On the Complexity of the Orbit Problem

VENTSISLAV CHONEV, University of Oxford
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 V may be reached from a starting point x under repeated applications of a linear transformation A. Answering two questions posed by Kannan and Lipton in the 1980s, we show that when V has dimension one, this problem is solvable in polynomial time, and when 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. *ACM* 0, 0, Article 0 (0), 33 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

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

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

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

Author's address: V. Chonev, J. Ouaknine and J. Worrell, Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, UK

DOI: http://dx.doi.org/10.1145/0000000.0000000

the target space V has dimension $m-1$. The sequence of numbers $u_n := y^T A^n x$ 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 1974]. The main difficulty in deciding Skolem's 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], 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 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 results are the following. We show that the higher-dimensional Orbit Problem can be solved in polynomial time 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 Skolem's Problem, our results, in contrast, are independent of the dimension of the matrix A.

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:

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

Then with target space

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

it can be shown that $A^n x \in V$ if and only if n has residue 2 modulo 6. Such periodic behaviour can be analysed in terms of the eigenvalues of the matrix 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 $\lambda = (-1 + i\sqrt{39})/2$. The key observation is that the eigenvalues of 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 $A^n x \in V$. Computable bounds on such an n are obtained utilising the work of [Mignotte et al. 1984; Vereshchagin 1985], quantifying and strengthening some of the

[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.

bounds given for Skolem's Problem. In the case of a one-dimensional target subspace 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 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 $A^n x \in V$ for a guessed value of n. Finally, the case in which the eigenvalues of 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 $A^n x \in V$. 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 using the guessing power of an $\mathbf{NP}$ machine.

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 *chamber hitting problem*, which replaces the target space with an intersection of half-spaces. In [Tarasov and Vyalyi 2010], the chamber hitting problem is related to decision problems in formal language theory. 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 A, B and C, whether there exist integers i and j such that $A^i B^j = C$.

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

1.2. Paper Outline

The structure of the paper is as follows. In Section 2, we reduce the higher-dimensional Orbit Problem to the *matrix power problem* in polynomial time, using standard techniques from linear algebra. We also describe a crucial element of our approach: a *Master System* of equations based on the eigenvalues of the input matrix. Then we proceed to solve the Orbit Problem with a target subspace of dimension one, two and three in Sections 3, 4 and 5, respectively. Whilst these fixed-dimensional versions of the Orbit Problem are closely related to Skolem's Problem for linear recurrence sequences of order two, three and four, in the interest of clarity, we have avoided referring to Skolem's Problem in the main text, instead consigning all technical lemmas concerning it to Appendices D, E and F. These lemmas rely crucially on two theorems from transcendence theory, due to Baker-Wüstholz and van der Poorten, presented in Appendix C, as well as some standard results from algebraic number theory concerning the efficient manipulation of algebraic numbers and p-adic valuations in a fixed number field, recounted in Appendices A and B. Finally, we give concluding remarks in Section 6.

2. REDUCTION

2.1. Matrix power problem

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

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

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

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

It is easy to check that $DM = AD$. Thus, if some vector z has coordinates z' with respect to $\mathcal{B}$, so that $z = Dz'$, then Az has coordinates Mz' with respect to $\mathcal{B}$, so that $Az = DMz'$. By induction, for all $n \in \mathbb{N}$, $A^n x = DM^n x'$, where $x' = [1, 0, \ldots, 0]^T$. Next we calculate a basis for $W = U \cap V$, let this basis be $\{w_1, \ldots, w_t\}$ and let $w_i = Dw_i'$ for all i. Now,

$$A^n x \in V \iff A^n x \in W \iff M^n x' \in span\{w_1', \ldots, w_t'\}$$

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

Define the matrices $T_1, \ldots, T_t$ by

$$T_i = [\, w_i' \;\; Mw_i' \;\; \ldots \;\; M^\nu w_i' \,]$$

We will show that $M^n x' \in span\{w_1', \ldots, w_t'\}$ if and only if $M^n \in span\{T_1, \ldots, T_t\}$. If for some coefficients a_i we have

