On the Complexity of the Orbit Problem

VENTSISLAV CHONEV, Institute of Science and Technology Austria
JOËL OUAKNINE, University of Oxford
JAMES WORRELL, University of Oxford

We consider higher-dimensional versions of Kannan and Lipton's Orbit Problem—determining whether a target vector space $\mathcal{V}$ may be reached from a starting point $\boldsymbol{x}$ under repeated applications of a linear transformation $\boldsymbol{A}$. Answering two questions posed by Kannan and Lipton in the 1980s, we show that when $\mathcal{V}$ has dimension one, this problem is solvable in polynomial time, and when $\mathcal{V}$ has dimension two or three, the problem is in $\mathbf{NP}^{\mathbf{RP}}$.

Categories and Subject Descriptors: F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems – Computations on matrices, Number-theoretic computations; G.2.1 [**Discrete Mathematics**]: Combinatorics – Recurrences and difference equations

General Terms: Algorithms, Theory, Verification

Additional Key Words and Phrases: Linear transformations, matrix orbits, linear recurrence sequences, Skolem's Problem, termination of linear programs

ACM Reference Format:
Ventsislav Chonev, Joël Ouaknine, and James Worrell, 2014. On the Complexity of the Orbit Problem. *J. ACM* 0, 0, Article 0 (2016), 40 pages.
DOI: http://dx.doi.org/10.1145/2857050

1. INTRODUCTION

The *Orbit Problem* was introduced by Harrison in [Harrison 1969] as a formulation of the reachability problem for linear sequential machines. The problem is stated as follows:

> Given a square matrix $\boldsymbol{A} \in \mathbb{Q}^{m\times m}$ and vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Q}^m$, decide whether there exists a non-negative integer n such that $\boldsymbol{A}^n\boldsymbol{x} = \boldsymbol{y}$.

The decidability of this problem remained open for over ten years, until it was shown to be decidable in polynomial time by Kannan and Lipton [Kannan and Lipton 1980]. In the conclusion of the journal version of their work [Kannan and Lipton 1986], the authors discuss a higher-dimensional extension of the Orbit Problem, as follows:

> Given a square matrix $\boldsymbol{A} \in \mathbb{Q}^{m\times m}$, a vector $\boldsymbol{x} \in \mathbb{Q}^m$, and a subspace $\mathcal{V}$ of $\mathbb{Q}^m$, decide whether there exists a non-negative integer n such that $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$.

As Kannan and Lipton point out, the higher-dimensional Orbit Problem is closely related to the *Skolem Problem*: given a square matrix $\boldsymbol{A} \in \mathbb{Q}^{m\times m}$ and vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{Q}^m$, decide whether there exists a non-negative integer n such that $\boldsymbol{y}^{\boldsymbol{T}}\boldsymbol{A}^n\boldsymbol{x} = 0$. Indeed, the Skolem Problem is the special case of the higher-dimensional Orbit Problem in

Authors' addresses: Ventsislav Chonev, IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria
{Joël Ouaknine, James Worrell}, Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, UK

DOI: http://dx.doi.org/10.1145/2857050

which the target space $\mathcal{V}$ has dimension $m-1$. The sequence of numbers $\langle u_n \rangle_{n=0}^{\infty}$ given by $u_n = \boldsymbol{y}^{\boldsymbol{T}} \boldsymbol{A}^n \boldsymbol{x}$ is a linear recurrence sequence. A well-known result, the Skolem-Mahler-Lech Theorem, states that the set $\{n : u_n = 0\}$ of zeros of any linear recurrence is the union of a finite set and finitely many arithmetic progressions [Mahler 1935; Lech 1953; Skolem 1934; Hansel 1986]. Moreover, it is known how to compute effectively the arithmetic progressions in question [Berstel and Mignotte 1976]. The main difficulty in deciding the Skolem Problem is thus to determine whether the finite component of the set of zeros is empty.

The decidability of the Skolem Problem has been open for many decades [Halava et al. 2005; Tao 2008], and it is therefore unsurprising that there has been virtually no progress on the higher-dimensional Orbit Problem since its introduction in [Kannan and Lipton 1986]. In fact, decidability of the Skolem Problem for matrices of dimension three and four [Mignotte et al. 1984; Vereshchagin 1985] was only established slightly prior to the publication of [Kannan and Lipton 1986], and there has been no substantial progress on this front since.[1] In terms of lower bounds, the strongest known result for the Skolem Problem is **NP**-hardness [Blondel and Portier 2002], which therefore carries over to the unrestricted version of the higher-dimensional Orbit Problem.

Kannan and Lipton speculated in [Kannan and Lipton 1986] that for target spaces of dimension one the Orbit Problem might be solvable, "hopefully with a polynomial-time bound". They moreover observed that the cases in which the target space $\mathcal{V}$ has dimension two or three seem "harder", and proposed this line of research as an approach towards the Skolem Problem. In spite of this, to the best of our knowledge, no progress has been recorded on the higher-order Orbit Problem in the intervening two-and-a-half decades.

Our main result is the following. We show that the higher-dimensional Orbit Problem is in **PTIME** if the target space has dimension one, and in $\mathbf{NP}^{\mathbf{RP}}$ if the target space has dimension two or three. While we make extensive use of the techniques of [Mignotte et al. 1984; Vereshchagin 1985] on the Skolem Problem, our results, in contrast, are independent of the dimension of the matrix $\boldsymbol{A}$.

Strictly speaking, Kannan and Lipton's original work on the Orbit Problem concerned the case that the target was an affine subspace of dimension 0. Our main results entail an $\mathbf{NP}^{\mathbf{RP}}$ complexity bound in case the target is an affine subspace of dimension 1 or 2, simply by embedding into a vector-space problem one dimension higher.

The following example illustrates some of the phenomena that emerge in the Orbit Problem for two-dimensional target spaces. Consider the following matrix and initial vector:

$$\boldsymbol{A} = \begin{bmatrix} 4 & 6 & 14 & 21 \\ -8 & -2 & -28 & -7 \\ -2 & -3 & -6 & -9 \\ 4 & 1 & 12 & 3 \end{bmatrix} \quad \boldsymbol{x} = \begin{bmatrix} 28 \\ -14 \\ -10 \\ 5 \end{bmatrix}$$

Then with target space

$$\mathcal{V} = \{(u_1, u_2, u_3, u_4) \in \mathbb{Q}^4 : 4u_1 + 7u_3 = 0, 4u_2 + 7u_4 = 0\}$$

it can be shown that $\boldsymbol{A}^n \boldsymbol{x} \in \mathcal{V}$ if and only if n has residue 2 modulo 6. Such periodic behaviour can be analysed in terms of the eigenvalues of the matrix $\boldsymbol{A}$. These are $\lambda\omega$, $\overline{\lambda}\omega$, $\lambda\overline{\omega}$ and $\overline{\lambda\omega}$, where $\omega = e^{\pi i/3}$ is a primitive 6-th root of unity and $\lambda = (-1 + i\sqrt{39})/2$.

[1] A proof of decidability of the Skolem Problem for linear recurrence sequences of order five was announced in [Halava et al. 2005]. However, as pointed out in [Ouaknine and Worrell 2012], the proof seems to have a serious gap.

The key observation is that the eigenvalues of $\boldsymbol{A}$ fall into only two classes under the equivalence relation $\sim$, defined by $\alpha \sim \beta$ if and only if α/β is a root of unity.

We handle such instances by analysing the equivalence classes of $\sim$. We show that, provided $\sim$ has sufficiently many equivalence classes, there is at most one exponent n such that $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$. Computable bounds on such an n are obtained utilising the work of [Mignotte et al. 1984; Vereshchagin 1985], quantifying and strengthening some of the bounds given for the Skolem Problem. In the case of a one-dimensional target subspace $\mathcal{V}$, the resulting bound is polynomial in the size of the problem representation, allowing for all exponents n up to the bound to be checked directly and yielding a polynomial-time algorithm. Unfortunately, when $\mathcal{V}$ has dimension two or three, the bounds on n are exponential in the size of the input, leading to an $\mathbf{NP}^{\mathbf{RP}}$ *guess-and-check* procedure, in which an $\mathbf{RP}$ oracle is used to check whether $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$ for a guessed value of n. Finally, the case in which the eigenvalues of $\boldsymbol{A}$ have fewer equivalence classes under $\sim$ is handled explicitly using a case analysis on the residue of n modulo the least common multiple of the orders of all ratios of eigenvalues which are roots of unity. For each such residue class, we show how to determine whether it contains exponents n for which $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$. Noting that there are at most exponentially many such residue classes, we can directly incorporate this case analysis into an $\mathbf{NP}^{\mathbf{RP}}$ algorithm.

1.1. Related Work

Aside from its connection to the Skolem Problem, the higher-dimensional Orbit Problem is closely related to termination problems for linear programs (see, e.g., [Ben-Amram et al. 2012; Braverman 2006; Tiwari 2004]) and to reachability questions for discrete linear dynamical systems (cf. [Halava et al. 2005]). Another related problem is the *Polyhedron Hitting Problem*, which replaces the target space with an intersection of half-spaces. In [Tarasov and Vyalyi 2011], the Polyhedron Hitting Problem is related to decision problems in formal language theory. Some partial decidability results for this problem are given in [Chonev et al. 2015]. Let us also mention the more recent work of Arvind and Vijayaraghavan [Arvind and Vijayaraghavan 2011] which places the original Orbit Problem in the logspace counting hierarchy $\mathbf{GapLH}$.

Another generalisation of the Orbit Problem was considered in [Cai et al. 2000] and shown to be decidable in polynomial time. This asks, given commuting rational matrices $\boldsymbol{A}$, $\boldsymbol{B}$, and $\boldsymbol{C}$, whether there exist integers i and j such that $\boldsymbol{A}^i\boldsymbol{B}^j = \boldsymbol{C}$.

A continuous version of the Orbit Problem is considered in [Hainry 2008]. Here one studies linear differential equations of the form $\boldsymbol{x}'(t) = \boldsymbol{A}\boldsymbol{x}(t)$ for a rational matrix $\boldsymbol{A}$. The problem is to decide, for a given initial condition $\boldsymbol{x}(0)$ and target vector $\boldsymbol{v}$, whether there exists t such that $\boldsymbol{x}(t) = \boldsymbol{v}$. The main result of [Hainry 2008] shows decidability of this problem.

2. PAPER OUTLINE

This work is based on our conference paper [Chonev et al. 2013]. The main technical results are the following theorems:

THEOREM 2.1. *Suppose we are given an instance of the Orbit Problem, comprising a square matrix $\boldsymbol{A} \in \mathbb{Q}^{m\times m}$, a vector $\boldsymbol{x} \in \mathbb{Q}^m$ and a subspace $\mathcal{V} \subseteq \mathbb{Q}^m$ with $\dim(\mathcal{V}) \leq 3$. Let $\|I\|$ be the length of the description of the input data. There exists a bound $N = 2^{\mathcal{O}(\|I\|)}$ such that if the instance is positive, then there exists a witness (that is, $n \in \mathbb{N}$ with $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$) such that $n < N$.*

THEOREM 2.2. *The Orbit Problem with $\dim(\mathcal{V}) \leq 3$ is in $\mathbf{NP}^{\mathbf{RP}}$. Further, if $\dim(\mathcal{V}) = 1$, then the problem is in $\mathbf{PTIME}$.*

In this section we give a high-level overview of the argument.

Firstly, we must emphasise that the fixed-dimensional versions of the Orbit Problem referred to by Theorems 2.1 and 2.2 are closely related to the Skolem Problem for linear recurrence sequences of order at most four. In the interest of clarity, we have confined our treatment of the Skolem Problem to the Appendix. We employ two powerful tools from transcendence theory, due to Baker-Wüstholz and van der Poorten, as well as standard results from algebraic number theory, to prove our main result on the Skolem Problem, Theorem C.1, which shows the existence of effective bounds on the zeros of LRS of order at most four and upon which our results on the Orbit Problem build.

We now outline the structure of the argument which establishes Theorems 2.1 and 2.2. The first step is a reduction to a similar problem, a polynomial version of the *matrix power problem:* given a rational square matrix $\boldsymbol{A}$ and polynomials $P_1, \ldots, P_d \in \mathbb{Q}[x]$ such that $P_1(\boldsymbol{A}), \ldots, P_d(\boldsymbol{A})$ are linearly independent over $\mathbb{Q}$, determine whether there exists n such that $\boldsymbol{A}^n$ lies in the $\mathbb{Q}$-vector space $span\{P_1(\boldsymbol{A}), \ldots, P_d(\boldsymbol{A})\}$. The reduction does not increase the dimension of the target space, so we will always have $d \leq 3$. The reduction can be carried out in polynomial time and rests entirely on standard techniques from linear algebra.

For the second step, we construct a *Master System.* This is a system of equations, based on the eigenvalues of $\boldsymbol{A}$ and the polynomials $P_1, \ldots, P_d$. It has $d+1$ unknowns: the exponent n and the coefficients $\kappa_1, \ldots, \kappa_d$ which witness the membership of $\boldsymbol{A}^n$ in $span\{P_1(\boldsymbol{A}), \ldots, P_d(\boldsymbol{A}))\}$. The solutions $(n, \kappa_1, \ldots, \kappa_d)$ of the Master System will be exactly the solutions of the matrix equation $\boldsymbol{A}^n = \kappa_1 P_1(\boldsymbol{A}) + \cdots + \kappa_d P_d(\boldsymbol{A})$. The domain of n is $\mathbb{N}$ throughout. Since the input data is rational, any solution $(n, \kappa_1, \ldots, \kappa_d)$ of the Master System will necessarily have $\kappa_1, \ldots, \kappa_d \in \mathbb{Q}$.

Next, in Section 4, we give a polynomial-time decision procedure to determine whether the Master System for an instance with a one-dimensional target space has a solution. The algorithm explicitly manipulates the equations in the system, preserving the set of solutions at every step, to determine the existence of a solution in polynomial time, settling the one-dimensional case of Theorem 2.2. The section rests critically on Theorem C.1 for non-degenerate linear recurrence sequences of order 2, which allows us to bound the exponent in all cases when $\boldsymbol{A}$ has two eigenvalues whose ratio is not a root of unity. In all other situations, the given Orbit instance essentially reduces to a system of linear congruences, easily solved using the Chinese Remainder Theorem. The solution method yields the full set of witness exponents n when this set is finite, or a description of the witness set as an arithmetic progression when it is infinite. Thus, if the problem instance is positive, a witness exponent which is at most exponentially large is automatically guaranteed to exist, as promised by Theorem 2.1, by virtue of our ability to write it down using polynomially many bits.

An important concept for the cases of two- and three-dimensional target spaces is the notion of *degeneracy*. An instance $(\boldsymbol{A}, \boldsymbol{x}, \mathcal{V})$ of the Orbit Problem is defined as degenerate if there exist two distinct eigenvalues of $\boldsymbol{A}$ whose quotient is a root of unity, otherwise the instance is non-degenerate. In general, it is possible to reduce an arbitrary Orbit Problem instance to a set of non-degenerate instances. Let L be the least common multiple of the orders of all quotients of eigenvalues of $\boldsymbol{A}$ which are roots of unity. For each $j \in \{0, \ldots, L-1\}$, consider separately the problem of deciding whether there exists $n \in \mathbb{N}$ such that $(\boldsymbol{A}^L)^n \left(\boldsymbol{A}^j \boldsymbol{x}\right) \in \mathcal{V}$. These instances are all non-degenerate,[2] and the original problem instance is positive if and only if at least one of these L non-degenerate instances is positive. Unfortunately, this reduction to the non-degenerate

[2]Indeed, the eigenvalues of $\boldsymbol{A}^L$ are exactly λ_i^L where λ_i are the eigenvalues of $\boldsymbol{A}$. If for any two distinct such eigenvalues, say $\lambda_i^L \neq \lambda_j^L$, we have $\left(\lambda_i^L / \lambda_j^L\right)^t = 1$, then λ_i / λ_j must also be a root of unity. Then by the definition of L, $\lambda_i^L / \lambda_j^L = 1$, which gives the contradiction $\lambda_i^L = \lambda_j^L$.

case carries an exponential overhead, as L is, in general, exponentially large in the size of the input data.

Instead, we adopt the following strategy for solving the Orbit Problem for possibly degenerate instances. Assume that as part of the input, we are given the residue $r = n \bmod L$. Thus, we are interested in determining whether the Master System has a solution $(n, \kappa_1, \ldots, \kappa_d)$ with exponent n such that $r = n \bmod L$. We will prove that for any r, there exists a bound N_r such that if there exists such an exponent with residue r, then one exists which does not exceed the bound N_r. Furthermore, $N_r = 2^{\mathcal{O}(||I'||)}$, where $\|I'\| = \|I\| + \|r\|$ is the length of the input augmented with the binary representation of r. This is clearly sufficient to prove Theorem 2.1: simply take $N = \max\{N_r : r \in \{0, \ldots, L-1\}\}$. The case analysis on r simplifies the Master System considerably, effectively eliminating degeneracy as a concern, and allowing us to derive the existence of N_r using our results on the Skolem Problem for LRS of order 3 and 4. For each fixed r, algebraic manipulation yields either a 'small' witness n of the correct residue, or a non-degenerate linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of low order such that if the Master System has a solution with exponent n with the desired residue r, then $u_n = 0$. The description of this linear recurrence sequence is computable in polynomial time from the input instance and r. Since $\|r\| = ||I||^{\mathcal{O}(1)}$, it follows that the length of the description of $\langle u_n \rangle_{n=0}^{\infty}$ is $\|u\| = ||I'||^{\mathcal{O}(1)} = ||I||^{\mathcal{O}(1)}$, so by Theorem C.1, the desired bound N_r exists and $N_r = 2^{\mathcal{O}(||I||)}$.

We must emphasise that this algebraic manipulation of the Master System and the calculation of the description of $\langle u_n \rangle_{n=0}^{\infty}$ is not part of the decision procedure for the Orbit Problem. Rather its purpose is to prove the existence of the desired bounds N_r, and hence of N. We make use of the observation that this manipulation can, in principle, be carried out in polynomial time, to conclude $N_r = 2^{\mathcal{O}(||I'||)}$ and $N = 2^{\mathcal{O}(||I||)}$, and hence establish Theorem 2.1.

Given the bound N of Theorem 2.1, we employ a *guess-and-check* procedure to obtain the complexity upper bounds of Theorem 2.2. Since N is at most exponentially large in the size of the input, an **NP** procedure can guess an exponent n such that $n < N$. Then we compute $\boldsymbol{A}^n\boldsymbol{x}$ by iterated squaring, thereby using polynomially many arithmetic operations. Moreover, all integers that occur in this algorithm have a polynomial-sized representation via arithmetic circuits. Now, to verify $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$, we compute the determinant of $\boldsymbol{B}^{\boldsymbol{T}}\boldsymbol{B}$, where $\boldsymbol{B}$ is the matrix whose columns are $\boldsymbol{A}^n\boldsymbol{x}$ and the basis vectors specifying $\mathcal{V}$, also as an arithmetic circuit. Clearly, n is a witness to the problem instance if and only if this determinant is zero. This is easy to determine with an **EqSLP** oracle, so we have membership in $\mathbf{NP}^{\mathbf{EqSLP}}$. It is known that $\mathbf{EqSLP} \subseteq \mathbf{coRP}$ [Schönhage 1979], so we have membership in $\mathbf{NP}^{\mathbf{RP}}$, thereby establishing Theorem 2.2.

3. REDUCTION

3.1. Matrix power problem

Suppose we are given a matrix $\boldsymbol{A} \in \mathbb{Q}^{m \times m}$, a vector $\boldsymbol{x} \in \mathbb{Q}^m$ and a target vector space $\mathcal{V} \subseteq \mathbb{Q}^m$ specified by a basis of rational vectors $\boldsymbol{y_1}, \ldots, \boldsymbol{y_k}$. We wish to decide whether there exists $n \in \mathbb{N}$ such that $\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V}$.

