurrence Sequences. In *Mathematical surveys and monographs*, 2003.

16 Vesa Halava, Tero Harju, Mika Hirvensalo, and Juhani Karhumäki. Skolem's problem – On the border between decidability and undecidability. Technical Report 683, Turku Centre for Computer Science, 2005.

17 A Khintchine. Neuer beweis und verallgemeinerung eines hurwitzschen satzes. *Mathematische Annalen*, 111:631–637, 1935.

18 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.

19 Richard Lipton, Florian Luca, Joris Nieuwveld, Joël Ouaknine, David Purser, and James Worrell. On the skolem problem and the skolem conjecture. In *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, LICS '22, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3531130.3533328.

20 M. Mignotte. Some useful bounds. In *Computer Algebra*, 1982.

21 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.

22 Eike Neumann. Decision problems for linear recurrences involving arbitrary real numbers. *Logical Methods in Computer Science*, 17(3), 2021.

23 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.

24 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.

25 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.

26 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.

27 Ashish Tiwari. Termination of linear programs. In *Computer-Aided Verification, CAV*, volume 3114 of *LNCS*, pages 70–82. Springer, July 2004.

28 Nikolai Vereshchagin. The problem of appearance of a zero in a linear recurrence sequence. *Mat. Zametki*, 38(2):609–615, 1985.

A Appendix: Notation and Prerequisites

For the purposes of discussing robustness, we shall use $\mathcal{B}$ to denote the unit Euclidean ball in $\mathbb{R}^\kappa$, centred at the origin. Similarly, we use $\mathcal{B}_{\mathbf{S}}$ to denote the set of $\mathbf{d}$ such that $\mathbf{d}^T\mathbf{S}\mathbf{d} \leq 1$. For real column vectors $\mathbf{x}, \mathbf{y}$, we use $\langle \mathbf{x}, \mathbf{y} \rangle$ to denote the inner product $\mathbf{x}^T\mathbf{y} = \mathbf{y}^T\mathbf{x}$. The notation $||\mathbf{x}||$ denotes the standard ℓ^2-norm $\sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$.

Throughout this paper, $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{C}$ respectively denote the natural numbers, integers, rationals, reals, and complex numbers. $\alpha \in \mathbb{C}$ is said to be algebraic if it is a root of a polynomial with integer coefficients. Algebraic numbers form an algebraically closed field, denoted by $\overline{\mathbb{Q}}$. We denote the field of real algebraic numbers by $\mathbb{A}$.

This Appendix contains a brief initiation to this number field $\mathbb{A}$ and $\overline{\mathbb{Q}}$. The key takeaways are that the usual arithmetic as well as polynomial root computation can be carried out with perfect precision, and that the First Order Theory of the Reals $\langle \mathbb{R}; +, \cdot, \geq, 0, 1 \rangle$ is a decidable logical system powerful enough to fit our purposes.

A.1 Algebraic Numbers: Arithmetic

For an algebraic number α, its defining polynomial p_α is the unique polynomial in $\mathbb{Z}[X]$ of least degree 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 by size(p); its height, denoted by $H(p)$, is the maximum absolute value of the coefficients of p; $d(p)$ denotes the degree of p. The height $H(\alpha)$ and degree $d(\alpha)$ of α are defined to be the height and degree of p_α.

For any $p \in \mathbb{Z}[X]$, the distance between distinct roots is effectively lower bounded in terms of its degree and height [20]. This bound allows one to represent an algebraic number α as a 4-tuple (p, a, b, r) where p is the defining polynomial, and $a + bi$ is a rational approximation of sufficient precision $r \in \mathbb{Q}$. We use size(α) to denote the size of this representation, i.e., number of bits needed to write down this 4-tuple.

Given a polynomial $p \in \mathbb{Z}[X]$, one can compute its roots in polynomial time [6]. Recently, implementations of algorithms to factor polynomials in $\overline{\mathbb{Q}}[X]$ have been verified [14]. Given α, β two algebraic numbers, one can always compute the representations of $\alpha + \beta$, $\alpha\beta$, $\frac{1}{\alpha}$, $\Re(\alpha)$, $\Im(\alpha)$, $|\alpha|$, and decide $\alpha = \beta$, $\alpha > \beta$ in polynomial time with respect to the size of their representations. [6, 11].

A.2 First Order Theory of the Reals

This logical theory reasons about the universe of real numbers, and is denoted $\langle \mathbb{R}; +, \cdot, \geq, 0, 1 \rangle$. That is, variables take real values; terms can be added and multiplied, we have the comparison predicate, and direct access to the constants 0 and 1. Thus, our propositional atoms are inequalities involving polynomials with integer coefficients. With existential quantifiers and polynomials, we can thus express algebraic constants too. Formally, we have access to only the existential quantifier, negation, and disjunction; however, this can express the universal quantifier and all other Boolean connectives as well.