$$M^n = \sum_{i=0}^{t} a_i T_i$$

then considering the first column of both sides, we have

$$M^n x' = \sum_{i=0}^{t} a_i w_i'$$

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

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 out of the space

$T = span\{T_1, \ldots, T_t\}$ it suffices to consider only matrices of the shape $p(M)$ where p is a polynomial. We find a basis for the space $P = \{p(M) \mid p \in \mathbb{Q}[x]\}$ and then a basis $\{p_1(M), \ldots, p_s(M)\}$ for $P \cap T$. Then $M^n \in T \iff M^n \in P \cap T$. We call this the *polynomial version* of the matrix power problem. Observe that $\dim(V) \geq \dim(T) \geq \dim(T \cap P)$, 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.

2.2. Degenerate and non-degenerate problem instances

An instance (A, x, V) of the Orbit Problem is defined as *non-degenerate* if no quotient of two distinct eigenvalues of A is a root of unity. In general, it is possible to reduce an arbitrary Orbit Problem instance to a set of non-degenerate instances as follows.

Let L be the least common multiple of the orders of all eigenvalue quotients which are roots of unity. For each $i \in \{0, \ldots, L-1\}$, we consider separately the problem of deciding whether there exists $n \in \mathbb{N}$ such that $\left(A^L\right)^n \left(A^i x\right) \in V$. The original problem instance is positive if and only if at least one of these L instances is positive.

It is easy to see that these instances are all non-degenerate. The eigenvalues of A^L are exactly λ_i^L where λ_i are the eigenvalues of 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 $\lambda_i^L = \lambda_j^L$.

Observe that the reduction to the matrix power problem in the previous section preserves non-degeneracy, since the eigenvalues of the invertible matrix M are a subset of the eigenvalues of A. However, the non-degeneracy comes at the cost of higher complexity, as L may be exponentially large in the size of A. A deterministic machine would have to examine L problem instances to decide the original instance. Even if we allow non-determinism, we must still calculate A^L, whose entries require space exponential in the size of the original problem instance.

We show this reduction to make the point that if one only aims to prove decidability for the Orbit Problem, without regard for complexity, non-degeneracy of the problem instance may be assumed without loss of generality. Below we give tighter complexity upper bounds for the Orbit Problem with a target space of dimension at most 3 and explicitly handle degenerate problem instances, but if we only aimed at proving decidability, the argument could be shortened significantly.

Finally, we point out that a similar device was used by Vereshchagin [Vereshchagin 1985] to prove decidability of Skolem's Problem for recurrences of order 3 and 4. Given a linear recurrence sequence $u(n)$, one may take

$$L = \text{lcm}\{m \mid \lambda_i / \lambda_j \text{ is an } m\text{-th root of unity}\}$$

where λ_i are the roots of the characteristic polynomial of the sequence. Then each sequence $v_i(n) = u(Ln + i)$ for $i \in \{0, \ldots, L-1\}$ is non-degenerate in the sense that dividing two of its characteristic roots never yields a root of unity.

2.3. Towards a system of equations

Suppose now we have an instance $(A, p_1, \ldots, p_s)$ of the polynomial version of the matrix power problem. Calculate the minimal polynomial $f_A(x)$ of A and obtain canonical representations of its roots $\alpha_1, \ldots, \alpha_k$, that is, the eigenvalues of A. This may be done in polynomial time, see Appendix A. Throughout this paper, 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 $a_1, \ldots, a_s \in \mathbb{C}$ and define the polynomials $P(x) = \sum_{i=1}^{s} a_i p_i(x)$ and $Q(x) = x^n$. It is easy to see that

$$Q(A) = P(A)$$

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 A if and only if $f_A(x)$ divides $P - Q$, that is, each α_i is a root of $P - Q$ with multiplicity at least $mul(\alpha_i)$. This is equivalent to saying that 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 a_i such that $A^n = \sum_{i=1}^{s} a_i p_i(A)$, it is sufficient to solve a system of equations (1) where the unknowns are $n \in \mathbb{N}$ and $a_1, \ldots, a_s \in \mathbb{C}$. 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 $f_A(x)$ has roots α_1, α_2, α_3 with multiplicities $mul(\alpha_i) = i$ and the target space is $span\{p_1(A), p_2(A)\}$ then the system contains six equations:

$$\begin{aligned}
\alpha_1^n &= a_1 p_1(\alpha_1) + a_2 p_2(\alpha_1) \\
\alpha_2^n &= a_1 p_1(\alpha_2) + a_2 p_2(\alpha_2) \\
n\alpha_2^{n-1} &= a_1 p_1'(\alpha_2) + a_2 p_2'(\alpha_2) \\
\alpha_3^n &= a_1 p_1(\alpha_3) + a_2 p_2(\alpha_3) \\
n\alpha_3^{n-1} &= a_1 p_1'(\alpha_3) + a_2 p_2'(\alpha_3) \\
n(n-1)\alpha_3^{n-2} &= a_1 p_1''(\alpha_3) + a_2 p_2''(\alpha_3)
\end{aligned}$$

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

$$\alpha_3^n = a_1 p_1(\alpha_3) + a_2 p_2(\alpha_3)$$

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

$$\alpha_2^n = a_1 p_1(\alpha_2) + a_2 p_2(\alpha_2)$$

$$n\alpha_2^{n-1} = a_1 p_1'(\alpha_2) + a_2 p_2'(\alpha_2)$$

3. ONE-DIMENSIONAL VERSION

Suppose we are given a one-dimensional matrix power problem instance (A, p) and wish to decide whether $A^n \in span\{p(A)\}$ for some n. We have constructed a system of equations in the exponent n and the coefficient a as in (1). For example, if the roots of $f_A(x)$ are $\alpha_1, \alpha_2, \alpha_3$ with multiplicities $mul(\alpha_i) = i$, the system is:

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

In this section we will describe how such systems may be solved in polynomial time. We allow the problem instance to be degenerate, that is, the ratios of eigenvalues of A may be roots of unity.

The strategy is to consider quotients of equations. This eliminates the unknown coefficient a and leaves only the exponent n. First, we perform some preliminary calculations.

(1) We check whether $a = 0$ has a corresponding n which solves the matrix equation $A^n = ap(A)$. This may be done using Kannan and Lipton's algorithm for the original Orbit Problem. If this is the case, we are done. Otherwise, assume $a \neq 0$.
(2) Let $c = \max_i\{mul(\alpha_i)\}$. We check for all $n < c$ whether A^n is a multiple of $p(A)$. If so, we are done. Otherwise, assume $n \geq c$.
(3) We check whether $\alpha_i = 0$ for some i. If so, then all of the equations $Eq(\alpha_i)$ are of the form $0 = ap^{(t)}(0)$, 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 $\alpha_i \neq 0$ for all i.
(4) Finally, we check whether the right-hand side $ap^{(t)}(\alpha_i)$ of some equation is equal to 0, using a polynomial division of $p^{(t)}(x)$ by the minimal polynomial of α_i. 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 by dividing $eq(\alpha_i, k)$ by $eq(\alpha_j, t)$, 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 1. 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 Appendix A, 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 Appendix D, n is bounded by a polynomial in the input. We check $A^n \in span\{p(A)\}$ for all n up to the bound and we are done.

Case 2. All quotients α_i/α_j are roots of unity, and all roots of $f_A(x)$ are simple. Then the system is equivalent to

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

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 3. 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 $a = \alpha_1/p(\alpha_1)$.

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

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 Appendix A, 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

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

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

4. TWO-DIMENSIONAL VERSION

Suppose we are given (A, p_1, p_2) where A is a square matrix and p_1, p_2 are polynomials. We wish to decide whether there exists $n \in \mathbb{N}$ such that $A^n \in span\{p_1(A), p_2(A)\}$. In this section we will show that this problem is in the complexity class $\mathbf{NP^{EqSLP}}$, and hence in $\mathbf{NP^{RP}}$, since $\mathbf{EqSLP} \subseteq \mathbf{coRP}$ by reference [Schönhage 1979]. The instance is allowed to be degenerate, that is, there may exist distinct eigenvalues of A whose ratio is a root of unity. We have derived a Master System of equations (1).