Observe that we can rescale $\boldsymbol{A}$ in polynomial time by the least common multiple of all denominators appearing in $\boldsymbol{A}$. This reduces the general problem to the sub-problem in which $\boldsymbol{A}$ is an integer matrix.

Let $\nu = \max\{m \mid \boldsymbol{x}, \boldsymbol{A}\boldsymbol{x}, \ldots, \boldsymbol{A}^m\boldsymbol{x} \text{ are linearly independent}\}$, $B = \{\boldsymbol{x}, \boldsymbol{A}\boldsymbol{x}, \ldots, \boldsymbol{A}^\nu\boldsymbol{x}\}$, $\mathcal{U} = \mathit{span}(B)$ and $\boldsymbol{D} = [\, \boldsymbol{x} \;\; \boldsymbol{A}\boldsymbol{x} \;\; \ldots \;\; \boldsymbol{A}^\nu\boldsymbol{x} \,]$. It is clear that $\mathcal{U}$ is invariant under the linear transformation $\boldsymbol{A}$, so consider the restriction of $\boldsymbol{A}$ to $\mathcal{U}$. Suppose $\boldsymbol{b} = (b_0, \ldots, b_\nu)^T$ are the coordinates of $\boldsymbol{A}^{\nu+1}\boldsymbol{x}$ with respect to B, that is, $\boldsymbol{A}^{\nu+1}\boldsymbol{x} = \boldsymbol{D}\boldsymbol{b}$. The restriction of $\boldsymbol{A}$

to $\mathcal{U}$ with respect to the basis B is described by the matrix

$$\boldsymbol{M} = \begin{bmatrix} 0 & 0 & \dots & 0 & b_0 \\ 1 & 0 & \dots & 0 & b_1 \\ 0 & 1 & \dots & 0 & b_2 \\ 0 & 0 & \ddots & 0 & \vdots \\ 0 & 0 & \dots & 1 & b_\nu \end{bmatrix}.$$

It is easy to check that $\boldsymbol{DM} = \boldsymbol{AD}$. Thus, if some vector $\boldsymbol{z}$ has coordinates $\boldsymbol{z'}$ with respect to B, so that $\boldsymbol{z} = \boldsymbol{Dz'}$, then $\boldsymbol{Az}$ has coordinates $\boldsymbol{Mz'}$ with respect to B, so that $\boldsymbol{Az} = \boldsymbol{DMz'}$. By induction, for all $n \in \mathbb{N}$, $\boldsymbol{A}^n\boldsymbol{x} = \boldsymbol{DM}^n\boldsymbol{x'}$, where $\boldsymbol{x'} = (1, 0, \dots, 0)^T$. Next we calculate a basis $\{\boldsymbol{w_1}, \dots, \boldsymbol{w_t}\}$ for $\mathcal{W} \stackrel{\text{def}}{=} \mathcal{U} \cap \mathcal{V}$. Then, if $\boldsymbol{w_i}$ are such that $\boldsymbol{w_i} = \boldsymbol{Dw'_i}$ for all i, we have

$$\boldsymbol{A}^n\boldsymbol{x} \in \mathcal{V} \iff \boldsymbol{A}^n\boldsymbol{x} \in \mathcal{W} \iff \boldsymbol{M}^n\boldsymbol{x'} \in span\{\boldsymbol{w'_1}, \dots, \boldsymbol{w'_t}\}.$$

Notice that the matrix $\boldsymbol{M}$ describes a restriction of the linear transformation denoted by $\boldsymbol{A}$, so its eigenvalues are a subset of the eigenvalues of $\boldsymbol{A}$. In particular, since $\boldsymbol{A}$ was rescaled to an integer matrix, the eigenvalues of $\boldsymbol{M}$ are algebraic integers as well.

Define the matrices $\boldsymbol{T_1}, \dots, \boldsymbol{T_t}$ by

$$\boldsymbol{T_i} = [\, \boldsymbol{w'_i} \;\; \boldsymbol{Mw'_i} \;\; \dots \;\; \boldsymbol{M}^\nu\boldsymbol{w'_i} \,].$$

We will show that $\boldsymbol{M}^n\boldsymbol{x'} \in span\{\boldsymbol{w'_1}, \dots, \boldsymbol{w'_t}\}$ if and only if $\boldsymbol{M}^n \in span\{\boldsymbol{T_1}, \dots, \boldsymbol{T_t}\}$. If for some coefficients κ_i we have

$$\boldsymbol{M}^n = \sum_{i=0}^{t} \kappa_i \boldsymbol{T_i},$$

then considering the first column of both sides, we have

$$\boldsymbol{M}^n\boldsymbol{x'} = \sum_{i=0}^{t} \kappa_i \boldsymbol{w'_i}.$$

Conversely, suppose $\boldsymbol{M}^n\boldsymbol{x'} = \sum_{i=0}^{t} \kappa_i \boldsymbol{w'_i}$. Then note that $\boldsymbol{x'}, \boldsymbol{Mx'}, \dots, \boldsymbol{M}^\nu\boldsymbol{x'}$ are just the unit vectors of size $\nu + 1$. Multiplying by $\boldsymbol{M}^j$ for $j = 0, \dots, \nu$ gives $\boldsymbol{M}^{n+j}\boldsymbol{x'} = \sum_{i=0}^{t} \kappa_i \boldsymbol{M}^j\boldsymbol{w'_i}$. The left-hand side is exactly the $(j+1)$-th column of $\boldsymbol{M}^n$, whereas $\boldsymbol{M}^j\boldsymbol{w'_i}$ on the right-hand side is exactly the $(j+1)$-th column of $\boldsymbol{T_i}$. So we have $\boldsymbol{M}^n = \sum_{i=0}^{t} \kappa_i \boldsymbol{T_i}$.

Thus, we have reduced the Orbit Problem to the *Matrix Power Problem:* determining whether some power of a given matrix lies inside a given vector space of matrices. Now we will perform a further reduction step. It is clear that within the space $\mathcal{T} \stackrel{\text{def}}{=} span\,\{\boldsymbol{T_1}, \dots, \boldsymbol{T_t}\}$ it suffices to consider only matrices of the shape $P(\boldsymbol{M})$ where $P \in \mathbb{Q}[x]$. We find a basis for the space $\mathcal{P} \stackrel{\text{def}}{=} \{P(\boldsymbol{M}) \mid P \in \mathbb{Q}[x]\}$ and then a basis $\{P_1(\boldsymbol{M}), \dots, P_s(\boldsymbol{M})\}$ for $\mathcal{P} \cap \mathcal{T}$. Then $\boldsymbol{M}^n \in \mathcal{T} \iff \boldsymbol{M}^n \in \mathcal{P} \cap \mathcal{T}$. We call the problem of determining, given $\boldsymbol{M}$ and $P_1, \dots, P_s$, whether there exists $n \in \mathbb{N}$ such that $\boldsymbol{M}^n \in span\{P_1(\boldsymbol{M}), \dots, P_s(\boldsymbol{M})\}$, the *polynomial version* of the matrix power problem. Observe that $\dim(\mathcal{V}) \geq \dim(\mathcal{T}) \geq \dim(\mathcal{T} \cap \mathcal{P})$, so the dimension of the target vector space does not grow during the described reductions. All described operations may be performed in polynomial time using standard techniques from linear algebra.

3.2. Master System of equations

Suppose now we have an instance $(\boldsymbol{A}, P_1, \ldots, P_s)$ of the polynomial version of the matrix power problem. Calculate the minimal polynomial of $\boldsymbol{A}$ and obtain canonical representations of its roots $\alpha_1, \ldots, \alpha_k$, that is, the eigenvalues of $\boldsymbol{A}$. This may be done in polynomial time, see Section A.1. Throughout, for an eigenvalue α_i we will denote by $mul(\alpha_i)$ the multiplicity of α_i as a root of the minimal polynomial of the matrix.

Fix an exponent $n \in \mathbb{N}$ and coefficients $\kappa_1, \ldots, \kappa_s \in \mathbb{C}$ and define the polynomials $P(x) = \sum_{i=1}^{s} \kappa_i P_i(x)$ and $Q(x) = x^n$. It is easy to see that

$$Q(\boldsymbol{A}) = P(\boldsymbol{A})$$

if and only if

$$\forall i \in \{1, \ldots, k\}. \forall j \in \{0, \ldots, mul(\alpha_i) - 1\}. P^{(j)}(\alpha_i) = Q^{(j)}(\alpha_i). \tag{1}$$

Indeed, $P - Q$ is zero at $\boldsymbol{A}$ if and only if the minimal polynomial of $\boldsymbol{A}$ divides $P - Q$, that is, each α_i is a root of $P - Q$ with multiplicity at least $mul(\alpha_i)$, or equivalently, each α_i is a root of $P - Q$ and its first $mul(\alpha_i) - 1$ derivatives.

Thus, in order to decide whether there exists an exponent n and coefficients κ_i such that $\boldsymbol{A}^n = \sum_{i=1}^{s} \kappa_i P_i(\boldsymbol{A})$, it is sufficient to solve the system of equations (1) where the unknowns are $n \in \mathbb{N}$ and $\kappa_1, \ldots, \kappa_s \in \mathbb{C}$. Each eigenvalue α_i contributes $mul(\alpha_i)$ equations which specify that $P(x)$ and its first $mul(\alpha_i) - 1$ derivatives all vanish at α_i.

For brevity in what follows, we will denote by $eq(\alpha_i, j)$ the j-th derivative equation contributed to the system by α_i, that is, $P^{(j)}(\alpha_i) = Q^{(j)}(\alpha_i)$. This notation is defined only for $0 \leq j < mul(\alpha_i)$. We will also denote by $Eq(\alpha_i)$ the set of equations contributed by α_i to the system:

$$Eq(\alpha_i) = \{eq(\alpha_i, 0), \ldots, eq(\alpha_i, mul(\alpha_i) - 1)\}.$$

For example, if the minimal polynomial of $\boldsymbol{A}$ has roots α_1, α_2, α_3 with multiplicities $mul(\alpha_i) = i$ and the target space is $span\{P_1(\boldsymbol{A}), P_2(\boldsymbol{A})\}$ then the system contains six equations:

$$\begin{aligned}
\alpha_1^n &= \kappa_1 P_1(\alpha_1) + \kappa_2 P_2(\alpha_1) \\
\alpha_2^n &= \kappa_1 P_1(\alpha_2) + \kappa_2 P_2(\alpha_2) \\
n\alpha_2^{n-1} &= \kappa_1 P_1'(\alpha_2) + \kappa_2 P_2'(\alpha_2) \\
\alpha_3^n &= \kappa_1 P_1(\alpha_3) + \kappa_2 P_2(\alpha_3) \\
n\alpha_3^{n-1} &= \kappa_1 P_1'(\alpha_3) + \kappa_2 P_2'(\alpha_3) \\
n(n-1)\alpha_3^{n-2} &= \kappa_1 P_1''(\alpha_3) + \kappa_2 P_2''(\alpha_3)
\end{aligned}$$

Then $eq(\alpha_3, 0)$ is the equation

$$\alpha_3^n = \kappa_1 P_1(\alpha_3) + \kappa_2 P_2(\alpha_3)$$

and $Eq(\alpha_2)$ is the two equations

$$\begin{aligned}
\alpha_2^n &= \kappa_1 P_1(\alpha_2) + \kappa_2 P_2(\alpha_2) \\
n\alpha_2^{n-1} &= \kappa_1 P_1'(\alpha_2) + \kappa_2 P_2'(\alpha_2)
\end{aligned}$$

4. ONE-DIMENSIONAL TARGET SPACE

Suppose we are given a one-dimensional matrix power problem instance $(\boldsymbol{A}, P)$ and wish to decide whether $\boldsymbol{A}^n \in span\{P(\boldsymbol{A})\}$ for some n. We have constructed a system of equations in the exponent n and the coefficient κ as in (1). For example, if the roots of

the minimal polynomial of $\boldsymbol{A}$ are $\alpha_1, \alpha_2, \alpha_3$ with multiplicities $mul(\alpha_j) = j$, the system is:

$$\begin{aligned}
\alpha_1^n &= \kappa P(\alpha_1) \\
\alpha_2^n &= \kappa P(\alpha_2) \\
n\alpha_2^{n-1} &= \kappa P'(\alpha_2) \\
\alpha_3^n &= \kappa P(\alpha_3) \\
n\alpha_3^{n-1} &= \kappa P'(\alpha_3) \\
n(n-1)\alpha_3^{n-2} &= \kappa P''(\alpha_3)
\end{aligned}$$

In this section we will describe how such systems may be solved in polynomial time. First, we perform some preliminary calculations.

(1) We check whether $\kappa = 0$ has a corresponding n which solves the matrix equation $\boldsymbol{A}^n = \kappa P(\boldsymbol{A})$, that is, whether $\boldsymbol{A}$ is nilpotent. Otherwise, assume $\kappa \neq 0$.
(2) Let $k = \max_j\{mul(\alpha_j)\}$. We check for all $n < k$ whether $\boldsymbol{A}^n$ is a multiple of $P(\boldsymbol{A})$. If so, we are done. Otherwise, assume $n \geq k$.
(3) We check whether $\alpha_j = 0$ for some j. If so, then all of the equations $Eq(\alpha_i)$ are of the form $0 = \kappa P^{(t)}(0)$, which is equivalent to $0 = P^{(t)}(0)$. We can easily check whether these equations are satisfied. If so, we dismiss them from the system without changing the set of solutions. If not, then there is no solution and we are done. Now we assume $\alpha_j \neq 0$ for all j.
(4) Finally, we check whether the right-hand side $\kappa P^{(t)}(\alpha_j)$ of some equation is equal to 0, by dividing $P^{(t)}(x)$ by the minimal polynomial of α_j. If this is the case, then the problem instance is negative, because the left-hand sides are all non-zero.

Let $eq(\alpha_i, k)/eq(\alpha_j, t)$ denote the equation obtained from $eq(\alpha_i, k)$ and $eq(\alpha_j, t)$ by asserting that the ratio of the left-hand sides equals the ratio of the right-hand sides, that is,

$$\frac{n(n-1)\dots(n-k+1)\alpha_i^{n-k}}{n(n-1)\dots(n-t+1)\alpha_j^{n-t}} = \frac{P^{(k)}(\alpha_i)}{P^{(t)}(\alpha_j)}.$$

We compute representations of all quotients α_i/α_j, and consider three cases.

Case I. Some quotient α_i/α_j is not a root of unity. Then $eq(\alpha_i, 0)$ and $eq(\alpha_j, 0)$ together imply $eq(\alpha_i, 0)/eq(\alpha_j, 0)$, that is,

$$\left(\frac{\alpha_i}{\alpha_j}\right)^n = \frac{P(\alpha_i)}{P(\alpha_j)}.$$

In Section A.1, we discuss the efficient representation and manipulation of algebraic numbers. By Lemma A.1, we can compute representations of $P(\alpha_i)/P(\alpha_j)$ and α_i/α_j in polynomial time. Then by Lemma D.1 in Section D, n is bounded by a polynomial in the input. We check $\boldsymbol{A}^n \in span\{P(\boldsymbol{A})\}$ for all n up to the bound and we are done.

Case II. All quotients α_i/α_j are roots of unity, and all roots of the minimal polynomial of $\boldsymbol{A}$ are simple. Then the system is equivalent to

$$\kappa = \frac{\alpha_1^n}{P(\alpha_1)} \wedge \bigwedge_{i<j} \frac{eq(\alpha_i, 0)}{eq(\alpha_j, 0)}.$$

It is sufficient to determine whether there exists some n which satisfies

$$\bigwedge_{i<j} \frac{eq(\alpha_i, 0)}{eq(\alpha_j, 0)}. \tag{2}$$

Consider each equation $eq(\alpha_i, 0)/eq(\alpha_j, 0)$:

$$\left(\frac{\alpha_i}{\alpha_j}\right)^n = \frac{P(\alpha_i)}{P(\alpha_j)}. \tag{3}$$

Suppose α_i/α_j is an r-th root of unity. If the right-hand side of (3) is also an r-th root of unity, then the solutions of (3) are $n \equiv t \bmod r$ for some t. If not, then (3) has no solution, so the entire system (1) has no solution, and the problem instance is negative. By Lemma A.1, we can determine in polynomial time whether the right-hand side of (3) is a root of unity, and if so, calculate t. We transform each equation in (2) into an equivalent congruence in n. This gives a system of congruences in n which is equivalent to (2). We solve it using the Chinese Remainder Theorem. The problem instance is positive if and only if the system of congruences has a solution.

Case III. All quotients α_i/α_j are roots of unity, and $f_A(x)$ has repeated roots. We transform the system into an equivalent one in the following way. First, we include in the new system all the quotients of equations $eq(\alpha_i, 0)$ as in Case 2. Second, for each repeated root α_i of $f_A(x)$, we take the quotients $\bigwedge_{j=0}^{mul(\alpha_i)-2} eq(\alpha_i, j)/eq(\alpha_i, j+1)$. Third, we include the equation $\kappa = \alpha_1/P(\alpha_1)$.

$$\bigwedge_{i<j} \frac{eq(\alpha_i, 0)}{eq(\alpha_j, 0)} \wedge \bigwedge_i \bigwedge_{j=0}^{mul(\alpha_i)-2} \frac{eq(\alpha_i, j)}{eq(\alpha_i, j+1)} \wedge \kappa = \frac{\alpha_1}{P(\alpha_1)}.$$

We solve the first conjunct as in Case 2. If there is no solution, then we are done. Otherwise, the solution is some congruence $n \equiv t_1 \bmod t_2$. For the remainder of the system, each ratio $eq(\alpha_i, j)/eq(\alpha_i, j+1)$ contributed by a repeated root α_i has the shape

$$\frac{\alpha_i}{n-j} = \frac{P^{(j)}(\alpha_i)}{P^{(j+1)}(\alpha_i)},$$

which is equivalent to

$$n = j + \frac{P^{(j+1)}(\alpha_i)}{P^{(j)}(\alpha_i)}\alpha_i. \tag{4}$$

For each such equation (4), we calculate the right-hand side in polynomial time, using the methods outlined in Section A.1, and check whether it is in $\mathbb{N}$. If not, then the system has no solution. Otherwise, (4) points to a single candidate n_0. We do this for all equations where n appears outside the exponent. If they point to the same value of n, then the system is equivalent to

$$\begin{aligned} n &\equiv t_1 \bmod t_2 \\ n &= n_0 \\ \kappa &= \alpha_1^n/P(\alpha_1) \end{aligned}$$

We check whether n_0 satisfies the congruence and we are done.

5. TWO-DIMENSIONAL TARGET SPACE

Suppose we are given a rational square matrix $\boldsymbol{A}$ and polynomials P_1, P_2 with rational coefficients such that $P_1(\boldsymbol{A})$ and $P_2(\boldsymbol{A})$ are linearly independent over $\mathbb{Q}$. We want to decide whether there exists $n \in \mathbb{N}$ such that $\boldsymbol{A}^n$ lies in the $\mathbb{Q}$-vector space $span\{P_1(\boldsymbol{A}), P_2(\boldsymbol{A})\}$. We have derived a Master System of equations (1) in the unknowns (n, κ_1, κ_2) whose solutions are precisely the solutions of the matrix equation $\boldsymbol{A}^n = \kappa_1 P_1(\boldsymbol{A}) + \kappa_2 P_2(\boldsymbol{A})$.

In this section, we will show that there exists a bound N, exponentially large in the size of the input, such that if the problem instance is positive, then there exists a witness exponent n with $n < N$. This will be sufficient to show that the problem is in the complexity class $\mathbf{NP}^{\mathbf{RP}}$, as outlined earlier.