Variables are either quantified or free. Remarkably, the First Order Theory of the Reals admits quantifier elimination: for any formula $\chi(\mathbf{x})$, whose free variables are $\mathbf{x}$, there exists an **equivalent** formula $\psi(\mathbf{x})$ that does not contain any quantified variables. The following result is relevant to us.

▶ **Theorem 17** (Renegar [26]). *Let $\chi(\mathbf{x})$ be a first order formula interpreted over the theory of the reals. There exists a procedure that returns an equivalent quantifier-free formula $\psi(\mathbf{x})$ in disjunctive normal form. Moreover, if the total number of variables in χ is bounded a priori, this procedure runs in time polynomial in the size of the representation of χ.*

B Appendix: Ostrowski Numeration System

In this appendix, we prove Lemma 8. We state number-theoretic properties of the continued fraction representation and Ostrowski Numeration System without proof. We refer the reader to [8] for a more detailed exposition, and we closely follow the discussion surrounding [7, Propositions 1.1, 2.1] in our own proof. We first prove a slightly simpler statement.

▶ **Lemma 18.** *For every irrational number x, strictly decreasing real positive function ψ, and interval $\mathcal{I} = [\alpha, \beta] \subset [0, 1]$, $\alpha \neq \beta$, there exists $y \in \mathcal{I}$ such that $[nx - y] < \psi(n)$ for infinitely many n.*

Proof. Without loss of generality, we can assume that $x \in (0, 1)$. Consider the continued fraction representation of x: $[0; a_1, a_2, a_3, \dots]$

$$x = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{\ddots}}}}$$

where $a_1, a_2, a_3, \dots \in \mathbb{N}$. Let the rational approximation of x obtained by truncating the expansion at the k^{th} level be $\frac{p_k}{q_k}$, i.e. $\frac{p_1}{q_1} = \frac{1}{a_1}$, and so on. Let $\theta_k = q_k x - p_k$. We have that $|\theta_k| = (-1)^k \theta_k$. It is well known that $|\theta_k| < 1/q_k$. We define $q_{-1} = p_0 := 0$, and $p_{-1} = q_0 := 1$, so that for $k \geq 1$, the following recurrences hold:

$$p_k = a_k p_{k-1} + p_{k-2},\ q_k = a_k q_{k-1} + q_{k-2}.$$

We thus have that $q_k \geq \left(\frac{1+\sqrt{5}}{2}\right)^k = \phi^k$.

▶ **Proposition 19** ([7]). *Let irrational x and its continued fraction representation $[0; a_1, a_2, a_3, \dots]$ be as above. The infinite series*

$$\sum_{i=1}^{\infty} a_i |\theta_{i-1}|$$

converges.

▶ **Proposition 20** (Ostrowski Numeration System, [7]). *Every real number $y \in [0, 1)$ can be written uniquely in the form*

$$y = \sum_{i=1}^{\infty} b_i |\theta_{i-1}| = \sum_{i=1}^{\infty} (-1)^{i-1} b_i \theta_{i-1}$$

where $b_i \in \mathbb{N}$ $b_i \leq a_i$ for all $i \geq 1$. If for some i, $a_i = b_i$, then $b_{i+1} = 0$. $a_i \neq b_i$ for infinitely many odd, and infinitely many even indices i.

We prove Lemma 18 by using the free choice of b_i in this system to construct appropriate y. We first handle the issue of placing y in the correct interval $[\alpha, \beta]$. Let $\beta - \alpha = \delta$. We use Proposition 19 to argue that there exists a suffix of the infinite series, such that changing the suffix does not change the real number it represents by more than $\delta/2$. Then, we can simply fix the corresponding prefix of $(\alpha + \beta)/2$ to be the prefix of y.

Once this prefix is locked in, our strategy is to set b_i to 0 in even positions, and 1 in some odd positions, to ensure that for sufficiently large k, $n_k = \sum_{i=1}^{k} b_i(-1)^{i-1}q_{i-1}$ is positive, and increasing in k.