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

$$0 = a_1 p_1(0) + a_2 p_2(0)$$

If at least one of $p_1(0)$, $p_2(0)$ is non-zero, then we have a linear dependence between a_1, a_2, so we express one in terms of the other and proceed to solve a one-dimensional

problem instance. Otherwise, $eq(0,0)$ is trivially satisfied for all n, a_1, a_2, so we dismiss it from the Master System. We examine in this way all equations contributed by 0, either dismissing them or obtaining a lower-dimensional system.

The broad strategy for the two-dimensional Orbit Problem will be to choose a tuple of equations from the system (1), obtain a Skolem instance of order 3 from this tuple and hence obtain a bound m on the exponent n such that if

$$A^n \in span\{p_1(A), p_2(A)\}$$

then $n < m$. This bound will be at most exponential in the size of the input, so an **NP** machine will be able to guess an exponent n up to the bound. Then the machine calculates A^n, representing numbers as arithmetic circuits, and expresses membership in the target vector space as an instance of the **EqSLP** problem: determining whether a given arithmetic circuit evaluates to 0.

For example, suppose A has three distinct eigenvalues α, β, γ, and consider the tuple of equations $eq(\alpha, 0)$, $eq(\beta, 0)$, $eq(\gamma, 0)$:

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\beta) \\ p_2(\gamma) \end{pmatrix}$$

If the vectors on the right-hand side are linearly independent over $\mathbb{C}$, then this triple states that the point in $\mathbb{C}^3$ described by the left-hand side lies on the plane in $\mathbb{A}^3$ described by the right-hand side. We can calculate the normal $(A_1, A_2, A_3)^T$ of the plane to obtain

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

It is a classical result [Everest et al. 2003] that the left-hand side as a function of n satisfies a linear recurrence of order 3 over $\mathbb{A}$. Provided that the ratios of α, β, γ are not roots of unity, Lemmas E.1, E.2, E.3 give a bound on n which is at most exponential in the input size, as desired. Problems arise, however, when this linear recurrence is allowed to be degenerate. For example, suppose that $A_3 = 0$, and let α/β and $-A_2/A_1$ be roots of unity of the same order. Then the zeros of the recurrence are a full arithmetic progression, and this linear recurrence sequence fails to give a bound on n.

Similarly, if the matrix A has only two eigenvalues α, β, but one of them, say α, is repeated, then we may consider the triple $eq(\alpha, 0)$, $eq(\alpha, 1)$, $eq(\beta, 0)$. Using similar reasoning, we see that

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

must hold for some effective algebraic constants A_1, A_2, A_3. This corresponds to a linear recurrence sequence of order 3 where one of the characteristic roots is repeated. Now provided that α/β is not a root of unity, we have an exponential bound on n from Lemma E.4. However, if α/β and $-A_3/A_1$ are roots of unity, and $A_2 = 0$, then the linear recurrence sequence could have infinitely many solutions, failing to give a bound on n.

In order to explicitly handle degenerate Orbit instances, we will consider the relation $\sim$ on the eigenvalues of 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 $\mathcal{C}_1, \ldots, \mathcal{C}_k$ of $\sim$ are of two kinds. First, a class can be its own image under complex conjugation:

$$\mathcal{C}_i = \{\overline{\alpha} \mid \alpha \in \mathcal{C}_i\}$$

Each such self-conjugate class $\{\alpha_1, \ldots, \alpha_s\}$ has the form $\{\alpha\omega_1, \ldots, \alpha\omega_s\}$ where ω_i are roots of unity, and $|\alpha_j| = \alpha \in \mathbb{R} \cap \mathbb{A}$. Call this α the *stem* of the equivalence class