Notice that we may freely assume that the eigenvalues of $\boldsymbol{A}$ are non-zero. Indeed, if 0 is an eigenvalue, then consider $eq(0,0)$:

$$0 = \kappa_1 P_1(0) + \kappa_2 P_2(0).$$

If at least one of $P_1(0)$, $P_2(0)$ is non-zero, then we have a linear dependence between κ_1, κ_2. Then we express one of the coefficients κ_1, κ_2 in terms of the other, obtaining a Master System of dimension 1, and then the claim follows inductively. Otherwise, if $P_1(0) = P_2(0) = 0$, then $eq(0,0)$ is trivially satisfied for all n, κ_1, κ_2, so we remove $eq(0,0)$ from the Master System without altering the set of solutions. We examine in this way all equations contributed by 0, either removing them from the system, or obtaining a lower-dimensional system which then yields the required bound N inductively.

As outlined in Section 2, we show the existence of the bound N by performing a case analysis on $n \bmod L$, where

$$L = \text{lcm}\{\text{order}(\lambda_i/\lambda_j) : \lambda_i, \lambda_j \text{ eigenvalues of } \boldsymbol{A} \text{ and } \lambda_i/\lambda_j \text{ root of unity}\}.$$

We will show that for any fixed value $r \in \{0, \dots, L-1\}$, there exists a bound N_r, exponentially large in the size of the input, such that if the Master System has a solution with exponent of residue r modulo L, then it has a solution with exponent n such that $n < N_r$. To obtain the bounds N_r, we show how the Master System can be manipulated algebraically in polynomial time to yield a non-degenerate linear recurrence sequence of order 3 whose zeros are a superset of the exponents n which solve the Master System. This manipulation is a proof technique to show the existence of the bound N_r, not a feature of the algorithm. The decision method is instead the guess-and-check procedure explained in Section 2.

Thus, from here onwards, we assume we are given a fixed r, which increases the input size only polynomially, and are interested solely in exponents n with $n \bmod L = r$. Since we admit degenerate problem instances, we need to consider the relation $\sim$ on the eigenvalues of $\boldsymbol{A}$, defined by

$$\alpha \sim \beta \text{ if and only if } \alpha/\beta \text{ is a root of unity.}$$

It is clear that $\sim$ is an equivalence relation. The equivalence classes $C_1, \dots, C_k$ of $\sim$ are of two kinds. First, a class can be its own image under complex conjugation:

$$C_i = \{\overline{\alpha} \mid \alpha \in C_i\}$$

Each such self-conjugate class $\{\alpha_1, \dots, \alpha_s\}$ has the form $\{\alpha\omega_1, \dots, \alpha\omega_s\}$ where ω_i are roots of unity, and $|\alpha_j| = \alpha \in \mathbb{R} \cap \mathbb{A}$. Call this α the *representative* of the equivalence class C_i. Second, if an equivalence class is not self-conjugate, then its image under complex conjugation must be another equivalence class of $\sim$. Thus, the remaining equivalence classes of $\sim$ are grouped into pairs (C_i, C_j) such that $C_i = \{\overline{x} \mid x \in C_j\} = \overline{C_j}$. In this case, we can write C_i and C_j as

$$C_i = \{\lambda\omega_1, \dots, \lambda\omega_s\}$$

$$C_j = \{\overline{\lambda\omega_1}, \dots, \overline{\lambda\omega_s}\}$$

where ω_i are roots of unity, $\lambda \in \mathbb{A}$ and $\arg(\lambda)$ is an irrational multiple of 2π. Call λ the representative of C_i and $\overline{\lambda}$ the representative of C_j.

Observe that the representatives of self-conjugate classes are distinct positive real numbers, and that no ratio of representatives can be a root of unity. Recall also that

we can assume the eigenvalues of $\boldsymbol{A}$ are algebraic integers, as a by-product of the reduction from the Orbit Problem. Since roots of unity and their multiplicative inverses are algebraic integers, it follows that the representatives of equivalence classes must also be algebraic integers.

Let

$$Eq(C) = \bigcup_{\alpha \in C} Eq(\alpha)$$

denote the set of equations contributed to the system by the eigenvalues in C, and let

$$Eq(C,i) = \bigcup_{\substack{\alpha \in C \\ mul(\alpha) > i}} \{eq(\alpha, i)\}$$

denote the set of i-th derivative equations contributed by the roots in C.

To show the existence of the required bound N_r, we will perform a case analysis on the number of equivalence classes of $\sim$.

Case I. Suppose $\sim$ has exactly one equivalence class $C = \{\alpha\omega_1, \ldots, \alpha\omega_s\}$, necessarily self-conjugate, with representative α. Consider the set of equations $Eq(C, 0)$:

$$\begin{aligned}(\alpha\omega_1)^n &= \kappa_1 P_1(\alpha\omega_1) + \kappa_2 P_2(\alpha\omega_1) \\ &\vdots \\ (\alpha\omega_s)^n &= \kappa_1 P_1(\alpha\omega_s) + \kappa_2 P_2(\alpha\omega_s)\end{aligned}$$

For our fixed r, the values of $\omega_1^n, \ldots, \omega_s^n$ are easy to calculate in polynomial time, since ω_i are roots of unity whose order divides L. Then the equations $Eq(C, 0)$ are equivalent to

$$\begin{bmatrix} \alpha^n \\ \vdots \\ \alpha^n \end{bmatrix} = \boldsymbol{B} \begin{bmatrix} \kappa_1 \\ \kappa_2 \end{bmatrix}, \tag{5}$$

where $\boldsymbol{B}$ is an $s \times 2$ matrix over $\mathbb{A}$ which, given r, is computable in polynomial time. Next we subtract the first row of (5) from rows $2, \ldots, s$, obtaining

$$\alpha^n = c_1\kappa_1 + c_2\kappa_2 \wedge \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} = \boldsymbol{B'} \begin{bmatrix} \kappa_1 \\ \kappa_2 \end{bmatrix}.$$

Here (c_1, c_2) is the first row of the matrix $\boldsymbol{B}$, and $\boldsymbol{B'}$ is the result of subtracting (c_1, c_2) from each of the bottom $s-1$ rows of $\boldsymbol{B}$. Thus, $Eq(C, 0)$ is equivalent to $\alpha^n = c_1\kappa_1 + c_2\kappa_2$ together with the constraint that $(\kappa_1, \kappa_2)^T$ must lie in the nullspace of $\boldsymbol{B'}$. We now consier the nullspace of $\boldsymbol{B'}$. If its dimension is less than 2, then we have a linear constraint on κ_1, κ_2. This constraint is of the form $\kappa_1 = \chi\kappa_2$ when the nullspace of $\boldsymbol{B'}$ has dimension 1, and is $\kappa_1 = \kappa_2 = 0$ when the nullspace is of dimension 0. In both cases, the Master System is equivalent to a lower-dimensional one which may be computed in polynomial time, so the existence of the bound N_r follows inductively. In the case when the nullspace of $\boldsymbol{B'}$ has dimension 2, then the linear constraint is vacuous, and $Eq(C, 0)$ is equivalent to $\alpha^n = c_1\kappa_1 + c_2\kappa_2$.

In the same way, for this fixed r, $Eq(C, 1)$ reduces to a single first-derivative equation:

$$n\alpha^{n-1} = c_3\kappa_1 + c_4\kappa_2.$$

We do this for all $Eq(C, i)$, obtaining a system of equations equivalent to (1) based on the representative of C, rather than the actual eigenvalues in C. Denote the resulting set of equations by $\mathcal{F}(Eq(C))$.

If some eigenvalue $x \in C$ has $mul(x) \geq 3$, then $\mathcal{F}(Eq(C))$ contains the following triple of equations:

$$\begin{bmatrix} \alpha^n \\ n\alpha^{n-1} \\ n(n-1)\alpha^{n-2} \end{bmatrix} = \kappa_1 \begin{bmatrix} c_1 \\ c_3 \\ c_5 \end{bmatrix} + \kappa_2 \begin{bmatrix} c_2 \\ c_4 \\ c_6 \end{bmatrix}. \tag{6}$$

If the vectors on the right-hand side of (6) are linearly independent over $\mathbb{A}$, then they specify a plane in $\mathbb{A}^3$, and the triple states that the point on the left-hand side must lie on this plane. Letting $(A_1, A_2, A_3)^T$ be the normal of the plane, we obtain

$$\begin{aligned} &A_1\alpha^n + A_2 n\alpha^{n-1} + A_3 n(n-1)\alpha^{n-2} = 0 \\ \iff &A_1\alpha^2 + A_2 n\alpha + A_3 n(n-1) = 0. \end{aligned}$$

This is a quadratic equation in n. It has at most two roots, both at most exponentially large in the size of the input, so we just take N_r to be the greater root. If the vectors on the right-hand side of (6) are linearly dependent over $\mathbb{A}$, then the exponents n which solve (6) are precisely those which solve:

$$\begin{bmatrix} \alpha^n \\ n\alpha^{n-1} \\ n(n-1)\alpha^{n-2} \end{bmatrix} = \kappa_1 \begin{bmatrix} c_1 \\ c_3 \\ c_5 \end{bmatrix}.$$

We divide the first equation by the second to obtain

$$\frac{\alpha}{n} = \frac{c_1}{c_3},$$

which limits n at most one, exponentially large, candidate value $\alpha c_3/c_1$.

If all eigenvalues x in C have $mul(x) \leq 2$ and at least one has $mul(x) = 2$, then $\mathcal{F}(Eq(C))$ consists of exactly two equations:

$$\begin{bmatrix} \alpha^n \\ n\alpha^{n-1} \end{bmatrix} = \kappa_1 \begin{bmatrix} c_1 \\ c_3 \end{bmatrix} + \kappa_2 \begin{bmatrix} c_2 \\ c_4 \end{bmatrix}. \tag{7}$$

If $(c_1, c_3)^T$ and $(c_2, c_4)^T$ are linearly independent over $\mathbb{A}$, then the right-hand side of (7) spans all of $\mathbb{A}^2$ as κ_1, κ_2 range over $\mathbb{A}$. Then (7) is solved by all $n \in \mathbb{N}$, so we can take $N_r = L$. Otherwise, the exponents n which solve (7) are exactly those which solve

$$\begin{bmatrix} \alpha^n \\ n\alpha^{n-1} \end{bmatrix} = \kappa_1 \begin{bmatrix} c_1 \\ c_3 \end{bmatrix}.$$

This limits n to at most one candidate value $\alpha c_3/c_1$, which is exponentially large in the input size.

Finally, if all eigenvalues x in C have $mul(x) = 1$, then $\mathcal{F}(Eq(C))$ contains only the equation

$$\alpha^n = \kappa_1 c_1 + \kappa_2 c_2,$$

which is solved by all $n \in \mathbb{N}$ if at least one of c_1, c_2 is non-zero, and has no solutions if $c_1 = c_2 = 0$. Either way, we take $N_r = L$ and are done.

Case II. Suppose $\sim$ has exactly two equivalence classes, C_1 and C_2, with respective representatives α and β, so that

$$C_1 = \{\alpha\omega_1, \ldots, \alpha\omega_s\},$$

$$C_2 = \{\beta\omega_1', \ldots, \beta\omega_l'\}.$$

The classes could be self-conjugate, in which case $\alpha, \beta \in \mathbb{A} \cap \mathbb{R}$, or they could be each other's image under complex conjugation, in which case $\alpha = \overline{\beta}$. In both cases, α/β is not a root of unity.

As in *Case I,* we transform the system $Eq(C_1) \wedge Eq(C_2)$ into the equivalent system $\mathcal{F}(Eq(C_1)) \wedge \mathcal{F}(Eq(C_2))$. If all eigenvalues x of $\boldsymbol{A}$ have $mul(x) = 1$, then the resulting system consists of two equations, one for each equivalence class of $\sim$:

$$\alpha^n = \kappa_1 c_1 + \kappa_2 c_2$$
$$\beta^n = \kappa_1 c_3 + \kappa_2 c_4$$

If $(c_1, c_3)^T$ and $(c_2, c_4)^T$ are linearly independent over $\mathbb{A}$, then there is a solution for each n, so just take $N_r = L$. Otherwise, it suffices to look for n which satisfies

$$\alpha^n = \kappa_1 c_1$$
$$\beta^n = \kappa_1 c_3$$

and hence

$$\left(\frac{\alpha}{\beta}\right)^n = \frac{c_1}{c_3}.$$

A bound on n follows from Lemma D.1. This argument relies crucially on the fact that α/β is not a root of unity.

If some eigenvalue x of $\boldsymbol{A}$ has $mul(x) \geq 2$, say $x \in C_1$, then the system contains the following triple of equations:

$$\begin{bmatrix} \alpha^n \\ n\alpha^{n-1} \\ \beta^n \end{bmatrix} = \kappa_1 \begin{bmatrix} c_1 \\ c_3 \\ c_5 \end{bmatrix} + \kappa_2 \begin{bmatrix} c_2 \\ c_4 \\ c_6 \end{bmatrix}. \tag{8}$$

If the vectors on the right-hand side of (8) are linearly dependent over $\mathbb{A}$, so that the right-hand side describes a space of dimension 1, it suffices to look for solutions to

$$\begin{bmatrix} \alpha^n \\ n\alpha^{n-1} \\ \beta^n \end{bmatrix} = \kappa_1 \begin{bmatrix} c_1 \\ c_3 \\ c_5 \end{bmatrix}.$$

Then dividing we obtain

$$\frac{\alpha}{n} = \frac{c_1}{c_3},$$

which limits n to at most one, exponentially large candidate value $\alpha c_3/c_1$. Otherwise, if the vectors on the right-hand side of (8) are linearly independent over $\mathbb{A}$, we calculate the normal $(A_1, A_2, A_3)^T$ to the plane described by them and obtain

$$A_1\alpha^n + A_2 n\alpha^{n-1} + A_3\beta^n = 0.$$

A bound on n which is exponential in the size of the input follows from Lemma F.4. This again relies on the fact that α/β cannot be a root of unity.

Case III. Suppose $\sim$ has at least three equivalence classes. Then we can choose eigenvalues α, β, γ, each from a distinct equivalence class, and consider $eq(\alpha, 0)$, $eq(\beta, 0)$ and $eq(\gamma, 0)$:

$$\begin{bmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{bmatrix} = \kappa_1 \begin{bmatrix} P_1(\alpha) \\ P_1(\beta) \\ P_1(\gamma) \end{bmatrix} + \kappa_2 \begin{bmatrix} P_2(\alpha) \\ P_2(\beta) \\ P_2(\gamma) \end{bmatrix}.$$

If the vectors on the right-hand side are linearly independent over $\mathbb{A}$, we eliminate κ_1, κ_2 to obtain

$$A_1\alpha^n + A_2\beta^n + A_3\gamma^n = 0.$$

The left-hand side is a non-degenerate linear recurrence sequence of order 3, so a bound on n follows from Lemmas F.1, F.2, F.3. If the vectors on the right-hand side are not linearly independent over $\mathbb{A}$, then we may equivalently consider

$$\begin{bmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{bmatrix} = \kappa_1 \begin{bmatrix} P_1(\alpha) \\ P_1(\beta) \\ P_1(\gamma) \end{bmatrix},$$

which gives

$$\left(\frac{\alpha}{\beta}\right)^n = \frac{P_1(\alpha)}{P_1(\beta)}.$$

An exponential bound on n follows from Lemma D.1, because α/β is not a root of unity.

Thus, we have now shown that for any $r \in \{0, \dots, L-1\}$, the required bound N_r exists and is at most exponential in the size of the input. Then $N = \max\{N_r : r \in \{0, \dots, L-1\}\}$ exists and is exponentially large, so the Orbit Problem with two-dimensional target space is in $\mathbf{NP}^{\mathbf{RP}}$, by the complexity argument of Section 2.

6. THREE-DIMENSIONAL TARGET SPACE

Suppose we are given a rational square matrix $\boldsymbol{A}$ and polynomials P_1, P_2, P_3 with rational coefficients such that $P_1(\boldsymbol{A}), P_2(\boldsymbol{A}), P_3(\boldsymbol{A})$ are linearly independent over $\mathbb{Q}$. We want to decide whether there exists $n \in \mathbb{N}$ such that $\boldsymbol{A}^n$ lies in the $\mathbb{Q}$-vector space $span\{P_1(\boldsymbol{A}), P_2(\boldsymbol{A}), P_3(\boldsymbol{A})\}$. We have derived a Master System of equations (1) in the unknowns $(n, \kappa_1, \kappa_2, \kappa_3)$ whose solutions are precisely the solutions of the matrix equation $\boldsymbol{A}^n = \kappa_1 P_1(\boldsymbol{A}) + \kappa_2 P_2(\boldsymbol{A}) + \kappa_3 P_3(\boldsymbol{A})$.

In this section, we will show that there exists a bound N, exponentially large in the size of the input, such that if the problem instance is positive, then there exists a witness exponent n with $n < N$. This will be sufficient to show that the problem is in the complexity class $\mathbf{NP}^{\mathbf{RP}}$, as outlined earlier.

The eigenvalues of $\boldsymbol{A}$ may be assumed to be non-zero algebraic numbers: if 0 is an eigenvalue, then $eq(0, 0)$ gives a linear dependence between the coefficients $\kappa_1, \kappa_2, \kappa_3$, yielding a lower-dimensional Master System, so the existence of the bound N follows inductively.

Following the strategy of the two-dimensional case, we will perform a case analysis on the residue of n modulo L: let $n \bmod L = r$ be fixed throughout this section. To obtain the required bound N, it is sufficient to derive a bound N_r, also exponentially large in the size of the input, such that if there exists a witness exponent of residue r modulo L, then such a witness may be found which does not exceed N_r. As in the two-dimensional case, we will select tuples of equations and obtain a bound on n using the results for the Skolem Problem for recurrences of order 4 in Section G. We will again perform a case analysis on the equivalence classes of the relation $\sim$.

Case I. Suppose there are at least two pairs of classes $(C_i, \overline{C_i})$, $(C_j, \overline{C_j})$ which are not self-conjugate. Then let $\alpha \in C_i$, $\beta = \overline{\alpha} \in \overline{C_i}$, $\gamma \in C_j$, $\delta = \overline{\gamma} \in \overline{C_j}$. Then we consider the tuple of equations

$$\begin{bmatrix} \alpha^n \\ \beta^n \\ \gamma^n \\ \delta^n \end{bmatrix} = \kappa_1 \begin{bmatrix} P_1(\alpha) \\ P_1(\beta) \\ P_1(\gamma) \\ P_1(\delta) \end{bmatrix} + \kappa_2 \begin{bmatrix} P_2(\alpha) \\ P_2(\beta) \\ P_2(\gamma) \\ P_2(\delta) \end{bmatrix} + \kappa_3 \begin{bmatrix} P_3(\alpha) \\ P_3(\beta) \\ P_3(\gamma) \\ P_3(\delta) \end{bmatrix}. \tag{9}$$

If the vectors on the right-hand side are linearly dependent over $\mathbb{A}$, then we rewrite the right-hand side as a linear combination of at most two vectors and obtain the required bound on n by considering a linear recurrence sequence of order 2 or 3. If the vectors on the right-hand side of (9) are linearly independent over $\mathbb{A}$, then we calculate the normal of the three-dimensional subspace of $\mathbb{A}^4$ that they span, obtaining an equation

$$A_1\alpha^n + A_2\beta^n + A_3\gamma^n + A_4\delta^n = 0 \tag{10}$$

and hence an exponential bound on n from Lemmas G.3 and G.4. We are relying on the fact that the ratios of $\alpha, \beta, \gamma, \delta$ are not roots of unity. Notice that we need (α, β) and (γ, δ) to be pairwise complex conjugates in order to apply Lemma G.4. Notice also that we may assume without loss of generality that $\alpha, \beta, \gamma, \delta$ are algebraic integers, as Lemma G.4 requires. Indeed, as remarked at the beginning of Section 3, the input data may be assumed to be over $\mathbb{Z}$, instead of $\mathbb{Q}$, with the simple technique of scaling the input by an integer chosen so as to 'clear the denominators'. Then $\boldsymbol{A}$ is an integer matrix, so its eigenvalues are algebraic integers.