Now, notice that since b_i, p_i are all integers, for any y,

$$\begin{aligned}
[n_k x - y] &= \left[\sum_{i=1}^{k} b_i(-1)^{i-1}q_{i-1}x - \sum_{i=1}^{k} b_i(-1)^{i-1}p_{i-1} - y\right] \\
&= \left[\sum_{i=1}^{k} b_i(-1)^{i-1}\theta_{i-1} - y\right] \\
&= \left[-\sum_{i=k+1}^{\infty} b_i(-1)^{i-1}\theta_{i-1}\right] = \sum_{i=k+1}^{\infty} b_i|\theta_{i-1}| \\
&< \sum_{i=k+1}^{\infty} b_i\frac{1}{q_{i-1}} \leq \sum_{i=k+1}^{\infty} b_i\frac{1}{\phi^{i-1}} \leq \frac{c}{\phi^k}.
\end{aligned}$$

Note that the last constant c can be set independently of the choice of which b_i are 1, and which are 0: it comes from the convergence of the geometric sum. We now make the choice of where to set $b_i = 1$. To conclude the proof, we shall show that given a decreasing function ψ, we can ensure that for infinitely many distinct n_k,

$$[n_k x - y] < \frac{c}{\phi^k} \leq \psi(n_k) = \psi\left(\sum_{i=1}^{k} b_i(-1)^{i-1}q_{i-1}\right).$$

The first inequality is guaranteed. Suppose the second inequality does not hold. Then, from $i = k$ onwards, we keep assigning $b_i := 0$. This holds n_k constant as k increases, but decreases $\frac{c}{\phi^k}$. Eventually, the second inequality will indeed hold. After this point, for the next odd i, we can set b_i to 1, and get a new n_k. We continue this ad infinitum, and we are done. ◀

Now, to get infinitely many *even* n, apply Lemma 18 with $2x$, $[a, b]$, $\psi_0(n) = \psi(2n)$. For some choice of y, there will be infinitely many n such that $[2nx - y] < \psi_0(n) = \psi(2n)$. To get infinitely many *odd* n, we can take a subset of the interval, and shift it by x. Take $\psi_1(n) = \psi(2n + 1)$. For some choice of $y - x$, there will be infinitely many n such that $[n(2x) - (y - x)] = [(2n + 1)x - y] < \psi_1(n) = \psi(2n + 1)$.

C Appendix: Technical details of hardness proof

In this Appendix, we prove properties 25 and 26, which we restate:

$$\Psi(n, r) \Rightarrow n[nt - s] > \frac{(1 - \varepsilon)\sqrt{7r}}{4\pi}. \tag{29}$$

$$\neg\Psi(n, r) \Rightarrow n[nt - s] < \frac{\sqrt{7r}}{(1 - \varepsilon)4\pi}. \tag{30}$$

By definition, $\Psi(n, r)$ holds if and only if $rn^2 + \frac{r}{2} + 1 - \cos 2(\pi(nt - s)) \geq r\sqrt{n^4 + n^2 + 2}$. Through elementary algebraic manipulations, we can alternately group the terms as

$$1 - \cos(2\pi(nt - s)) \geq \frac{r}{2}\left(\frac{7n^2 + 14}{(n^2 + \sqrt{n^4 + n^2 + 2})(n^2 + 4 + \sqrt{n^4 + n^2 + 2})}\right) = r \cdot Q(n). \quad (31)$$

We note that in the limit, the ratio of $Q(n)$ to $7/8n^2$ tends to 1 from below. On the other hand, for small values of x, the expression $x^2/2$ is a close over-approximation for $1 - \cos x$. We capture the crucial interdependence in the following technical lemma.

▶ **Lemma 21.** *Let $r > 0$. For every $\varepsilon > 0$, we can compute N such that*

1. *For all $n \geq N$, $Q(n) > 7(1 - \varepsilon)^2/8n^2$.*
2. $1 - \cos x < 7r/8N^2 \;\Rightarrow\; 1 - \cos x \geq (1 - \varepsilon)^2 x^2/2$.

For some r, ε, let N be computed by Lemma 21. Consider $n \geq N$. In case $\Psi(n, r)$ holds, property 25 follows by considering the beginning and end of the chain of inequalities

$$2\pi^2[nt - s]^2 = \frac{[2\pi(nt - s)]_{2\pi}^2}{2} \geq 1 - \cos(2\pi(nt - s)) \geq r \cdot Q(n) > \frac{7r(1 - \varepsilon)^2}{8n^2}. \quad (32)$$

Similarly, if $\neg\Psi(n, r)$ holds, we can use Lemma 21 to construct the chain

$$2\pi^2(1-\varepsilon)^2[nt - s]^2 = \frac{(1 - \varepsilon)^2[2\pi(nt - s)]_{2\pi}^2}{2} \leq 1 - \cos(2\pi(nt - s)) < r \cdot Q(n) < \frac{7r}{8N^2}. \quad (33)$$