$\mathcal{C}_i$. 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 $(\mathcal{C}_i, \mathcal{C}_j)$ such that $\mathcal{C}_i = \{\overline{x} \mid x \in \mathcal{C}_j\} = \overline{\mathcal{C}_j}$. In this case, we can write $\mathcal{C}_i$ and $\mathcal{C}_j$ as

$$\mathcal{C}_i = \{\lambda\omega_1, \ldots, \lambda\omega_s\}$$
$$\mathcal{C}_j = \{\overline{\lambda\omega_1}, \ldots, \overline{\lambda\omega_s}\}$$

where ω_i are roots of unity, $\lambda \in \mathbb{A}$ and $\arg(\lambda)$ is an irrational multiple of 2π. Call λ the *stem* of $\mathcal{C}_i$ and $\overline{\lambda}$ the stem of $\mathcal{C}_j$. Observe that the stems of self-conjugate classes are distinct positive real numbers, and that the ratio $\lambda/\overline{\lambda}$ of the stems of paired classes cannot be a root of unity. Recall also that we can assume the eigenvalues of 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 stems of equivalence classes must also be algebraic integers.

Let

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

denote the set of equations contributed to the system by the eigenvalues in $\mathcal{C}$, and let

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

denote the set of i-th derivative equations contributed by the roots in $\mathcal{C}$. We will case split on the equivalence classes of $\sim$.

Case I. Suppose $\sim$ has exactly one equivalence class $\mathcal{C} = \{\alpha\omega_1, \ldots, \alpha\omega_s\}$, necessarily self-conjugate, with stem α. Let L be the least common multiple of the orders of $\omega_1, \ldots, \omega_s$ and notice that L is at most exponentially large in the size of the input, so it can be expressed using polynomially many bits. We proceed by case analysis on the residue of n modulo L. Suppose $n \bmod L = r$ and consider the set of equations $Eq\{\mathcal{C}, 0\}$:

$$(\alpha\omega_1)^n = a_1 p(\alpha\omega_1) + a_2 p(\alpha\omega_1)$$
$$\vdots$$
$$(\alpha\omega_s)^n = a_1 p(\alpha\omega_s) + a_2 p(\alpha\omega_s)$$

For a fixed r, we can easily calculate $\omega_1^n, \ldots, \omega_s^n$ in polynomial time, since ω_i are roots of unity whose order divides L. Then the equations $Eq(\mathcal{C}, 0)$ are equivalent to

$$\begin{pmatrix} \alpha^n \\ \vdots \\ \alpha^n \end{pmatrix} = B \begin{pmatrix} a_1 \\ \\ a_2 \end{pmatrix} \tag{5}$$

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

$$\alpha^n = \varphi_1 a_1 + \varphi_2 a_2 \wedge \begin{pmatrix} 0 \\ \vdots \\ 0 \end{pmatrix} = B' \begin{pmatrix} a_1 \\ \\ a_2 \end{pmatrix}$$

Here $(\ \varphi_1 \ \ \varphi_2\ )$ is the first row of the matrix B, and B' is the result of subtracting $(\ \varphi_1 \ \ \varphi_2\ )$ from each of the bottom $s-1$ rows of B. Thus, $Eq(\mathcal{C}, 0)$ is equivalent to

$$\alpha^n = \varphi_1 a_1 + \varphi_2 a_2$$

together with the constraint that $(a_1 \;\; a_2)^T$ must lie in the nullspace of B'. We calculate the nullspace of B' directly. If its dimension is less than 2, then we have a linear constraint on a_1, a_2. This constraint is of the form $a_1 = ka_2$ when the nullspace of B' has dimension 1, and is $a_1 = a_2 = 0$ when the nullspace is of dimension 0. In both cases, we substitute into the system (1), and solve the resulting lower-dimensional system using the algorithms for the one-dimensional Orbit Problem and Kannan and Lipton's original Orbit Problem. In the case when the nullspace of B' has dimension 2, then the linear constraint is vacuous, and $Eq(\mathcal{C}, 0)$ is equivalent to $\alpha^n = \varphi_1 a_1 + \varphi_2 a_2$.