Case II. Suppose now that there is exactly one pair of classes $(C_i, \overline{C_i})$ which are not self-conjugate. In general, for any eigenvalue x of $\boldsymbol{A}$ we must have $mul(x) = mul(\overline{x})$. Therefore, if any eigenvalue $\alpha \in C_i$ has $mul(\alpha) > 1$, we can select the tuple of equations $eq(\alpha, 0)$, $eq(\alpha, 1)$, $eq(\overline{\alpha}, 0)$, $eq(\overline{\alpha}, 1)$:

$$\begin{bmatrix} \alpha^n \\ \overline{\alpha}^n \\ n\alpha^{n-1} \\ n\overline{\alpha}^{n-1} \end{bmatrix} = \kappa_1 \begin{bmatrix} P_1(\alpha) \\ P_1(\overline{\alpha}) \\ P_1'(\alpha) \\ P_1'(\overline{\alpha}) \end{bmatrix} + \kappa_2 \begin{bmatrix} P_2(\alpha) \\ P_2(\overline{\alpha}) \\ P_2'(\alpha) \\ P_2'(\overline{\alpha}) \end{bmatrix} + \kappa_3 \begin{bmatrix} P_3(\alpha) \\ P_3(\overline{\alpha}) \\ P_3'(\alpha) \\ P_3'(\overline{\alpha}) \end{bmatrix}.$$

This gives a non-degenerate linear recurrence sequence of order 4 over $\mathbb{A}$ for a recurrence sequence with two repeated characteristic roots:

$$A_1\alpha^n + A_2\overline{\alpha}^n + A_3 n\alpha^{n-1} + A_4 n\overline{\alpha}^{n-1} = 0.$$

An exponential bound N on n follows from Lemma G.1, since $\alpha/\overline{\alpha}$ is not a root of unity.

We can now assume that eigenvalues in C_i and $\overline{C_i}$ contribute exactly one equation to the system. Now we use the fixed value of r to transform $Eq(C_i) \wedge Eq(\overline{C_i})$ into $\mathcal{F}(Eq(C_i)) \wedge \mathcal{F}(Eq(\overline{C_i}))$. Since all eigenvalues in C_i and $\overline{C_i}$ contribute one equation each, $\mathcal{F}(Eq(C_i)) \wedge \mathcal{F}(Eq(\overline{C_i}))$ is just

$$\begin{aligned} \lambda^n &= \kappa_1 c_1 + \kappa_2 c_2 + \kappa_3 c_3 \\ \overline{\lambda}^n &= \kappa_1 c_4 + \kappa_2 c_5 + \kappa_3 c_6 \end{aligned}$$

where $\lambda, \overline{\lambda}$ are the representatives of C_i and $\overline{C_i}$. We do the same to all self-conjugate classes as well, reducing the system of equations to an equivalent system based on the representatives of the equivalence classes, not the actual eigenvalues of $\boldsymbol{A}$. This is beneficial, because the representatives cannot divide to give roots of unity, so we can use 4-tuples of equations to construct non-degenerate linear recurrence sequences of order 4.

If there are at least two self-conjugate equivalence classes, with respective representatives α, β, we take the tuple

$$\begin{aligned} \lambda^n &= \kappa_1 c_1 + \kappa_2 c_2 + \kappa_3 c_3 \\ \overline{\lambda}^n &= \kappa_1 c_4 + \kappa_2 c_5 + \kappa_3 c_6 \\ \alpha^n &= \kappa_1 c_7 + \kappa_2 c_8 + \kappa_3 c_9 \\ \beta^n &= \kappa_1 c_{10} + \kappa_2 c_{11} + \kappa_3 c_{12} \end{aligned}$$

and obtain the following equation, where the left-hand side is a non-degenerate linear recurrence sequence:

$$A_1\lambda^n + A_2\overline{\lambda}^n + A_3\alpha^n + A_4\beta^n = 0.$$

Then we have an exponentially large bound N_r from Lemmas G.3 and G.4. Similarly, if there is only one self-conjugate equivalence class, with representative α, but some of its eigenvalues are repeated, we use the tuple

$$\begin{aligned}\lambda^n &= \kappa_1 c_1 + \kappa_2 c_2 + \kappa_3 c_3\\ \overline{\lambda}^n &= \kappa_1 c_4 + \kappa_2 c_5 + \kappa_3 c_6\\ \alpha^n &= \kappa_1 c_7 + \kappa_2 c_8 + \kappa_3 c_9\\ n\alpha^{n-1} &= \kappa_1 c_{10} + \kappa_2 c_{11} + \kappa_3 c_{12}\end{aligned}$$

to obtain the non-degenerate instance

$$A_1\lambda^n + A_2\overline{\lambda}^n + A_3\alpha^n + A_4 n\alpha^{n-1} = 0,$$

which gives an exponential bound N_r according to Lemma G.2. If there is exactly one self-conjugate class, with representative α, containing no repeated roots, then the system consists of three equations:

$$\begin{aligned}\lambda^n &= \kappa_1 c_1 + \kappa_2 c_2 + \kappa_3 c_3\\ \overline{\lambda}^n &= \kappa_1 c_4 + \kappa_2 c_5 + \kappa_3 c_6\\ \alpha^n &= \kappa_1 c_7 + \kappa_2 c_8 + \kappa_3 c_9\end{aligned}$$

Depending on whether the vectors $(c_1, c_4, c_7)^T$, $(c_2, c_5, c_8)^T$, $(c_3, c_6, c_9)^T$ are linearly independent over $\mathbb{A}$, either this triple is solved by all $n \in \mathbb{N}$ (in which case set $N_r = L$), or it reduces to a lower-dimensional Master System, yielding the claim inductively. Finally, if there are no self-conjugate classes, the system consists of only two equations:

$$\begin{aligned}\lambda^n &= \kappa_1 c_1 + \kappa_2 c_2 + \kappa_3 c_3\\ \overline{\lambda}^n &= \kappa_1 c_4 + \kappa_2 c_5 + \kappa_3 c_6\end{aligned}$$

Again, depending on the dimension of

$$span\left\{\begin{bmatrix} c_1\\ c_4\end{bmatrix}, \begin{bmatrix} c_2\\ c_5\end{bmatrix}, \begin{bmatrix} c_3\\ c_6\end{bmatrix}\right\},$$

we can either set the bound N_r to L (because the transformed Master System is solved by all $n \in \mathbb{N}$), or obtain N_r inductively from a lower-dimensional Master System.

Case III. All equivalence classes of $\sim$ are self-conjugate. The techniques used for this case are identical to the ones already presented. We use the fixed value of r to reduce to a non-degenerate system based on the representatives of the classes, with the number of equations contributed by each class determined by the maximum multiplicity of an eigenvalue in that class.

If there are less than four equations, then we study the dimension of the vector space spanned by the vectors on the right-hand side: if it has full dimension, then we see the Master System is satisfied by all n of the correct residue r, so we can just set $N_r = L$. Otherwise, we obtain the bound inductively from a lower-dimensional non-degenerate Master System.

On the other hand, if there are at least four equations, then we can choose four equations which have a solution for n if and only if an effectively computable non-degenerate LRS of order 4 vanishes at n. We then employ the bounds of Theorem C.1 concerning LRS of order 4 to obtain the desired N_r.

We remark here that it is only for this final case that we need the representatives of self-conjugate classes to be real, necessitating the choice of the magnitude of the eigenvalues in the class for representative, regardless of whether this magnitude is itself an eigenvalue. The reason for this technical point is that Lemma G.4, which gives a bound on the index of zeros of an LRS of order 4 with four distinct characteristic roots, requires that the characteristic roots be closed under complex conjugation. No strengthening of Lemma G.4 is known which avoids this precondition – as we remark in Section G, this is the reason why the Skolem Problem is open for LRS of order 4 over $\mathbb{A}$. If we had chosen the representative of a self-conjugate class to be an arbitrary (possibly complex) eigenvalue, we would obtain LRS of order 4 whose characteristic roots do not satisfy the precondition on Lemma G.4, and we would not be able to obtain our bound N_r here.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library.

Received September 2015; revised ; accepted

A. MATHEMATICAL TECHNIQUES

A.1. Algebraic numbers: representation and manipulation

A complex number α is *algebraic* if there exists a polynomial $P \in \mathbb{Q}[x]$ such that $P(\alpha) = 0$. The set of algebraic numbers, denoted by $\mathbb{A}$, is a subfield of $\mathbb{C}$. The *minimal polynomial* of α is the unique monic polynomial of least degree which vanishes at α. The *degree* of $\alpha \in \mathbb{A}$ is defined as the degree of its minimal polynomial and is denoted by $\deg(\alpha)$. The *height* of α, denoted by H_α, is defined as the maximum absolute value of the coefficients of the integer polynomial obtained by scaling the minimal polynomial of α by the least common multiple of the denominators of its coefficients. The roots of the minimal polynomial of α (including α) are called the *Galois conjugates* of α. The *absolute norm* of α, denoted $\mathcal{N}_{abs}(\alpha)$, is the product of the Galois conjugates of α. By Viete's laws, we have

$$\mathcal{N}_{abs}(\alpha) = (-1)^{\deg(\alpha)} \frac{a}{b}$$

where a, b are respectively the constant term and the leading coefficient of the minimal polynomial of α. It follows that $\mathcal{N}_{abs}(\alpha) \in \mathbb{Q}$. An *algebraic integer* is an algebraic number whose minimal polynomial has integer coefficients. The set of algebraic integers, denoted $\mathcal{O}_\mathbb{A}$, is a ring under the usual addition and multiplication. The algebraic integers are *integrally closed*, that is, the roots of any monic polynomial with coefficients in $\mathcal{O}_\mathbb{A}$ are all algebraic integers. For any $\alpha \in \mathbb{A}$, it is possible to find $\beta \in \mathcal{O}_\mathbb{A}$ and $m \in \mathbb{Z}$ such that $\alpha = \beta/m$.

The *canonical representation* of an algebraic number α is its minimal polynomial, along with a numerical approximation of $\Re(\alpha)$ and $\Im(\alpha)$ of sufficient precision to distinguish α from its Galois conjugates [Cohen 1993, Section 4.2.1]. More precisely, we represent α by the tuple

$$(P, x, y, R) \in (\mathbb{Q}[x] \times \mathbb{Q}^3)$$

meaning that α is the unique root of the irreducible (over $\mathbb{Q}$) polynomial P which lies inside the circle centred at (x, y) in the complex plane with radius R. A bound due to Mignotte [Mignotte 1982] states that for roots $\alpha_j \neq \alpha_k$ of a polynomial $P(x)$,

$$|\alpha_j - \alpha_k| > \frac{\sqrt{6}}{d^{(d+1)/2} H^{d-1}}, \tag{11}$$

where d and H are the degree and height of P, respectively. Thus, if R is restricted to be less than half the root separation bound, the representation is well-defined and allows for equality checking. Observe that given its minimal polynomial, the remaining data necessary to describe α is polynomial in the length of the input. Given $P \in \mathbb{Q}[x]$ and $k \in \mathbb{N}$, it is known how to obtain k bits of the roots of P in time polynomial in the length of the description of P and in the length of the binary representation of k [Pan 1996].

When we say an algebraic number α is given, we assume we have a canonical description of α. We will denote by $\|\alpha\|$ the length of this description, assuming that integers are expressed in binary and rationals are expressed as pairs of integers. Observe that $|\alpha|$ is an exponentially large quantity in $\|\alpha\|$ whereas $\log|\alpha|$ is polynomially large. Notice also that $1/\log|\alpha|$ is at most exponentially large in $\|\alpha\|$. For a rational a, $\|a\|$ is just the sum of the lengths of its numerator and denominator written in binary. For a polynomial $P \in \mathbb{Q}[x]$, $\|P\|$ will denote $\sum_{j=0}^{\deg(P)} \|p_j\|$, where p_j are the coefficients of P.

LEMMA A.1. *Given canonical representations of $\alpha, \beta \in \mathbb{A}$ and a polynomial $P \in \mathbb{Q}[x]$, it is possible to compute canonical descriptions of $\alpha \pm \beta$, $\alpha\beta^{\pm 1}$ and $P(\alpha)$ in time polynomial in the length of the input (that is, in $\|\alpha\| + \|\beta\| + \|P\|$).*

PROOF. Let R, Q be the minimal polynomials of α and β, respectively. Then the resultant of $R(x-y)$ and $Q(y)$, interpreted as polynomials in y with coefficients in $\mathbb{Q}[x]$, is a polynomial in x which vanishes at $\alpha+\beta$. We compute it in polynomial time using the Sub-Resultant algorithm (see Algorithm 3.3.7 in [Cohen 1993]) and factor it into irreducibles using the LLL algorithm [Lenstra et al. 1982]. Finally, we approximate the roots of each irreducible factor to identify the minimal polynomial of $\alpha+\beta$. The degree of $\alpha+\beta$ is at most $\deg(\alpha)\deg(\beta)$, while its height is bounded by $H_{\alpha+\beta} \le H_\alpha^{\deg(\alpha)} H_\beta^{\deg(\beta)}$ [Zippel 1997]. Therefore, by (11), a polynomial number of bits suffices to describe $\alpha+\beta$ unambiguously. Similarly, we can compute canonical representations of $\alpha-\beta$, $\alpha\beta$ and α/β in polynomial time using resultants, see [Cohen 1993].

To calculate $P(\alpha)$ we repeatedly use addition and multiplication. It suffices to prove that all intermediate results may be represented in polynomial space. It is clear that their degrees are at most $\deg(\alpha)$, but it is not obvious how quickly the coefficients of their minimal polynomials grow. However, there is a simple reason why their representation is polynomially bounded. Let $\boldsymbol{A}$ be the companion matrix of the minimal polynomial of α. Then $P(\alpha)$ is an eigenvalue of $P(\boldsymbol{A})$. We can calculate $P(\boldsymbol{A})$ using only polynomial space. Then from the Leibniz formula

$$\det(\lambda \boldsymbol{I} - P(\boldsymbol{A})) = \sum_{\sigma \in S_n} \operatorname{sign}(\sigma) \prod_{i=1}^{n} (\lambda \boldsymbol{I} - P(\boldsymbol{A}))_{i,\sigma(i)},$$

it is evident that the coefficients of the characteristic polynomial of $P(\boldsymbol{A})$ are exponentially large in the length of the input, so their representation requires only polynomial space. This characteristic polynomial may be factored into irreducibles in polynomial time, so the description of $P(\alpha)$ and of all intermediate results is polynomially bounded. □

It is trivial to check whether $\alpha = \beta$ and whether α belongs to one of $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$. It takes only polynomial time to determine whether α is a root of unity, and if so, to calculate its order and phase.

A.2. Number fields and ideals

In this section, we recall some terminology and results from algebraic number theory. For more details, see [Cohen 1993; Stewart and Tall 2002]. We also define the ideal-

counting function v_P, which is a notion of magnitude of algebraic numbers distinct from the usual absolute value. We follow the presentation of [Halava et al. 2005].

An *algebraic number field* is a field extension $\mathbb{K}$ of $\mathbb{Q}$ which, considered as a $\mathbb{Q}$-vector space, has finite dimension. This dimension is called the *degree* of the number field and is denoted by $[\mathbb{K} : \mathbb{Q}]$. Given two algebraic numbers α and β, the *Field Membership Problem* is to determine whether $\beta \in \mathbb{Q}(\alpha)$ and, if so, to return a polynomial P with rational coefficients such that $\beta = P(\alpha)$. This problem can be solved using the LLL algorithm, see [Cohen 1993, Section 4.5.4].

For any number field $\mathbb{K}$, there exists an element $\theta \in \mathbb{K}$ such that $\mathbb{K} = \mathbb{Q}(\theta)$. Such a θ is called a *primitive element* of $\mathbb{K}$ and satisfies $\deg(\theta) = [\mathbb{K} : \mathbb{Q}]$. The proof is constructive: there is always a primitive element for $\mathbb{Q}(\alpha_1, \alpha_2)$ of the form $\alpha_1 + l\alpha_2$ for some small integer $l \leq \deg(\alpha_1)\deg(\alpha_2)$. Thus, repeatedly using an algorithm for the Field Membership Problem for different l is guaranteed to yield a primitive element for $\mathbb{Q}(\alpha_1, \alpha_2)$, and therefore by induction, for any number field $\mathbb{K} = \mathbb{Q}(\alpha_1, \ldots, \alpha_k)$ specified by algebraic numbers $\alpha_1, \ldots, \alpha_k$. Also using an algorithm for the Field Membership Problem, one can represent each α_j as a polynomial in θ and thereby determine a maximal $\mathbb{Q}$-linearly independent subset of $\{\alpha_1, \ldots, \alpha_k\}$.

There exist exactly $\deg(\theta)$ monomorphisms from $\mathbb{K}$ into $\mathbb{C}$, given by $\theta \to \theta_j$, where θ_j are the Galois conjugates of the primitive element θ. If $\alpha \in \mathbb{K}$, then $\deg(\alpha) | \deg(\theta)$. Moreover, if $\sigma_1, \ldots, \sigma_{\deg(\theta)}$ are the monomorphisms from $\mathbb{K}$ into $\mathbb{C}$ then $\sigma_1(\alpha), \ldots, \sigma_{\deg(\theta)}(\alpha)$ are exactly the Galois conjugates of α, each repeated $\deg(\theta)/\deg(\alpha)$ times. The *norm of* α *relative to* $\mathbb{K}$ is defined as

$$\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha) = \prod_{j=1}^{\deg(\theta)} \sigma_j(\alpha) = (\mathcal{N}_{abs}(\alpha))^{\deg(\theta)/\deg(\alpha)}$$

For a number field $\mathbb{K}$, the set $\mathcal{O}_\mathbb{K} = \mathcal{O}_\mathbb{A} \cap \mathbb{K}$ of algebraic integers in $\mathbb{K}$ forms a ring under the usual addition and multiplication. The ideals of $\mathcal{O}_\mathbb{K}$ are finitely generated, and form a commutative ring under the operations

$$IJ = [\{xy \mid x \in I, y \in J\}]$$

$$I + J = \{x + y \mid x \in I, y \in J\},$$

with unit $\mathcal{O}_\mathbb{K}$ and zero $\{0\}$, where $[S]$ denotes the ideal generated by S. An ideal P is *prime* if $P = AB$ implies $A = P$ or $A = [1]$. The fundamental theorem of ideal theory states that each non-zero ideal may be represented uniquely (up to reordering) as a product of prime ideals.

This theorem gives rise to the following *ideal-counting function* $v_P : \mathcal{O}_\mathbb{K} \backslash \{0\} \to \mathbb{N}$. For a fixed prime ideal P, we define $v_P(\alpha)$ to be the number of times P appears in the factorisation into prime ideals of $[\alpha]$. That is,

$$v_P(\alpha) = k \text{ if and only if } P^k \mid [\alpha] \text{ and } P^{k+1} \nmid [\alpha]$$

We also define $v_P(0) = \infty$. The function satisfies the following properties:

— $v_P(\alpha\beta) = v_P(\alpha) + v_P(\beta)$
— $v_P(\alpha + \beta) \geq \min\{v_P(\alpha), v_P(\beta)\}$
— If $v_P(\alpha) \neq v_P(\beta)$, then $v_P(\alpha + \beta) = \min\{v_P(\alpha), v_P(\beta)\}$.