In the same way, a case analysis on n mod L reduces $Eq(\mathcal{C}, 1)$ into a single first-derivative equation:

$$n\alpha^{n-1} = \varphi_3 a_1 + \varphi_4 a_2$$

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

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

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \\ n(n-1)\alpha^{n-2} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \\ \varphi_5 \end{pmatrix} + a_2 \begin{pmatrix} \varphi_2 \\ \varphi_4 \\ \varphi_6 \end{pmatrix} \tag{6}$$

If the vectors on the right-hand side of (6) are linearly independent, 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. We calculate the normal $(A_1 \;\; A_2 \;\; A_3)^T$ of the plane and obtain

$$A_1\alpha^n + A_2 n\alpha^{n-1} + A_3 n(n-1)\alpha^{n-2} = 0$$

This is a quadratic equation in n. It has at most two solutions, both at most exponentially large in the size of the input, so we are done. If the vectors on the right-hand side of (6) are linearly dependent, then we may equivalently consider

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \\ n(n-1)\alpha^{n-2} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \\ \varphi_5 \end{pmatrix}$$

We divide the first equation by the second to obtain

$$\frac{\alpha}{n} = \frac{\varphi_1}{\varphi_3}$$

which limits n at most one, exponentially large, candidate value.

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

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \end{pmatrix} + a_2 \begin{pmatrix} \varphi_2 \\ \varphi_4 \end{pmatrix} \tag{7}$$

If $(\varphi_1 \;\; \varphi_3)^T$ and $(\varphi_2 \;\; \varphi_4)^T$ are linearly independent, then the right-hand side of (7) spans all of $\mathbb{A}^2$ as a_1, a_2 range over $\mathbb{A}$, so the problem instance is trivially positive. Otherwise, we may equivalently consider

$$\begin{pmatrix} \alpha^n \\ n\alpha^{n-1} \end{pmatrix} = a_1 \begin{pmatrix} \varphi_1 \\ \varphi_3 \end{pmatrix}$$

and limit n to at most one candidate value which is exponentially large in the input size.

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

$$\alpha^n = a_1\varphi_1 + a_2\varphi_2$$

which has a solution if and only if at least one of φ_1, φ_2 is non-zero.

Case II. Suppose $\sim$ has exactly two equivalence classes, $\mathcal{C}_1$ and $\mathcal{C}_2$, with respective stems α and β, so that

$$\mathcal{C}_1 = \{\alpha\omega_1, \ldots, \alpha\omega_s\}$$

$$\mathcal{C}_2 = \{\beta\omega'_1, \ldots, \beta\omega'_l\}$$

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 the previous case, we define L to be the least common multiple of the orders of $\omega_1, \ldots, \omega_s$, $\omega'_1, \ldots, \omega'_l$, and proceed by case analysis on $r = n \bmod L$. We transform the system $Eq(\mathcal{C}_1) \wedge Eq(\mathcal{C}_2)$ into the equivalent system $\mathcal{F}(Eq(\mathcal{C}_1)) \wedge \mathcal{F}(Eq(\mathcal{C}_2))$. If all eigenvalues x of A have $mul(x) = 1$, then the system consists of two equations, one for each equivalence class of $\sim$:

$$\begin{aligned}\alpha^n &= a_1\varphi_1 + a_2\varphi_2\\ \beta^n &= a_1\varphi_3 + a_2\varphi_4\end{aligned}$$

If $(\ \varphi_1\ \ \varphi_3\ )^T$ and $(\ \varphi_2\ \ \varphi_4\ )^T$ are linearly independent, then there is a solution for each n, so the problem instance is positive. Otherwise, it suffices to look for n which satisfies

$$\begin{aligned}\alpha^n &= a_1\varphi_1\\ \beta^n &= a_1\varphi_3\end{aligned}$$

and hence

$$\left(\frac{\alpha}{\beta}\right)^n = \frac{\varphi_1}{\varphi_3}$$

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 A has $mul(x) \geq 2$, say $x \in \mathcal{C}_1$, then the system contains the following triple of equations:

$$\begin{pmatrix}\alpha^n\\ n\alpha^{n-1}\\ \beta^n\end{pmatrix} = a_1\begin{pmatrix}\varphi_1\\ \varphi_3\\ \varphi_5\end{pmatrix} + a_2\begin{pmatrix}\varphi_2\\ \varphi_4\\ \varphi_6\end{pmatrix} \tag{8}$$