For any $\alpha \in \mathbb{K}$ we can find an algebraic integer $\beta \in \mathcal{O}_\mathbb{K}$ and a rational integer $n \in \mathbb{Z} \subseteq \mathcal{O}_\mathbb{K}$ such that $\alpha = \beta/n$. We extend v_P to $\mathbb{K}$ by defining $v_P(\alpha) = v_P(\beta) - v_P(n)$. The first of the three properties of v_P above guarantees that this value is independent of the choice of β, n, making the extension of v_P to $\mathbb{K}$ well-defined. Note that the extension preserves the above three properties.

For an ideal $I \neq \{0\}$, the quotient ring $\mathcal{O}_\mathbb{K}/I$ is finite. The *norm* of I, denoted $\mathcal{N}(I)$, is defined as $|\mathcal{O}_\mathbb{K}/I|$. We define also $\mathcal{N}([0]) = \infty$. Notice that $\mathcal{N}(I) = 1$ if and only if $I = \mathcal{O}_\mathbb{K}$, otherwise $\mathcal{N}(I) \geq 2$. Each prime ideal P contains a unique prime number p, and $\mathcal{N}(P) = p^f$ for some natural number $f \geq 1$. In general,

$$|\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha)| = \mathcal{N}([\alpha]) \geq 2^{v_P(\alpha)}$$

since $\mathcal{N}(P) \geq 2$ for any prime ideal P. Hence,

$$v_P(\alpha) \leq \log_2 |\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha)| \leq \log_2 |\mathcal{N}_{abs}(\alpha)|^d$$

where $d = [\mathbb{K} : \mathbb{Q}]$. Thus, if we are given $\mathbb{K} = \mathbb{Q}(\alpha_1, \ldots, \alpha_k)$ for canonically represented algebraic numbers α_j and a canonically represented $\alpha \in \mathbb{K}$, we can observe that d is at most polynomially large in the length of the input and $|\mathcal{N}_{abs}(\alpha)|$ is at most exponentially large in the length of the input. Therefore, $v_P(\alpha)$ is only polynomially large.

The following lemma is simple, but useful:

LEMMA A.2. *Let $\mathbb{K}$ be a number field and $\alpha \in \mathbb{K}$ with $\alpha \notin \mathcal{O}_\mathbb{K}$. Then there exists a prime ideal P of $\mathcal{O}_\mathbb{K}$ such that $v_P(\alpha) \neq 0$.*

PROOF. There exist $\beta \in \mathcal{O}_\mathbb{K}$ and $m \in \mathbb{Z}$ such that $\alpha = \beta/m$. If $[\beta] = [m]$, then β and m are associates, so α must be a unit of $\mathcal{O}_\mathbb{K}$. Since $\alpha \notin \mathcal{O}_\mathbb{K}$, it follows that $[\beta] \neq [m]$, so the factorisations of $[\beta]$ and $[m]$ into prime ideals must differ. Therefore, $v_P(\beta) \neq v_P(m)$ for some prime ideal P, so $v_P(\alpha) \neq 0$. □

A.3. Transcendental number theory

We now move to some techniques from Transcendental Number Theory on which our results depend in a critical way. First, we state a powerful result due to Baker on linear forms of logarithms of algebraic numbers.

THEOREM A.3. *[Baker 1975, Theorem 3.1] Let $\alpha_1, \ldots, \alpha_m$ be non-zero algebraic numbers with degrees at most d and heights at most A. Further, let $\beta_0, \ldots, \beta_m$ be algebraic numbers with degrees at most d and heights at most $B \geq 2$. Write*

$$\Lambda = \beta_0 + \beta_1 \log(\alpha_1) + \ldots + \beta_m \log(\alpha_m)\,.$$

Then either $\Lambda = 0$ or $|\Lambda| > B^{-C}$, where C is an effectively computable number depending only on m, d, A and the chosen branch of the complex logarithm.

Various quantitative versions of this theorem are known with explicit constants, as well as sharper lower bounds for restricted cases. Of these, in the present paper we make use of the following result, due to Baker and Wüstholz, concerning the homogeneous case with (rational) integer coefficients:

THEOREM A.4. *[Baker and Wüstholz 1993] With the notation as in Theorem A.3, suppose $\beta_0 = 0$, $\alpha_1, \ldots, \alpha_m \neq 1$, $\beta_1, \ldots, \beta_m \in \mathbb{Z}$ and $A, B \geq e$. Let also $\log$ be the principal branch of the natural logarithm, defined by $\log(z) = \log|z| + i\arg(z)$, where $-\pi < \arg(z) \leq \pi$. Let also D be the degree of the extension field $\mathbb{Q}(\alpha_1, \ldots, \alpha_m)$ over $\mathbb{Q}$. Then if $\Lambda \neq 0$, then*

$$\log|\Lambda| > -(16mD)^{2(m+2)}(\log(A))^m \log(B).$$

The next theorem, due to van der Poorten [van der Poorten 1977] is analogous to Baker's bound, but with respect to P-adic valuations instead of the usual Archimedean absolute value.

THEOREM A.5. *[van der Poorten 1977] Let $\alpha_1, \ldots, \alpha_m$ be algebraic numbers of degree at most d belonging to a number field $\mathbb{K}$ and with heights at most A. Let P be a*

prime ideal of $\mathbb{K}$ *containing the rational prime* p*. Let also* $\beta_1, \ldots, \beta_m$ *be rational integers with absolute values at most* $B \geq e^2$*. If* $\alpha_1^{\beta_1}\alpha_2^{\beta_2}\ldots\alpha_m^{\beta_m} \neq 1$*, then*

$$v_P(\alpha_1^{\beta_1}\ldots\alpha_m^{\beta_m} - 1) \leq (16(m+1)d)^{12(m+1)}(p^d/\log(p))(\log(A))^m(\log(B))^2.$$

A.4. Algebraic integers near the unit circle

Suppose $\alpha \neq 0$ is an algebraic integer. It is easy to see that it is impossible for all the Galois conjugates of α to be strictly within the unit circle: just notice that the product of all Galois conjugates of α must be a non-zero integer by Viete's laws. Further, an old result due to Kronecker [Kronecker 1875] establishes that unless α is a root of unity, then at least one of its Galois conjugates must be strictly outside the unit circle. In this paper, we make use of the following theorem, due to Blanksby and Montgomery [Blanksby and Montgomery 1971], which strengthens Kronecker's result by providing an effective separation between this Galois conjugate and the unit circle.

THEOREM A.6. *Let* α *be an algebraic integer of degree* $d \geq 2$*. Then there is a Galois conjugate* $\sigma(\alpha)$ *of* α *such that* $|\sigma(\alpha)| > 1 + 1/(30d^2\log(6d))$.

A.5. Linear recurrence sequences

We now recall some basic properties of linear recurrence sequences. For more details, we refer the reader to [Everest et al. 2003; Halava et al. 2005]. Let $\mathbb{F}$ be $\mathbb{R}$ or $\mathbb{C}$ throughout this section. A *linear recurrence sequence (LRS)* over $\mathbb{F}$ is an infinite sequence $\langle u_n\rangle_{n=0}^{\infty}$ of terms in $\mathbb{F}$ such that there exists a natural number k and numbers $a_1, \ldots, a_k \in \mathbb{F}$ such that $a_k \neq 0$ and $\langle u_n\rangle_{n=0}^{\infty}$ satisfies the linear recurrence equation

$$u_{n+k} = a_1u_{n+k-1} + a_2u_{n+k-2} + \cdots + a_ku_n. \tag{12}$$

The recurrence (12) is said to have order k. Note that the same sequence can satisfy different recurrence relations, but it satisfies a unique recurrence of minimum order. The *characteristic polynomial* of the sequence $\langle u_n\rangle_{n=0}^{\infty}$ is

$$P(x) = x^k - a_1x^{k-1} - a_2x^{k-2} - \cdots - a_k$$

and its roots are called the *characteristic roots* of the sequence.

If $\boldsymbol{A} \in \mathbb{F}^{k\times k}$ is a square matrix and $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{F}^k$ are column vectors, then it can be shown that the sequence $u_n = \boldsymbol{v^T}\boldsymbol{A}^n\boldsymbol{w}$ satisfies a linear recurrence of order k. Indeed, by the Cayley-Hamilton Theorem, $\boldsymbol{A}$ satisfies its own characteristic equation $\det(\boldsymbol{A} - x\boldsymbol{I}) = 0$, which gives a recurrence relation on $\langle u_n\rangle_{n=0}^{\infty}$ with coefficients matching those of the characteristic polynomial $\det(\boldsymbol{A} - x\boldsymbol{I})$ of $\boldsymbol{A}$. Conversely, any LRS may be expressed in this way. Given a linear recurrence relation (12), it is sufficient to take $\boldsymbol{A}$ to be:

$$\boldsymbol{A} = \begin{bmatrix} a_1 & a_2 & \ldots & a_{k-1} & a_k \\ 1 & 0 & \ldots & 0 & 0 \\ 0 & 1 & \ldots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \ldots & 1 & 0 \end{bmatrix}.$$

Then if $\boldsymbol{w}$ is the vector $(u_{k-1}, \ldots, u_0)^T$ of initial terms of $\langle u_n\rangle_{n=0}^{\infty}$ in reverse order and $\boldsymbol{v}$ is the unit vector $(0, \ldots, 0, 1)^T$, we have $u_n = \boldsymbol{v^T}\boldsymbol{A}^n\boldsymbol{w}$. The characteristic polynomial of the LRS is the characteristic polynomial of $\boldsymbol{A}$, and the characteristic roots of the LRS are precisely the eigenvalues of $\boldsymbol{A}$.

By converting to the Jordan form, from the matrix expression $u_n = \boldsymbol{v^T}\boldsymbol{A}^n\boldsymbol{w}$ we can obtain a closed-form solution for the n-th term of the linear recurrence sequence in

terms of the eigenvalues $\lambda_1, \ldots, \lambda_l$ of $\boldsymbol{A}$:

$$u_n = \sum_{j=0}^{l} P_j(n)\lambda_j^n \tag{13}$$

for all $n \geq 0$, where $P_j \in \mathbb{F}[x]$ are univariate polynomials of degree strictly less than the multiplicity of λ_j as a root of the characteristic polynomial of $\boldsymbol{A}$. In the case $\mathbb{F} = \mathbb{R}$, the set of characteristic roots is closed under complex conjugation. Thus, if $\rho_1, \ldots, \rho_l \in \mathbb{R}$ are the real roots of $P(x)$ and $\gamma_1, \overline{\gamma}_1, \ldots, \gamma_m, \overline{\gamma}_m \in \mathbb{C}$ are the complex ones, the sequence is given by

$$u_n = \sum_{j=1}^{l} A_j(n)\rho_j^n + \sum_{j=1}^{m} \left(C_j(n)\gamma_j^n + \overline{C_j}(n)\overline{\gamma_j}^n\right) \tag{14}$$

for all $n \geq 0$, where $A_j \in \mathbb{R}[x]$ and $C_j \in \mathbb{C}[x]$. The coefficients of P_j in (13) and of A_j, C_j in (14) are algebraic numbers, effectively computable in polynomial time from the description of the LRS.

A linear recurrence sequence is called *degenerate* if for some pair of distinct characteristic roots λ_1, λ_2 of its minimum-order recurrence, the ratio λ_1/λ_2 is a root of unity, otherwise the sequence is *non-degenerate*. As pointed out in [Everest et al. 2003], the study of arbitrary LRS can be reduced effectively to that of non-degenerate LRS by partitioning the original LRS into finitely many non-degenerate subsequences. Specifically, for a given degenerate linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ with characteristic roots λ_j and matrix form $u_n = \boldsymbol{v}^{\boldsymbol{T}}\boldsymbol{A}^n\boldsymbol{w}$, let L be the least common multiple of the orders of all ratios λ_i/λ_j which are roots of unity. Then for each $j \in \{0, \ldots, L-1\}$, consider the sequence

$$u_n^{(j)} = \boldsymbol{v}^{\boldsymbol{T}}\boldsymbol{A}^{nL+j}\boldsymbol{w} = \boldsymbol{v}^{\boldsymbol{T}}(\boldsymbol{A}^L)^n(\boldsymbol{A}^j\boldsymbol{w}).$$

Each of these sequences has characteristic roots λ_i^L and is therefore non-degenerate. Indeed, if λ_1^L and λ_2^L are distinct eigenvalues of $\boldsymbol{A}^L$, and $(\lambda_1/\lambda_2)^L$ is a k-th root of unity for some k, then λ_1/λ_2 is an Lk-th root of unity, so $Lk \,|\, \mathrm{order}(\lambda_1/\lambda_2)$. But by the definition of L, we also have $\mathrm{order}(\lambda_1/\lambda_2) \,|\, L$, so $k = 1$. This yields $\lambda_1^L = \lambda_2^L$, a contradiction.

From the crude lower bound $\varphi(r) \geq \sqrt{r/2}$ on Euler's totient function, it follows that if $\alpha \in \mathbb{A}$ has degree d and is a primitive r-th root of unity, then $r \leq 2d^2$. There are $||\boldsymbol{A}||^{\mathcal{O}(1)}$ ratios λ_i/λ_j to consider, and if a ratio is a root of unity then its order is $||\boldsymbol{A}||^{\mathcal{O}(1)}$, so it follows that $L = 2^{\mathcal{O}(||\boldsymbol{A}||)}$. Thus, non-degeneracy can be ensured by considering at most exponentially many subsequences of the original LRS.

B. SKOLEM PROBLEM: INTRODUCTION

In the rest of the Appendix, we study the *Skolem Problem*: given a linear recurrence sequence (LRS) $\langle u_n \rangle_{n=0}^{\infty}$, determine whether there exists a natural number n such that $u_n = 0$. The sequence may be real- or complex-valued, but to make the problem well-defined, we shall require that the sequence be given in some effective form. For this reason, we take all the linear recurrence sequences to be over the algebraic numbers, at times restricting further to real-valued or rational sequences.

The connection to the Orbit Problem becomes evident if we recall the matrix representation of linear recurrence sequences. If $\langle u_n \rangle_{n=0}^{\infty}$ is given by $u_n = \boldsymbol{y}^{\boldsymbol{T}}\boldsymbol{A}^n\boldsymbol{x}$ for a $d \times d$ matrix $\boldsymbol{A}$ and vectors $\boldsymbol{x}$ and $\boldsymbol{y}$, then $u_n = 0$ if and only if $\boldsymbol{A}^n\boldsymbol{x} \in \{\boldsymbol{y}\}^{\perp}$. That is, the orbit of $\boldsymbol{x}$ under $\boldsymbol{A}$ intersects the $(d-1)$-dimensional hyperplane $\{\boldsymbol{y}\}^{\perp}$ if and only if the linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order d contains zero as an element.

The Skolem Problem has a history dating back to the 1930s, as evidenced by the celebrated Skolem-Mahler-Lech Theorem, a powerful result which characterises the zero sets of linear recurrence sequences:

THEOREM B.1. *(Skolem-Mahler-Lech) Let $\langle u_n \rangle_{n=0}^{\infty}$ be a linear recurrence sequence over a field with characteristic 0. Then the zero set of $\langle u_n \rangle_{n=0}^{\infty}$, $Z(u) = \{n \in \mathbb{N} : u_n = 0\}$, is semilinear, that is, the union of a finite set and finitely many arithmetic progressions.*

This result was originally established in the case of rational LRS in [Skolem 1934], then strengthened to include LRS over the algebraic numbers in [Mahler 1935], and finally extended to any field of characteristic 0 [Lech 1953; Mahler and Cassels 1956]. These proofs rely heavily on p-adic analysis and unfortunately do not yield constructive methods to compute the zero set of a given linear recurrence, nor to determine its emptiness. Nonetheless, later work [Berstel and Mignotte 1976] established an effective procedure to explicitly calculate the arithmetic progressions mentioned in the theorem for LRS over the rationals. This immediately renders decidable the problem of deciding finiteness of the zero set of a rational LRS. In a similar vein, it is also decidable whether the zero set of a rational LRS is equal to $\mathbb{N}$, and whether it has a finite complement [Salomaa and Soittola 1978, Section II.12].

Whilst the computation of the infinite component of the zero set is a significant advancement, no effective method is known to compute the finite component or to decide its emptiness. Thus, the decidability of the Skolem Problem remains open and is an outstanding question in number theory and theoretical computer science; see, for example, the exposition of [Tao 2008, Section 3.9]. Efforts towards an upper complexity bound have yielded only partial results: decidability for LRS over $\mathbb{A}$ of order at most 3 and for LRS over $\mathbb{R} \cap \mathbb{A}$ of order 4 in references [Vereshchagin 1985; Mignotte et al. 1984]. The decision method relies crucially on sophisticated results in transcendental number theory, specifically, Baker's lower bounds on the magnitudes of linear forms in logarithms of algebraic numbers and van der Poorten's analogous results in the context of p-adic valuations. Recently, a proof of decidability for LRS of order 5 was announced in [Halava et al. 2005]. However, as pointed out in [Ouaknine and Worrell 2012], the proof incorrectly addresses the case of LRS of the form:

$$u_n = A\lambda_1^n + \overline{A\lambda_1^n} + B\lambda_2^n + \overline{B\lambda_2^n} + Cr^n,$$

with one real and four complex characteristic roots with magnitudes satisfying $|\lambda_1| = |\lambda_2| > |r|$. Another paper [Litow 1997] claims decidability for all orders, but is also flawed [Ouaknine and Worrell 2012].

In terms of lower bounds, the strongest known result for the Skolem Problem is **NP**-hardness [Blondel and Portier 2002]. Reference [Litow 1997] claims **PSPACE**-hardness, but this has also been shown incorrect [Ouaknine and Worrell 2012].

C. SKOLEM PROBLEM: MAIN RESULT AND OUTLINE

The main technical result of the Appendix of this paper is the following:

THEOREM C.1. *Let $\langle u_n \rangle_{n=0}^{\infty}$ be a non-degenerate LRS of order d over $\mathbb{A}$ which is not identically zero and whose description has size $\|u\|$.*

(1) If $d = 2$, then there exists a bound $N = ||u||^{\mathcal{O}(1)}$ such that if $u_n = 0$, then $n < N$.
(2) If $d = 3$, then there exists a bound $N = 2^{\mathcal{O}(||u||)}$ such that if $u_n = 0$, then $n < N$.
(3) If $d = 4$ and $\langle u_n \rangle_{n=0}^{\infty}$ is over $\mathbb{R} \cap \mathbb{A}$, then there exists a bound $N = 2^{\mathcal{O}(||u||)}$ such that if $u_n = 0$, then $n < N$.

References [Mignotte et al. 1984; Vereshchagin 1985] show the existence of similar bounds, but make no attempt to quantify them in terms of the description of the input,

thereby showing the problems decidable, but obtaining no more specific complexity upper bound. Our contribution is to show the bounds are at most exponential in the size of the input, and in fact, polynomial for LRS of order 2. This permits us to obtain the following complexity bounds for the Skolem Problem for rational LRS:

THEOREM C.2. *For LRS over $\mathbb{Q}$ of order at most 4, the Skolem Problem is in the complexity class $\mathbf{NP}^{\mathbf{RP}}$. Further, for LRS over $\mathbb{Q}$ of order 2, the problem is in $\mathbf{PTIME}$.*

Two points need to be addressed: how to reduce from arbitrary LRS to non-degenerate, non-zero LRS, and how to obtain the complexity results of Theorem C.2 from the bounds of Theorem C.1.

On the first point, as we showed in Section A.5, the study of arbitrary LRS can be reduced effectively to the non-degenerate case. This uses the technique of partitioning a given LRS into L non-degenerate subsequences, where

$$L = \operatorname{lcm}\{\operatorname{order}(\lambda_i/\lambda_j) \,:\, \lambda_i, \lambda_j \text{ characteristic roots and } \lambda_i/\lambda_j \text{ root of unity }\}. \tag{15}$$

Specifically, if $\langle u_n\rangle_{n=0}^{\infty}$ is given, then we consider the sequences $\langle v_n^{(j)}\rangle_{n=0}^{\infty}$ defined by $v_n^{(j)} = u_{Ln+j}$ for $j = 0, \dots, L-1$. These subsequences are non-degenerate, so for the purposes of showing decidability, non-degeneracy may be assumed without loss of generality. However, when attempting to establish a more precise complexity upper bound, the size of L needs to be taken into account.

Recall that if α is an algebraic number of degree d and a root of unity of order r, then $r \leq 2d^2$. In particular, if $\langle u_n\rangle_{n=0}^{\infty}$ is an LRS defined by $u_n = \boldsymbol{x}^{\boldsymbol{T}}\boldsymbol{A}^n\boldsymbol{y}$ described using $\|u\| = \|\boldsymbol{x}\| + \|\boldsymbol{A}\| + \|\boldsymbol{y}\|$ bits, and λ_i, λ_j are characteristic roots such that λ_i/λ_j is a root of unity, then $\operatorname{order}(\lambda_i/\lambda_j) = ||u||^{\mathcal{O}(1)}$. Since (15) takes the least common multiple of the orders of $\mathcal{O}(\|u\|^2)$ ratios λ_i/λ_j and each order is polynomially large in the size of the input, it follows that $L = 2^{\mathcal{O}(||u||)}$. Moreover, this is not an over-approximation: it is easy to construct LRS where the ratios λ_i/λ_j are roots of unity of mutually coprime orders, thereby making L at least exponential in the size of the input. Thus, for arbitrary LRS, applying this technique to eliminate non-degeneracy carries an exponential overhead.

However, in this paper, we restrict our attention to LRS of order at most 4. Therefore, in (15), the number of ratios considered is bounded by an absolute constant, so L is the least common multiple of a fixed number of polynomially large orders, hence $L = ||u||^{\mathcal{O}(1)}$.

Furthermore, if the LRS is over $\mathbb{Q}$, then the degree of each characteristic root λ_i is at most 4, since we know *a priori* that the characteristic polynomial of the sequence has rational coefficients. Then the degrees of all ratios λ_i/λ_j are also absolutely bounded, so $L = \mathcal{O}(1)$. Therefore, in our context of rational LRS of bounded order, non-degeneracy may be obtained by considering a constant number of subsequences, whose matrix representation may be computed from that of $\langle u_n\rangle_{n=0}^{\infty}$ in polynomial time.

Some of the non-degenerate subsequences $\langle v_n^{(j)}\rangle_{n=0}^{\infty}$ could potentially be identically zero, resulting in the zero set of the original degenerate sequence $\langle u_n\rangle_{n=0}^{\infty}$ containing an entire arithmetic progression. For each subsequence, we check directly whether this is the case by examining its first d terms, and if so, then the problem instance is immediately positive. Otherwise, the instance has been reduced to a polynomially-large (or constant, in the rational case) number of problem instances, each featuring a non-degenerate LRS which is not the zero sequence.

The second point is how to obtain the complexity upper bounds. For non-degenerate rational LRS $\langle u_n\rangle_{n=0}^{\infty}$ defined by $u_n = \boldsymbol{x}^{\boldsymbol{T}}\boldsymbol{A}^n\boldsymbol{y}$, let N denote the bound provided by Theorem C.1. If the sequence is of order 2, then N is only polynomial in the size of the input. Thus, we simply calculate u_n for all $n < N$. All intermediate results are rational

numbers, and we only ever raise $\boldsymbol{A}$ to a polynomially large power, so the representation of all intermediate results stays polynomially bounded. Thus, the **PTIME** upper bound for LRS of order 2 follows directly.

For rational LRS of order 3 or 4, the bound N is at most exponential in the size of the input. We argue the problem is in $\mathbf{NP}^{\mathbf{EqSLP}}$, where **EqSLP** is the complete class for the following problem: given a division-free straight-line program (or equivalently, an arithmetic circuit) producing an integer M, determine whether $M = 0$. Since the bound N is at most exponentially large in the size of the input, an **NP** algorithm can guess the index of a purported zero: $n \in \mathbb{N}$ with $n < N$. Thus, we only need to verify that $u_n = 0$. Direct calculation is not an option, since n is exponential in the size of the input, whilst the entries of $\boldsymbol{A}^n$ are doubly-exponential in magnitude, requiring an exponential number of bits to write down. However, we can easily represent the entries of $\boldsymbol{A}^n$ as polynomially-sized arithmetic circuits, using the technique of repeated squaring. Then verifying $u_n = 0$ reduces to checking whether a polynomially-large arithmetic circuit evaluates to 0, which can be solved by an **EqSLP** oracle. The bound $\mathbf{NP}^{\mathbf{EqSLP}}$ follows directly. Finally, it is known that $\mathbf{EqSLP} \subseteq \mathbf{coRP}$ [Schönhage 1979], so we also have membership in $\mathbf{NP}^{\mathbf{RP}}$, as Theorem C.2 claims.

Therefore, all that remains is to prove Theorem C.1. We devote the rest of the Appendix to the technical details of the proof. Sections D, F and G address LRS of order 2, 3 and 4, respectively. Section E shows two applications of Baker's Theorem which are crucially important for orders 3 and 4.

D. LRS OF ORDER TWO

In this section, we consider the problem of whether a linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order 2 over $\mathbb{A}$ contains zero as a term. The characteristic equation of the recurrence may have one repeated root θ, or two distinct roots θ_1, θ_2. Thus, the n-th term of the sequence is given by one of the following:

$$u_n = (A + Bn)\theta^n \qquad \text{(where } A, B, \theta \in \mathbb{A} \text{ and } B, \theta \neq 0\text{)} \tag{16}$$

$$u_n = A\theta_1^n + B\theta_2^n \qquad \text{(where } A, B, \theta_1, \theta_2 \in \mathbb{A} \text{ and } A, B, \theta_1, \theta_2 \neq 0\text{)} \tag{17}$$

Solving the Skolem Problem for LRS of the form (16) is trivial: simply determine whether the unique root of $A + Bx$ is a natural number. We therefore concentrate on LRS of the form (17). In this case, $u_n = 0$ if and only if $(\theta_1/\theta_2)^n = -B/A$.

Thus, the problem reduces to the *algebraic number power problem:* decide whether there exists $n \in \mathbb{N}$ such that

$$\alpha^n = \beta \tag{18}$$

for given $\alpha, \beta \in \mathbb{A}$. The assumption of non-degeneracy of $\langle u_n \rangle_{n=0}^{\infty}$ allows us to assume α is not a root of unity[3]. The algebraic number power problem is decidable [Halava et al. 2005]. Reference [Kannan and Lipton 1986] proved a polynomial bound on n when β has the form $P(\alpha)$ for a given $P \in \mathbb{Q}[x]$. We give a brief recapitulation of the decidability proof of [Halava et al. 2005] and sharpen it to extract a polynomial bound on n.

LEMMA D.1. *Suppose $\alpha, \beta \in \mathbb{A}$. If α is not a root of unity, then there exists a bound N such that if (18) holds, then $n < N$. Moreover, $N = ||I||^{\mathcal{O}(1)}$, where $\|I\| = \|\alpha\| + \|\beta\|$ is the length of the input.*

[3]Notice in passing that if α is a root of unity, then the algebraic number power problem is easy to decide: simply determine whether β is an r-th root of unity, where r is the order of α. If this is indeed the case, however, then there exists no bound of the kind promised by Theorem C.1, since $\alpha^n = \beta$ holds periodically.

PROOF. Let $\mathbb{K} = \mathbb{Q}(\alpha, \beta)$. If α is not an algebraic integer, then by Lemma A.2 there exists a prime ideal P in the ring $\mathcal{O}_{\mathbb{K}}$ such that $v_P(\alpha) \neq 0$. Then if $\alpha^n = \beta$, we have

$$v_P(\alpha^n) = nv_P(\alpha) = v_P(\beta).$$

If $v_P(\alpha)$ and $v_P(\beta)$ have different signs, then we are done. Otherwise,

$$n = \frac{v_P(\beta)}{v_P(\alpha)} \leq |v_P(\beta)| \leq \log_2 |\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\beta)| \leq \log_2 |\mathcal{N}_{abs}(\beta)|^d,$$

where $d = [\mathbb{Q}(\alpha, \beta) : \mathbb{Q}]$ is at most polynomially large in $\|\alpha\| + \|\beta\|$. It follows that the bound on n is polynomially large in the length of the input.

Suppose therefore that α is an algebraic integer. It is not a root of unity by the premise of the Lemma, so by Theorem A.6 (Blanksby and Montgomery), α has a Galois conjugate $\sigma(\alpha)$ such that

$$|\sigma(\alpha)| > 1 + \frac{1}{30d^2 \log(6d)},$$

where d is the degree of α. This implies

$$\frac{1}{\log |\sigma(\alpha)|} < 60d^2 \log(6d).$$

Recall that σ is a monomorphism on $\mathbb{K}$, so $\sigma(\alpha^n) = (\sigma(\alpha))^n$. Then if $\alpha^n = \beta$, we have

$$n = \frac{\log |\sigma(\beta)|}{\log |\sigma(\alpha)|} < \log |\sigma(\beta)| 60d^2 \log(6d).$$

Observe that if we are given canonical descriptions of α and β, then $60d^2 \log(6d)$ is at most polynomially large in $\|\alpha\|$, and $\log |\sigma(\beta)|$ is at most polynomially large in $\|\beta\|$. It follows that the bound on n is polynomial in the length of the input. □

E. APPLICATION OF BAKER'S THEOREM

Before we proceed to LRS of order 3 and 4, we make a brief diversion to show two pertinent applications of Baker's Theorem. They essentially capture the technically difficult core of the Skolem Problem for LRS of order 3 or 4, so for clarity, they are exhibited here first, prior to their use in the context of LRS.

The first application concerns powers λ^n ($n \in \mathbb{N}$) of an algebraic number λ on the unit circle. We show that for large n and any fixed $b \in \mathbb{A}$, the distance $|\lambda^n - b|$ cannot be 'too small', unless λ is a root of unity.

LEMMA E.1. *Let $\lambda, b \in \mathbb{A}$, where $|\lambda| = 1$ and λ is not a root of unity. Suppose $\phi(n)$ is a function from $\mathbb{N}$ to $\mathbb{C}$ for which there exist $a, \chi \in \mathbb{Q}$ such that $\chi \in (0, 1)$ and $|\phi(n)| \leq a\chi^n$. There exists a bound N such that if*

$$\lambda^n = \phi(n) + b, \tag{19}$$

then $n < N$. Moreover, N is at most exponential in the length of the input $\|I\| = \|\lambda\| + \|b\| + \|a\| + \|\chi\|$.

PROOF. The left-hand side of (19) describes points on the unit circle, whereas the right-hand side tends to b. If $|b| \neq 1$, then for n large enough, the right-hand side of (19) will always be off the unit circle. This happens when

$$n > \frac{\log(||b| - 1|/a)}{\log(\chi)}.$$

The difficult case is when b is on the unit circle. Here we will use Baker's Theorem to derive a bound on n. Consider the angle Λ between λ^n and b. Since λ is not a root

of unity, by Lemma D.1, this angle can be zero for at most one value of n, which is polynomially large in $\|I\|$. Otherwise, write

$$\Lambda = \log \frac{\lambda^n}{b} = n\log(\lambda) - \log(b) + 2k_n \log(-1) \neq 0,$$

where k_n is an integer chosen so that $\Lambda = i\tau$ for some $\tau \in [0, 2\pi)$. Then $2n$ is an upper bound on the height of the coefficients in front of the logarithms (because $k_n \leq n$), $H = \max\{H_\lambda, H_b, 3\}$ is a height bound for the arguments to the logarithms and $d = \max\{\deg(\lambda), \deg(b)\}$ is a bound on the degrees. Then by Theorem A.4 (Baker-Wüstholz), we have

$$\log|\Lambda| > -(48d)^{10}\log^2 H \log(2n),$$

which is equivalent to

$$|\Lambda| > (2n)^{-(48d)^{10}\log^2 H}.$$

This is a lower bound on the length of the arc between λ^n and b. The length of the chord is at least half of the bound: $|\lambda^n - b| \geq |\Lambda|/2$. So in the equation $\lambda^n - b = \phi(n)$, the left-hand side is bounded below by an inverse polynomial in n. However, the right-hand side shrinks exponentially quickly in n. For all n large enough, the right-hand side will be smaller in magnitude than the left-hand side.

We will now quantify the bound on n. Let $p_1 = (48d)^{10}\log^2 H$ and $p_2 = 2$. Observe that $p_1, p_2 = ||I||^{\mathcal{O}(1)}$. Then (19) cannot hold if

$$\frac{1}{2}(p_2 n)^{-p_1} \geq a\chi^n,$$

which is equivalent to

$$-\log(2) - \log(a) - p_1\log(p_2) - p_1\log(n) \geq n\log(\chi).$$

Define $p_3 = \log(2) + \log(a) + p_1\log(p_2)$ and $p_4 = \max\{p_3, p_1\} = ||I||^{\mathcal{O}(1)}$. Then it suffices to have

$$\frac{p_4}{-\log(\chi)} \leq \frac{n}{1+\log(n)},$$

which is guaranteed by

$$\sqrt{n} \geq \frac{p_4}{-\log(\chi)}.$$

Observe that $-1/\log(\chi)$ is at most exponentially large in $\|\chi\|$. Therefore, the bound on n is exponential in the size of the input. □

Continuing in the same line, we next consider two algebraic numbers, λ_1 and λ_2, whose powers define discrete trajectories embedded in two circles in the complex plane: $a\lambda_1^n$ and $b\lambda_2^n + c$ as n varies over $\mathbb{N}$. The following lemma shows that unless λ_1, λ_2 are roots of unity, then for large n, the n-th points of the two trajectories are never 'too close' to each other.

LEMMA E.2. *Suppose $\lambda_1, \lambda_2, a, b, c \in \mathbb{A}$ are non-zero, where $|\lambda_1| = |\lambda_2| = 1$ and λ_1, λ_2 are not roots of unity. Let $\phi(n)$ be a function from $\mathbb{N}$ to $\mathbb{C}$ such that $0 < |\phi(n)| \leq w\chi^n$ for some $w, \chi \in \mathbb{Q}$, $\chi \in (0,1)$. Then there exists a bound N such that if*

$$a\lambda_1^n = b\lambda_2^n + c + \phi(n), \tag{20}$$

then $n < N$. Moreover, $N = 2^{\mathcal{O}(||I||)}$, where $\|I\| = \|\lambda_1\| + \|\lambda_2\| + \|a\| + \|b\| + \|c\| + \|w\| + \|\chi\|$.

PROOF. Multiplying the equation by $\overline{c}/|c||a|$ allows us to assume that $|a| = 1$ and $c \in \mathbb{R}^+$.

Let $f(n) = a\lambda_1^n$, $g(n) = b\lambda_2^n + c$. It is clear that $f(n)$ describes points on the unit circle $\mathcal{O}_1$, whilst $g(n)$ describes points on the circle $\mathcal{O}_2$ with centre c on the real line and radius $|b|$.

If these circles do not intersect, then for n large enough, $|\phi(n)|$ will be forever smaller than the smallest distance between the circles. This happens when

$$n > \frac{\log(c - |b| - 1) - \log(w)}{\log(\chi)},$$

which is an exponential lower bound on n in the size of the input.

Suppose now the circles intersect in two points, z_1 and z_2. Let $\mathcal{L}_1$ be the horizontal line through z_1 and $\mathcal{L}_2$ the horizontal line through z_2. Let $\mathcal{L}_1 \cap \mathcal{O}_1 = \{x_1, z_1\}$, $\mathcal{L}_1 \cap \mathcal{O}_2 = \{y_1, z_1\}$, $\mathcal{L}_2 \cap \mathcal{O}_1 = \{x_2, z_2\}$ and $\mathcal{L}_2 \cap \mathcal{O}_2 = \{y_2, z_2\}$. It is trivial that $z_2 = \overline{z_1}$, $x_2 = \overline{x_1}$, $y_2 = \overline{y_1}$.

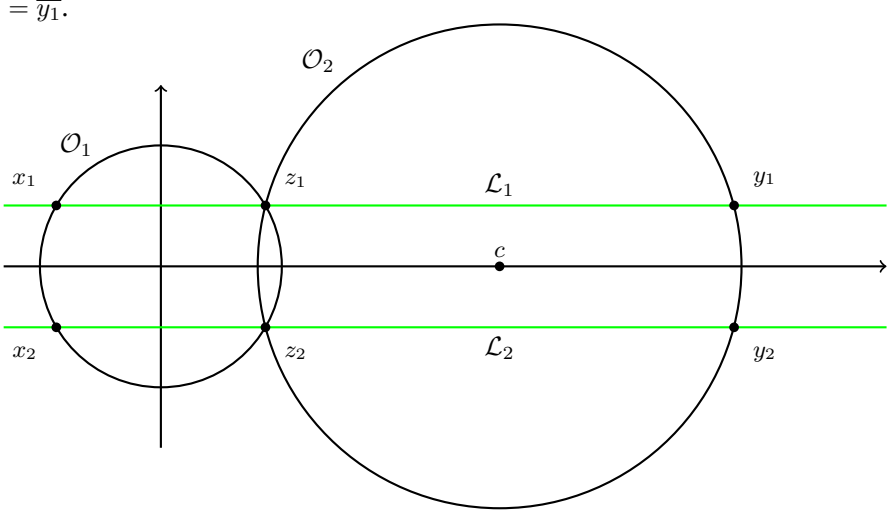

We first argue that for n large enough, (20) can hold only if for some intersection point z_i, $\Re(z_i)$ lies between $\Re(f(n))$ and $\Re(g(n))$, or $\Im(z_i)$ lies between $\Im(f(n))$ and $\Im(g(n))$. This can only be violated in two symmetric situations: either

(1) $f(n)$ is on the arc z_1z_2 of $\mathcal{O}_1$ which lies inside $\mathcal{O}_2$ and $g(n)$ is on the arc y_1y_2 of $\mathcal{O}_2$ which lies outside $\mathcal{O}_1$, or
(2) $f(n)$ is on the arc x_1x_2 of $\mathcal{O}_1$ which lies outside $\mathcal{O}_2$ and $g(n)$ is on the arc z_1z_2 of $\mathcal{O}_2$ which lies inside $\mathcal{O}_1$.

In the first situation, when $g(n)$ is on the arc y_1y_2 of $\mathcal{O}_2$ outside $\mathcal{O}_1$, we have

$$|f(n) - g(n)| \geq |g(n)| - 1 \geq |y_1| - 1.$$

Since the point y_1 is strictly to the right of 1 on the complex plane, this lower bound is positive, and moreover it is independent of n, so (20) cannot hold for n large enough because $\phi(n)$ tends to zero exponentially quickly. In particular, (20) does not hold if

$$n > \frac{\log(|y_1| - 1) - \log(w)}{\log(\chi)},$$

which is exponentially large in the size of the input. The second situation is analogous.