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

$$\begin{pmatrix}\alpha^n\\ n\alpha^{n-1}\\ \beta^n\end{pmatrix} = a_1\begin{pmatrix}\varphi_1\\ \varphi_3\\ \varphi_5\end{pmatrix}$$

Then dividing we obtain

$$\frac{\alpha}{n} = \frac{\varphi_1}{\varphi_3}$$

which limits n to at most one, exponentially large candidate value. Otherwise, if the vectors on the right-hand side of (8) are linearly independent, 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 E.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 $eq(\gamma, 0)$:

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\beta) \\ p_2(\gamma) \end{pmatrix}$$

If the vectors on the right-hand side are linearly independent, we calculate the normal $(A_1, A_2, A_3)^T$ of the plane on the right-hand side 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 E.1, E.2, E.3. If the vectors on the right-hand side are not linearly independent, then we may equivalently consider

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \end{pmatrix}$$

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 in all cases, an **NP** machine may compute a bound m such that if

$$A^n x \in span\{y, z\}$$

then $n < m$. This bound is at most exponential in the size of the input. From here it is easy to argue membership in $\mathbf{NP}^{\mathbf{EqSLP}}$. The machine guesses some n up to the bound, ensuring that this n is consistent with the guess for n mod L. Now we need to compute $A^n x$ and check whether it is in the target vector space. Since n has magnitude at most exponential in the size of the input, the entries of $A^n x$ are, in general, doubly-exponential in magnitude. That is, they require an exponential number of bits to write down. However, the entries of $A^n x$ may easily be represented as polynomial-sized arithmetic circuits. We consider all projections of $A^n x$, y and z to three coordinates, and for each projection we express the question of linear independence as the zeroness of a 3×3 determinant, also expressed as an arithmetic circuit. It is clear that n is a witness to the problem instance if and only if for any projection to three coordinates, $A^n x$, y and z are linearly dependent. 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 also have membership in $\mathbf{NP}^{\mathbf{RP}}$.

5. THREE-DIMENSIONAL VERSION

Suppose we have a problem instance (A, p_1, p_2, p_3) and wish to decide whether $A^n \in span\{p_1(A), p_2(A), p_3(A)\}$ for some n. As before, we have constructed a system (1) in n and the coefficients a_1, a_2, a_3. The eigenvalues of A are non-zero algebraic integers (if 0 is an eigenvalue, then $eq(0, 0)$ gives a linear dependence between the coefficients a_1, a_2, a_3, so we proceed to solve a lower-dimensional problem instance). In this section we will show that there exists an effective bound m which is at most exponentially

large in the size of the input, such that

$$A^n \in span\{p_1(A), p_2(A), p_3(A)\} \Rightarrow n < m$$

Then by the same reasoning as in the two-dimensional case, we will have membership in $\mathbf{NP}^{\mathbf{EqSLP}}$ and $\mathbf{NP}^{\mathbf{RP}}$ for the three-dimensional Orbit Problem.

Following the strategy of the two-dimensional case, we will select tuples of equations and obtain a bound on n using the lemmas for Skolem's Problem for recurrences of order 4 in Appendix F. 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 $(\mathcal{C}_i, \overline{\mathcal{C}_i})$, $(\mathcal{C}_j, \overline{\mathcal{C}_j})$ which are not self-conjugate. Then let $\alpha \in \mathcal{C}_i$, $\beta = \overline{\alpha} \in \overline{\mathcal{C}_i}$, $\gamma \in \mathcal{C}_j$, $\delta = \overline{\gamma} \in \overline{\mathcal{C}_j}$. Then we consider the tuple of equations

$$\begin{pmatrix} \alpha^n \\ \beta^n \\ \gamma^n \\ \delta^n \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\beta) \\ p_1(\gamma) \\ p_1(\delta) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\beta) \\ p_2(\gamma) \\ p_2(\delta) \end{pmatrix} + a_3 \begin{pmatrix} p_3(\alpha) \\ p_3(\beta) \\ p_3(\gamma) \\ p_3(\delta) \end{pmatrix} \tag{9}$$

If the vectors on the right-hand side are linearly dependent, then we rewrite the right-hand side as a linear combination of at most 2 vectors and obtain a bound on n as we did for tuples of equations in the one- and two-dimensional Orbit Problem. If the vectors on the right-hand side of (9) are linearly independent, 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 F.3 and F.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 F.4.