Therefore, we can assume that one of the intersection points z_i separates $f(n)$ and $g(n)$ horizontally or vertically in the figure. That is, z_i satisfies $\Re(f(n)) \leq \Re(z_i) \leq$

$\Re(g(n))$ or $\Im(f(n)) \leq \Im(z_i) \leq \Im(g(n))$. We will show a lower bound on $|f(n)-g(n)|$ which shrinks slower than exponentially. The real (horizontal) and imaginary (vertical) cases are completely analogous. We show the working for the real case. Assume that $\Re(z_i)$ lies between $\Re(f(n))$ and $\Re(g(n))$. Clearly,

$$|f(n) - g(n)| \geq |\Re(g(n) - f(n))| = |\Re(z_i - f(n))| + |\Re(g(n) - z_i)|.$$

Let $\alpha = \arg(\lambda_1)$, $\gamma = \arg(a)$ and $\beta = \arg(z_i)$. Then

$$|\Re(z_i - f(n))| = |\cos(n\alpha + \gamma) - \cos(\beta)| = 2\left|\sin\frac{\beta - n\alpha - \gamma}{2}\sin\frac{\beta + n\alpha + \gamma}{2}\right|.$$

Let u_n, v_n be appropriately chosen integers so that

$$\frac{\beta - n\alpha - \gamma}{2} + u_n\pi \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right],$$

$$\frac{\beta + n\alpha + \gamma}{2} + v_n\pi \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right].$$

Then using the inequality

$$|\sin(x)| \geq \frac{|x|}{\pi} \text{ for } x \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right],$$

we have

$$\left|\sin\frac{\beta - n\alpha - \gamma}{2}\right| \geq \frac{1}{\pi}\left|\frac{\beta - n\alpha - \gamma}{2} + \pi u_n\right|,$$

$$\left|\sin\frac{\beta + n\alpha + \gamma}{2}\right| \geq \frac{1}{\pi}\left|\frac{\beta + n\alpha + \gamma}{2} + \pi v_n\right|.$$

Both of these expressions are sums of logarithms of algebraic numbers, non-zero for n exceeding a polynomially large bound in $\|I\|$ by Lemma D.1, so we can give lower bounds for them using Theorem A.4 (Baker-Wüstholz) as in Lemma E.1:

$$|\Re(z_i - f(n))| \geq (p_1 n)^{-p_2}$$

for some $p_1, p_2 > 0$ which are independent of n and at most polynomially large in the input. A similar lower bound holds for $|\Re(g(n) - z_i)|$. If $\delta = \arg(\lambda_2)$, $\eta = \arg(b)$ and $\theta = \arg(z_i - c)$, we have

$$|\Re(g(n) - z_i)| = |b|(\cos(n\delta + \eta) - \cos(\theta)) \geq (p_3 n)^{-p_4},$$

where $p_3, p_4 > 0$ are independent of n and have at most polynomial size in the input. Hence we have

$$|f(n) - g(n)| \geq 2(p_5 n)^{-p_6},$$

where $p_5 = \max\{p_1, p_3\}$ and $p_6 = \max\{p_2, p_4\}$. Since $\phi(n)$ shrinks exponentially quickly, a bound on n follows past which (20) cannot hold. In the manner of Lemma E.1, we can show that this bound is exponentially large in the size of the input. The vertical case is analogous, except that considering imaginary parts gives sines instead of cosines, so we shift all angles by $\pi/2$ and proceed as above. If the circles are tangent and neither lies inside the other, then the intersection point separates $f(n)$ and $g(n)$ horizontally, so we are done by the above analysis.

Finally, suppose that the circles are tangent and one lies inside the other: $|b| + c = 1$. The argument of $f(n)$ is $\gamma + n\alpha$. By the law of cosines applied to the triangle with vertices $f(n)$ and the centres of the circles, we have

$$|f(n) - c|^2 = c^2 + 1 - 2c\cos(\gamma + n\alpha).$$

Therefore, the shortest distance from $f(n)$ to a point on $\mathcal{O}_2$ is

$$h(n) = \sqrt{c^2 + 1 - 2c\cos(\gamma + n\alpha)} - (1 - c).$$

Let $A(n) = \sqrt{c^2 + 1 - 2c\cos(\gamma + n\alpha)}$ and $B = 1 - c$. Since $A \leq 1 + c$, we have $A + B \leq 2$, so

$$h(n) = A - B = \frac{A^2 - B^2}{A + B} \geq c(1 - \cos(\gamma + n\alpha))$$

Let k_n be an integer, so that

$$\gamma + n\alpha + k_n 2\pi \in [-\pi, \pi).$$

By Lemma D.1, this is zero for at most one, polynomially large in $\|I\|$, value of n. For larger n, a lower bound on this angle follows from Theorem A.4 (Baker-Wüstholz):

$$|\gamma + n\alpha + k_n 2\pi| \geq (p_7 n)^{-p_8}$$

for some constants $p_7, p_8 > 0$ which are polynomially large in the input. Then

$$\cos(\gamma + n\alpha) \leq \cos((p_7 n)^{-p_8}),$$

so

$$h(n) \geq c(1 - \cos((p_7 n)^{-p_8})).$$

From the Taylor expansion of $\cos(x)$, it follows easily that

$$1 - \cos(x) \geq \frac{11}{24}x^2 \text{ for } x \leq 1.$$

Since $p_7, p_8 \geq 1$, we have $(p_7 n)^{-p_8} \leq 1$. Therefore,

$$h(n) \geq c\frac{11}{24}(p_7 n)^{-2p_8}.$$

This lower bound on $h(n)$ shrinks inverse-polynomially as n grows. Recall that $h(n)$ is the smallest distance from $f(n)$ to $\mathcal{O}_2$. It follows that for n large enough, $|\phi(n)| < h(n)$ forever, so $f(n) = g(n) + \phi(n)$ cannot hold. In the manner of Lemma E.1, we can show that the bound on n is exponentially large in the input. □

F. LRS OF ORDER THREE

We now move to the problem of determining whether a linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order 3 over $\mathbb{A}$ contains zero as an element. The characteristic equation of such a sequence may have either three distinct (real or complex) roots, or one repeated real root and one simple real root, or one real root of multiplicity 3. Thus, the n-th element of the sequence is given by one of the following:

$$u_n = A\alpha^n + B\beta^n + C\gamma^n \qquad \text{(where } A, B, C, \alpha, \beta, \gamma \in \mathbb{A} \text{ are all non-zero)} \tag{21}$$

$$u_n = (A + Bn)\alpha^n + C\beta^n \qquad \text{(where } A, B, C, \alpha, \beta \in \mathbb{A} \text{ with } B, C, \alpha, \beta \neq 0) \tag{22}$$

$$u_n = (Cn^2 + Bn + A)\alpha^n \qquad \text{(where } A, B, C, \alpha \in \mathbb{A} \text{ with } C, \alpha \neq 0) \tag{23}$$

Finding the zeros of LRS of the form (23) is trivial: simply check whether the quadratic $Cn^2 + Bn + A$ has roots which are natural numbers. Thus, we focus on the remaining two possibilities. We will consider only non-degenerate sequences: the ratios of the roots α, β, γ are not roots of unity.

First we consider $\langle u_n \rangle_{n=0}^{\infty}$ given by (21). Notice that $A, B, C, \alpha, \beta, \gamma$ are all non-zero, otherwise the sequence satisfies a recurrence relation of lower order. Thus, we can

rearrange $u_n = 0$ to obtain:

$$\left(\frac{\beta}{\alpha}\right)^n = -\frac{C}{B}\left(\frac{\gamma}{\alpha}\right)^n - \frac{A}{B}. \tag{24}$$

Assume without loss of generality $|\alpha| \geq |\beta| \geq |\gamma|$. In Lemmas F.1, F.2, F.3 below, we consider separately the cases $|\alpha| > |\beta|$, $|\alpha| = |\beta| > |\gamma|$ and $|\alpha| = |\beta| = |\gamma|$, and obtain a bound on n which is exponential in the length of the description of the sequence and beyond which $u_n = 0$ cannot hold.

LEMMA F.1. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is given by (21). If $|\alpha| > |\beta|$, then there exists a bound N such that if $u_n = 0$, then $n < N$. Moreover, $N = 2^{\mathcal{O}(||I||)}$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\| + \|\gamma\|$.*

PROOF. This follows straightforwardly from the dominance of α. If

$$n > \max\left\{\frac{\log|A/2B|}{\log|\beta/\alpha|}, \frac{\log|A/2C|}{\log|\gamma/\alpha|}\right\},$$

then

$$\left|-\frac{B}{A}\left(\frac{\beta}{\alpha}\right)^n - \frac{C}{A}\left(\frac{\gamma}{\alpha}\right)^n\right| \leq \left|\frac{B}{A}\left(\frac{\beta}{\alpha}\right)^n\right| + \left|\frac{C}{A}\left(\frac{\gamma}{\alpha}\right)^n\right| < \frac{1}{2} + \frac{1}{2} = 1.$$

□

LEMMA F.2. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is given by (21). If $|\alpha| = |\beta| > |\gamma|$, then there exists a bound N such that if $u_n = 0$, then $n < N$. Moreover, $N = 2^{\mathcal{O}(||I||)}$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\| + \|\gamma\|$.*

PROOF. This is a direct application of Lemma E.1 to equation (24). □

LEMMA F.3. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is given by (21). If $|\alpha| = |\beta| = |\gamma|$, there exist at most two values of n such that $u_n = 0$. Moreover, they are at most exponential in the length of the input $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\| + \|\gamma\|$ and are computable in polynomial time.*

PROOF. The left-hand side of (24) as a function of n describes points on the unit circle in the complex plane, whereas the right-hand side describes points on a circle centred at $-A/B$ with radius $|C/B|$. Note these circles do not coincide, because $A \neq 0$. We can obtain their equations and compute their intersection point(s). If they do not intersect, then equation (24) can never hold. Otherwise, the equation can only hold if the two sides are simultaneously equal to the same intersection point. For each of the (at most two) intersection points θ, let

$$S_1 = \left\{n \;\middle|\; \left(\frac{\beta}{\alpha}\right)^n = \theta\right\},$$

$$S_2 = \left\{n \;\middle|\; -\frac{C}{B}\left(\frac{\gamma}{\alpha}\right)^n - \frac{A}{B} = \theta\right\}.$$

Observe that $|S_i| \leq 1$, because β/α and γ/α are not roots of unity. We compute S_1 and S_2 from the bound in Lemma D.1 and check whether $S_1 \cap S_2$ is non-empty. □

Next, we consider LRS of the form (22). We will assume that B, C, α, β are all non-zero, otherwise the sequence satisfies a linear recurrence of lower order.

LEMMA F.4. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is given by (22). There exists a bound N such that if $u_n = 0$, then $n < N$. Moreover, $N = 2^{\mathcal{O}(||I||)}$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\|$.*

PROOF. We wish to solve for $n \in \mathbb{N}$ the equation:

$$(A + Bn)\alpha^n + C\beta^n = 0. \tag{25}$$

If $|\alpha| \geq |\beta|$, then for

$$n > \frac{|A| + |C|}{|B|},$$

we have

$$|C| < |B|n - |A| \leq |A + Bn|,$$

so

$$|C\beta^n| < |(A + Bn)\alpha^n|,$$

therefore (25) cannot hold. Now suppose $|\alpha| > |\beta|$ and rewrite (25) as

$$\frac{A + Bn}{C} = -\left(\frac{\beta}{\alpha}\right)^n.$$

Equation (25) implies

$$\left|\frac{\beta}{\alpha}\right|^n = \left|\frac{A + Bn}{C}\right| \leq \left|\frac{A}{C}\right| + \left|\frac{B}{C}\right| n.$$

However, we will show that for all n large enough, this fails to hold. Indeed, the inequality

$$\left|\frac{\beta}{\alpha}\right|^n > \left|\frac{A}{C}\right| + \left|\frac{B}{C}\right| n$$

is implied by

$$d\,(n+1) < \left|\frac{\beta}{\alpha}\right|^n,$$

where $d = \max\{|A/C|, |B/C|\}$. Taking logarithms, we see that it suffices to have

$$\frac{n}{1 + \log(n+1)} > \frac{f}{\log|\beta/\alpha|},$$

where $f = \max\{\log(d), 1\}$. Noting that $1 + \log(n+1) < 2\sqrt{n}$ for all $n \geq 1$, we see that it suffices to have

$$n > 4f^2/\log^2|\beta/\alpha|$$

to guarantee that (25) cannot hold. This is an exponential bound on n in the length of the input. □

G. LRS OF ORDER FOUR

We now proceed to the problem of determining whether a linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order 4 over $\mathbb{A}$ contains zero as an element. As before, we assume non-degeneracy of the sequence. Depending on the roots of the characteristic polynomial, the n-th term of the sequence is given by one of the following (where $A, B, C, D, \alpha, \beta, \gamma, \delta$ are algebraic):

$$u_n = A\alpha^n + B\beta^n + C\gamma^n + D\delta^n \qquad \text{(where } A, B, C, D \neq 0\text{)} \tag{26}$$

$$u_n = (A + Bn)\alpha^n + C\beta^n + D\gamma^n \qquad \text{(where } B, C, D \neq 0\text{)} \tag{27}$$

$$u_n = (A + Bn)\alpha^n + (C + Dn)\beta^n \qquad \text{(where } B, D \neq 0\text{)} \tag{28}$$

$$u_n = (A + Bn + Cn^2)\alpha^n + D\beta^n \qquad \text{(where } C, D \neq 0\text{)} \tag{29}$$

$$u_n = (A + Bn + Cn^2 + Dn^3)\alpha^n \qquad \text{(where } D \neq 0\text{)} \tag{30}$$

Solving $u_n = 0$ in the case of $\langle u_n \rangle_{n=0}^{\infty}$ given by (30) is trivial: just calculate canonical descriptions of the roots of $A + Bx + Cx^2 + Dx^3$ and check whether any are natural numbers.

In the case of $\langle u_n \rangle_{n=0}^{\infty}$ given by (29), rearrange $u_n = 0$ as

$$(A + Bn + Cn^2)\left(\frac{\alpha}{\beta}\right)^n = -D.$$

The left-hand side tends to 0 or ∞ in magnitude, depending on whether $|\alpha| < |\beta|$. In both cases, since $C, D \neq 0$, a bound on n follows which is at most exponential in the size of the input.

The remaining three cases, where $\langle u_n \rangle_{n=0}^{\infty}$ is of the form (26), (27) or (28) are more involved. They are the subject of Lemmas G.1, G.2, G.3 and G.4, which show the existence of a bound N which is at most exponentially large in the size of the input and beyond which $u_n = 0$ cannot hold.

Note that for complex algebraic LRS given by (26) with characteristic roots all of the same magnitude, the Skolem Problem is not known to be decidable. Thus, our final technical result, Lemma G.4 will require the simplifying assumption that $u_n \in \mathbb{R} \cap \mathbb{A}$ for all n. This is the only reason why Theorem C.1 insists that LRS of order 4 be real algebraic. In all other cases, as shown by Lemmas G.1, G.2 and G.3, an exponential bound on n exists even for complex algebraic LRS.

LEMMA G.1. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is non-degenerate and is given by (28). There exists a bound $N = 2^{\mathcal{O}(||I||)}$ such that if $u_n = 0$, then $n < N$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|D\| + \|\alpha\| + \|\beta\|$.*

PROOF. We wish to solve for $n \in \mathbb{N}$ the equation:

$$(A + Bn)\alpha^n + (C + Dn)\beta^n = 0 \text{ (where } B, D \neq 0). \tag{31}$$

Rearrange (31) as

$$\lambda^n = -\frac{(C + Dn)}{(A + Bn)}, \tag{32}$$

where $\lambda = \alpha/\beta$ is not a root of unity. The right-hand side of (32) tends to $-D/B$ as n tends to infinity.

If λ is an algebraic integer, then by Theorem A.6 (Blanksby and Montgomery), it has a Galois conjugate $\sigma(\lambda)$ such that

$$|\sigma(\lambda)| > 1 + \frac{1}{30d^2 \log(6d)},$$

where d is the degree of λ. Assume the monomorphism σ has been applied to both sides of (32), so $|\lambda|$ is bounded away from 1 by an inverse polynomial in the size of the input. By the triangle inequality, if

$$n \geq \frac{|BC| + |AD| + |AB|}{|B|^2} \stackrel{\text{def}}{=} N_1 = 2^{\mathcal{O}(||I||)},$$

then

$$\left|\frac{C + Dn}{A + Bn}\right| \leq \frac{|D|n + |C|}{|B|n - |A|} \leq \left|\frac{D}{B}\right| + 1.$$

Following the reasoning of Lemma D.1 and relying on the Blansky and Montgomery bound, we see there exists a bound $N_2 \in ||I||^{\mathcal{O}(1)}$ such that if $n > N_2$, then $|\lambda^n| > |D/B| + 1$. Therefore, for $n > \max\{N_1, N_2\} = 2^{\mathcal{O}(||I||)}$, equation (32) cannot hold.

Second, suppose λ is not an algebraic integer. Then by Lemma A.2 there exists a prime ideal P in the ring of integers of $\mathbb{K} = \mathbb{Q}(\alpha, \beta, A, B, C, D)$ such that $v_P(\lambda) \neq 0$. Without loss of generality, we can assume $v_P(\lambda) > 0$ (if $v_P(\lambda) < 0$, swap α with β, A with C, and B with D). Applying v_P to (32) gives

$$\begin{aligned} v_P(\lambda^n) &= n v_P(\lambda) \\ &= v_P\left(-\frac{C+Dn}{A+Bn}\right) \\ &\leq \log\left|\mathcal{N}_{\mathbb{K}/\mathbb{Q}}\left(-\frac{C+Dn}{A+Bn}\right)\right| \\ &\leq [\mathbb{K}:\mathbb{Q}]\log\left|\mathcal{N}_{abs}\left(-\frac{C+Dn}{A+Bn}\right)\right| \\ &= [\mathbb{K}:\mathbb{Q}]\log\prod_{i=1}^{[\mathbb{K}:\mathbb{Q}]}\left|\frac{\sigma_i(C)+\sigma_i(D)n}{\sigma_i(A)+\sigma_i(B)n}\right|, \end{aligned}$$

where $\sigma_1, \ldots, \sigma_{[\mathbb{K}:\mathbb{Q}]}$ are the monomorphisms from $\mathbb{K}$ into $\mathbb{C}$. As in the previous case, if

$$n > \frac{|\sigma_i(BC)| + |\sigma_i(AD)| + |\sigma_i(AB)|}{|\sigma_i(B)|^2} \stackrel{\text{def}}{=} N_i = 2^{\mathcal{O}(||I||)},$$

then we have

$$\left|\frac{\sigma_i(C)+\sigma_i(D)n}{\sigma_i(A)+\sigma_i(B)n}\right| \leq \left|\frac{\sigma_i(D)}{\sigma_i(B)}\right| + 1 \stackrel{\text{def}}{=} e_i = 2^{\mathcal{O}(||I||)}.$$

It follows therefore that if $n > \max_i\{N_i\}$, we have

$$v_P\left(-\frac{C+Dn}{A+Bn}\right) \leq [\mathbb{K}:\mathbb{Q}]\sum_{i=1}^{[\mathbb{K}:\mathbb{Q}]}\log e_i \stackrel{\text{def}}{=} M = ||I||^{\mathcal{O}(1)}.$$

Then for $n > \max_i\{N_i\}$ and $n > M$, we have

$$v_P(\lambda^n) = n v_P(\lambda) \geq n > M,$$

whereas

$$v_P\left(-\frac{C+Dn}{A+Bn}\right) \leq M,$$

so equation (32) cannot hold. □

LEMMA G.2. *Suppose $\langle u_n\rangle_{n=0}^{\infty}$ is non-degenerate and is given by (27). There exists a bound $N = 2^{\mathcal{O}(||I||)}$ such that if $u_n = 0$, then $n < N$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|D\| + \|\alpha\| + \|\beta\| + \|\gamma\|$.*

PROOF. We wish to solve for $n \in \mathbb{N}$ the equation:

$$(A+Bn)\alpha^n + C\beta^n + D\gamma^n = 0 \text{ (where } B, C, D \neq 0). \tag{33}$$

First suppose $|\alpha| \geq |\beta|, |\gamma|$. Then the term $(A+Bn)\alpha^n$ is dominant. More precisely, rewrite (33) as

$$A + Bn = -C\left(\frac{\beta}{\alpha}\right)^n - D\left(\frac{\gamma}{\alpha}\right)^n$$

and observe that if

$$n > \frac{|A| + |C| + |D|}{|B|},$$

then

$$|A + Bn| \geq |B|n - |A| > |C| + |D| \geq \left| -C \left(\frac{\beta}{\alpha}\right)^n - D \left(\frac{\gamma}{\alpha}\right)^n \right|,$$

so (33) cannot hold due to the strictness of the above inequality.

Second, suppose that $|\beta| > |\alpha|, |\gamma|$. Then the term $C\beta^n$ is dominant. More precisely, rewrite (33) as

$$(A + Bn) \left(\frac{\alpha}{\beta}\right)^n + D \left(\frac{\gamma}{\beta}\right)^n = -C. \tag{34}$$

We show that for n sufficiently large, the inequalities

$$\left| D \left(\frac{\gamma}{\beta}\right)^n \right| < \frac{|C|}{2}$$

and

$$\left| (A + Bn) \left(\frac{\alpha}{\beta}\right)^n \right| < \frac{|C|}{2}$$

both hold, rendering (34) impossible. The former inequality holds for $n > \log|C/2D| / \log|\gamma/\beta|$, which is at most exponentially large in the input. The latter inequality is implied by

$$\left| (n+1) \left(\frac{\alpha}{\beta}\right)^n \right| < \frac{|C|}{2M},$$

where $M = \max\{|A|, |B|\}$. Now let $r = \lceil -\log(2)/\log(\alpha/\beta) \rceil$, so that

$$\left(\frac{\alpha}{\beta}\right)^r \leq \frac{1}{2},$$

and consider only n of the form $n = kr$ for $k \in \mathbb{Z}^+$. If

$$k > \frac{\log|C/4Mr|}{\log(7/8)}$$

and $k \geq 5$, we have

$$\left(\frac{\alpha}{\beta}\right)^{kr} k < \left(\frac{1}{2}\right)^k (k+1) < \left(\frac{7}{8}\right)^k < \frac{|C|}{4Mr},$$

so

$$\left(\frac{\alpha}{\beta}\right)^n (n+1) \leq \left(\frac{\alpha}{\beta}\right)^n 2n < \frac{|C|}{2M}.$$

It is clear that r is at most exponentially large in the size of the input, whereas the bound on k is polynomial. Therefore, the bound on n is exponential.

Finally, suppose $|\beta| = |\gamma| > |\alpha|$. Rewrite (33) as

$$\left(\frac{\beta}{\gamma}\right)^n = -\frac{D}{C} - \frac{A + Bn}{C} \left(\frac{\alpha}{\gamma}\right)^n.$$

Then an exponential bound on n follows from Lemma E.1, because the right-hand side is a constant plus an exponentially decaying term, whereas the left-hand side is on unit circle. □

LEMMA G.3. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is non-degenerate and is given by (26). Suppose that $\alpha, \beta, \gamma, \delta$ do not all have the same magnitude. There exists a bound $N = 2^{\mathcal{O}(||I||)}$ such that if $u_n = 0$, then $n < N$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|D\| + \|\alpha\| + \|\beta\| + \|\gamma\| + \|\delta\|$.*

PROOF. We wish to solve for $n \in \mathbb{N}$ the equation:

$$A\alpha^n + B\beta^n + C\gamma^n + D\delta^n = 0 \text{ (where } A, B, C, D \neq 0). \tag{35}$$

Let $|\alpha| \geq |\beta| \geq |\gamma| \geq |\delta|$. First, if $|\alpha| > |\beta|$, then $A\alpha^n$ is the dominant term in (35). Rewrite the equation as

$$\frac{B}{A}\left(\frac{\beta}{\alpha}\right)^n + \frac{C}{A}\left(\frac{\gamma}{\alpha}\right)^n + \frac{D}{A}\left(\frac{\delta}{\alpha}\right)^n = -1$$

and observe that if

$$n > \max\left\{\frac{\log|3B/A|}{\log|\alpha/\beta|}, \frac{\log|3C/A|}{\log|\alpha/\gamma|}, \frac{\log|3D/A|}{\log|\alpha/\delta|}\right\},$$

then

$$\left|\frac{B}{A}\left(\frac{\beta}{\alpha}\right)^n + \frac{C}{A}\left(\frac{\gamma}{\alpha}\right)^n + \frac{D}{A}\left(\frac{\delta}{\alpha}\right)^n\right| < \frac{1}{3} + \frac{1}{3} + \frac{1}{3} = 1.$$

Second, if $|\alpha| = |\beta| > |\gamma|$, then rewrite (35) as

$$\left(\frac{\beta}{\alpha}\right)^n = -\frac{A}{B} - \frac{C}{B}\left(\frac{\gamma}{\alpha}\right)^n - \frac{D}{B}\left(\frac{\delta}{\alpha}\right)^n. \tag{36}$$

The left-hand side of (36) is on the unit circle, whereas the right is a constant plus exponentially decaying terms. An exponential bound on n follows from Lemma E.1.

Finally, if $|\alpha| = |\beta| = |\gamma| > |\delta|$, then an exponential bound on n follows from Lemma E.2 applied to equation (36). □

Thus, the only outstanding problem is to solve $u_n = 0$ in the case of $\langle u_n \rangle_{n=0}^{\infty}$ given by (26) when $|\alpha| = |\beta| = |\gamma| = |\delta|$. This case is difficult for general algebraic $\alpha, \beta, \gamma, \delta$: it is in fact the reason why the Skolem Problem is open for LRS of order 4 over $\mathbb{A}$. However, for real LRS, the set of characteristic roots is closed under complex conjugation, so complex roots come in conjugate pairs.

Another simplifying observation which is helpful for this last outstanding case is that for any LRS $\langle u_n \rangle_{n=0}^{\infty}$ over $\mathbb{A}$, one can find another LRS $\langle v_n \rangle_{n=0}^{\infty}$ over $\mathcal{O}_{\mathbb{A}}$ such that $u_n = 0$ if and only if $v_n = 0$. Indeed, recall that for any algebraic number α, it is possible to find an algebraic integer β and a rational integer M such that $\alpha = \beta/M$: it is sufficient to choose M to be the least common multiple of all denominators of the coefficients of the minimal polynomial of α. Then suppose the sequence $\langle u_n \rangle_{n=0}^{\infty}$ has initial terms $u_0, \ldots, u_{d-1} \in \mathbb{A}$ and satisfies a recurrence equation $u_n = \sum_{j=0}^{d-1} a_j u_{n-j-1}$ with $a_0, \ldots, a_{d-1} \in \mathbb{A}$. Let $M \in \mathbb{Z}$ be chosen so that $Ma_j \in \mathcal{O}_{\mathbb{A}}$ and $Mu_j \in \mathcal{O}_{\mathbb{A}}$ for $j = 0, \ldots, d-1$. Then it is easy to see that the sequence $\langle v_n \rangle_{n=0}^{\infty}$ defined by $v_n = M^{n+1}u_n$ has the same zero set as $\langle u_n \rangle_{n=0}^{\infty}$, has algebraic integer initial terms and satisfies a linear recurrence relation of order d with algebraic integer coefficients. Since M can be written down using only polynomial space, this reduction to the integer case can be

carried out in polynomial time. Therefore, by the integral closure of $\mathcal{O}_{\mathbb{A}}$, we can assume the characteristic roots $\alpha, \beta, \gamma, \delta$ are algebraic integers.

With these two observations in place, we proceed to the final technical result concerning the Skolem Problem for LRS of order 4 over $\mathbb{R} \cap \mathbb{A}$:

LEMMA G.4. *Suppose $\langle u_n \rangle_{n=0}^{\infty}$ is non-degenerate and is given by (26). Suppose that $\alpha, \beta, \gamma, \delta$ are algebraic integers with $|\alpha| = |\beta| = |\gamma| = |\delta|$. Suppose also $\{\alpha, \beta, \gamma, \delta\}$ is closed under complex conjugation. There exists a bound $N = 2^{\mathcal{O}(||I||)}$ such that if $u_n = 0$, then $n < N$, where $\|I\|$ is the length of the input $\|A\| + \|B\| + \|C\| + \|D\| + \|\alpha\| + \|\beta\| + \|\gamma\| + \|\delta\|$.*

PROOF. Let $\mathbb{K} = \mathbb{Q}(\alpha, \beta, \gamma, \delta, A, B, C, D)$. We have to solve for $n \in \mathbb{N}$ the equation:

$$A\alpha^n + B\beta^n + C\gamma^n + D\delta^n = 0 \text{ (where } A, B, C, D \neq 0). \tag{37}$$

The closure of $\{\alpha, \beta, \gamma, \delta\}$ under complex conjugation, the equality $|\alpha| = |\beta| = |\gamma| = |\delta|$ and the non-degeneracy of the LRS imply that the characteristic roots are two pairs of complex conjugates, so assume without loss of generality that $\beta = \overline{\alpha}$ and $\gamma = \overline{\delta}$. If α/β is an algebraic integer, then since it is not a root of unity, there exists a monomorphism σ from $\mathbb{K}$ to $\mathbb{C}$ such that $|\sigma(\alpha)| \neq |\sigma(\beta)|$. Applying σ to (37) leads to a Skolem instance of order 4 with roots whose magnitudes are not all the same. A bound on n follows from Lemma G.3.

Suppose then that α/β is not an algebraic integer. By the reasoning of Lemma A.2, there exists a prime ideal P in $\mathcal{O}_{\mathbb{K}}$ such that $v_P(\alpha) \neq v_P(\beta)$ and at least one of $v_P(\alpha)$ and $v_P(\beta)$ is strictly positive. Assume without loss of generality that

$$v_P(\alpha) > v_P(\beta) \geq 0.$$

Since $\alpha\beta = \gamma\delta = |\alpha|^2$, we have

$$v_P(\alpha) + v_P(\beta) = v_P(\gamma) + v_P(\delta).$$

Therefore, at most two of the roots are smallest under the valuation v_P.

If one root, say β, is strictly smaller under v_P than the rest, then rewrite (37) as

$$A\alpha^n + B\beta^n = -C\gamma^n - D\delta^n \tag{38}$$

Since $v_P(\beta) < v_P(\alpha)$, for $n > v_P(A/B)/v_P(\beta/\alpha)$ we have

$$v_P(A\alpha^n + B\beta^n) = v_P(B) + nv_P(\beta),$$

whereas

$$v_P(-C\gamma^n - D\delta^n) \geq v_P(C) + nv_P(\gamma).$$

Therefore, for $n > v_P(B/C)/v_P(\gamma/\beta)$, we have that the left-hand side of (38) is strictly smaller under v_P than the right-hand side, so (37) cannot hold. This bound on n is polynomial in the input size.

Now suppose that there are two roots with strictly smallest valuation with respect to v_P:

$$0 \leq v_P(\beta) = v_P(\gamma) < v_P(\alpha) = v_P(\delta).$$

Then rewrite (37) as

$$B\beta^n \left(\left(-\frac{C}{B} \right) \left(\frac{\gamma}{\beta} \right)^n - 1 \right) = A\alpha^n + D\delta^n. \tag{39}$$

Since γ/β is not a root of unity, the term $(-C/B)(\gamma/\beta)^n - 1$ can be zero for at most one value of n. This value is at most polynomially large in the input size (by Lemma D.1).

For all other n, we use Theorem A.5 to this term. Let p be the unique prime rational integer in the ideal P, and let $d = [\mathbb{K} : \mathbb{Q}]$. Let H be an upper bound for the heights of $-C/B$ and γ/β. Then by Theorem A.5 (van der Poorten), we have

$$v_P\left(\left(-\frac{C}{B}\right)\left(\frac{\gamma}{\beta}\right)^n - 1\right) \leq (48d)^{36}\frac{p^d}{\log p}(\log H)^2(\log n)^2. \tag{40}$$

It is classical that $\mathcal{N}(P) = p^f$ for some positive integer f, so $\mathcal{N}(P) \geq p$. Moreover, since α is an algebraic integer, all prime ideals $P_1, \ldots, P_s$ in the factorisation of $[\alpha]$ appear with positive exponents $k_1, \ldots, k_s$:

$$[\alpha] = P_1^{k_1} \ldots P_s^{k_s}.$$

Since $\mathcal{N}(P_i) \geq 2$ for all P_i, we have

$$|\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha)| = \mathcal{N}([\alpha]) \geq \mathcal{N}(P) \geq p.$$

Therefore, p is at most exponentially large in the input size. Then we can write (40) as

$$v_P\left(\left(-\frac{C}{B}\right)\left(\frac{\gamma}{\beta}\right)^n - 1\right) \leq E_1(\log n)^2,$$

where E_1 is exponentially large in the input size and independent of n. Now we apply v_P to both sides of equation (39):

$$v_P(LHS) \leq v_P(B) + nv_P(\beta) + E_1(\log n)^2$$

and

$$v_P(RHS) \geq v_P(A) + nv_P(\alpha).$$

Equation (37) cannot hold if

$$v_P(B) + nv_P(\beta) + E_1(\log n)^2 < v_P(A) + nv_P(\alpha),$$

which is implied by

$$v_P(B/A) + E_1(\log n)^2 < n,$$

since $v_P(\alpha) > v_P(\beta)$. Let $E_2 = \max\{v_P(B/A), E_1\}$, then this is implied by

$$E_2((\log n)^2 + 1) < n.$$

Since

$$(\log n)^2 + 1 < \frac{5\sqrt{n}}{2}$$

for all $n \geq 1$, it suffices to have

$$n > \left(\frac{5}{2}E_2\right)^2.$$

This bound on n is exponential in the size of the input. □

REFERENCES

V. Arvind and T. Vijayaraghavan. 2011. The orbit problem is in the GapL hierarchy. *J. Comb. Optim.* 21, 1 (2011), 124–137.

Alan Baker. 1975. *Transcendental number theory.* Cambridge University Press, Cambridge.

A. Baker and G. Wüstholz. 1993. Logarithmic Forms and Group Varieties. *Jour. Reine Angew. Math.* 442 (1993), 19–62.

Amir M Ben-Amram, Samir Genaim, and Abu Naser Masud. 2012. On the termination of integer loops. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 34, 4 (2012), 16.

Jean Berstel and Maurice Mignotte. 1976. Deux propriétés décidables des suites récurrentes linéaires. *Bulletin de la Société Mathématique de France* 104 (1976), 175–184.

P. Blanksby and H. Montgomery. 1971. Algebraic integers near the unit circle. *Acta Arith.* (1971), 355–369.

Vincent D. Blondel and Natacha Portier. 2002. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. *Linear Algebra Appl.* 351 (2002), 91 – 98.

M. Braverman. 2006. Termination of Integer Linear Programs. In *Proceedings of the 18th International Conference on Computer Aided Verification (CAV, LNCS 4144)*. Springer, 372–385.

Jin-yi Cai, Richard J Lipton, and Yechezkel Zalcstein. 2000. The complexity of the ABC problem. *SIAM J. Comput.* 29, 6 (2000), 1878–1888.

Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2013. The orbit problem in higher dimensions. In *STOC*. ACM, 941–950.

Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2015. The Polyhedron-hitting Problem. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '15)*. SIAM, 940–956. http://dl.acm.org/citation.cfm?id=2722129.2722193

H. Cohen. 1993. *A Course in Computational Algebraic Number Theory*. Springer.

Graham Everest, Alf van der Poorten, Thomas Ward, and Igor Shparlinski. 2003. *Recurrence Sequences*. American Mathematical Society.

Emmanuel Hainry. 2008. Reachability in linear dynamical systems. In *Logic and Theory of Algorithms*. Springer, 241–250.

V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. 2005. *Skolem's Problem – On the Border between Decidability and Undecidability*. Technical Report 683. Turku Centre for Computer Science.

G. Hansel. 1986. Une démonstration simple du théorème de Skolem-Mahler-Lech. *Theoretical Computer Science* 43 (1986), 91 – 98.

Michael A. Harrison. 1969. *Lectures on Linear Sequential Machines*. New York: Academic Press.

R. Kannan and R. Lipton. 1986. Polynomial-Time Algorithm for the Orbit Problem. *J. ACM* 33, 4 (1986), 808–821.

Ravindran Kannan and Richard J. Lipton. 1980. The orbit problem is decidable. In *Proceedings of the twelfth annual ACM symposium on Theory of computing (STOC)*. ACM, 252–261.

L. Kronecker. 1875. Zwei Sätze über Gleichungen mit ganzzahligen Koeffizienten. *J. Reine Angew. Math.* 53 (1875), 173–175.

Christer Lech. 1953. A note on recurring series. *Arkiv för Matematik* 2 (1953), 417–421.

A. K. Lenstra, H. W. Lenstra, and L. Lovász. 1982. Factoring polynomials with rational coefficients. *Math. Ann.* 261 (1982), 515–534.

B Litow. 1997. A decision method for the rational sequence problem. In *Electronic Colloquium on Computational Complexity (ECCC)*, Vol. 4.

K. Mahler. 1935. Eine arithmetische Eigenschaft der Taylor-Koeffizienten rationaler Funktionen. *Proc. Akad. Wet. Amsterdam* 38 (1935), 51–60.

K. Mahler and J. W. S. Cassels. 1956. On the Taylor coefficients of rational functions. *Mathematical Proceedings of the Cambridge Philosophical Society* 52 (1 1956), 39–48. Issue 01. DOI: http://dx.doi.org/10.1017/S0305004100030966

M. Mignotte. 1982. Some Useful Bounds. *Computer Algebra* (1982), 259–263.

M. Mignotte, T. Shorey, and R. Tijdeman. 1984. The distance between terms of an algebraic recurrence sequence. *Jour. Reine Angew. Math.* 349 (1984), 63 – 76.

Joël Ouaknine and James Worrell. 2012. Decision Problems for Linear Recurrence Sequences. In *Reachability Problems*, Alain Finkel, Jérôme Leroux, and Igor Potapov (Eds.). Lecture Notes in Computer Science, Vol. 7550. Springer Berlin Heidelberg, 21–28. DOI: http://dx.doi.org/10.1007/978-3-642-33512-9_3

V. Pan. 1996. Optimal and nearly optimal algorithms for approximating polynomial zeros. *Computers & Mathematics with Applications* 31, 12 (1996), 97 – 138.

Arto Salomaa and Matti Soittola. 1978. *Automata-theoretic aspects of formal power series*. Springer-Verlag Berlin.

Arnold Schönhage. 1979. On the power of random access machines. In *Automata, Languages and Programming*, Hermann Maurer (Ed.). Lecture Notes in Computer Science, Vol. 71. Springer Berlin / Heidelberg, 520–529.

Th. Skolem. 1934. Ein Verfahren zur Behandlung gewisser exponentialer Gleichungen und diophantischer Gleichungen. *Skand. Mat. Kongr.* 8 (1934), 163–188.

I. Stewart and D. Tall. 2002. *Algebraic Number Theory and Fermat's Last Theorem* (3rd ed.). A. K. Peters.

T. Tao. 2008. *Structure and randomness: pages from year one of a mathematical blog*. American Mathematical Society.

Sergey Tarasov and Mikhail Vyalyi. 2011. Orbits of Linear Maps and Regular Languages. In *Computer Science – Theory and Applications*, Alexander Kulikov and Nikolay Vereshchagin (Eds.). Lecture Notes in Computer Science, Vol. 6651. Springer Berlin Heidelberg, 305–316. DOI: http://dx.doi.org/10.1007/978-3-642-20712-9_24

Ashish Tiwari. 2004. Termination of linear programs. In *Computer Aided Verification*. Springer, 70–82.

Alfred Jacobus van der Poorten. 1977. Linear forms in logarithms in the p-adic case. *Transcendence Theory: Advances and Applications* (1977), 29–57.

N. Vereshchagin. 1985. Occurrence of zero in a linear recursive sequence. *Mathematical Notes* 38 (1985), 609–615.

Richard Zippel. 1997. Zero testing of algebraic functions. *Inform. Process. Lett.* 61, 2 (1997), 63 – 67.