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

$$\begin{pmatrix} \alpha^n \\ \overline{\alpha}^n \\ n\alpha^{n-1} \\ n\overline{\alpha}^{n-1} \end{pmatrix} = a_1 \begin{pmatrix} p_1(\alpha) \\ p_1(\overline{\alpha}) \\ p_1'(\alpha) \\ p_1'(\overline{\alpha}) \end{pmatrix} + a_2 \begin{pmatrix} p_2(\alpha) \\ p_2(\overline{\alpha}) \\ p_2'(\alpha) \\ p_2'(\overline{\alpha}) \end{pmatrix} + a_3 \begin{pmatrix} p_3(\alpha) \\ p_3(\overline{\alpha}) \\ p_3'(\alpha) \\ p_3'(\overline{\alpha}) \end{pmatrix}$$

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 on n follows from Lemma F.1, since $\alpha/\overline{\alpha}$ is not a root of unity.

We can now assume that eigenvalues in $\mathcal{C}_i$ and $\overline{\mathcal{C}_i}$ contribute exactly one equation to the system. Now we proceed again by case analysis on $r = n \bmod L$, transforming $Eq(\mathcal{C}_i) \wedge Eq(\overline{\mathcal{C}_i})$ into the equivalent $\mathcal{F}(Eq(\mathcal{C}_i)) \wedge \mathcal{F}(Eq(\overline{\mathcal{C}_i}))$. Since all eigenvalues in $\mathcal{C}_i$ and $\overline{\mathcal{C}_i}$ contribute one equation each, $\mathcal{F}(Eq(\mathcal{C}_i)) \wedge \mathcal{F}(Eq(\overline{\mathcal{C}_i}))$ is just

$$\begin{aligned} \lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3 \\ \overline{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6 \end{aligned}$$

where $\lambda, \overline{\lambda}$ are the stems of $\mathcal{C}_i$ and $\overline{\mathcal{C}_i}$. We do the same to all self-conjugate classes as well, reducing the system of equations to an equivalent system based on the stems of the equivalence classes, not the actual eigenvalues of A. This is beneficial, because

the stems 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 stems α, β, we take the tuple

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3\\ \overline{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6\\ \alpha^n &= a_1\varphi_7 + a_2\varphi_8 + a_3\varphi_9\\ \beta^n &= a_1\varphi_{10} + a_2\varphi_{11} + a_3\varphi_{12}\end{aligned}$$

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 a bound on n from Lemmas F.3 and F.4. Similarly, if there is only one self-conjugate equivalence class, with stem α, but some of its eigenvalues are repeated, we use the tuple

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3\\ \overline{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6\\ \alpha^n &= a_1\varphi_7 + a_2\varphi_8 + a_3\varphi_9\\ n\alpha^{n-1} &= a_1\varphi_{10} + a_2\varphi_{11} + a_3\varphi_{12}\end{aligned}$$

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 a bound on n according to Lemma F.2. If there is exactly one self-conjugate class, with stem α, containing no repeated roots, then the system consists of three equations:

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3\\ \overline{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6\\ \alpha^n &= a_1\varphi_7 + a_2\varphi_8 + a_3\varphi_9\end{aligned}$$

Depending on whether the vectors $(\ \varphi_1\ \ \varphi_4\ \ \varphi_7\ )^T$, $(\ \varphi_2\ \ \varphi_5\ \ \varphi_8\ )^T$, $(\ \varphi_3\ \ \varphi_6\ \ \varphi_9\ )^T$ are linearly independent, this is either a trivially positive instance, or a lower-dimensional non-degenerate instance. Finally, if there are no self-conjugate classes, the system consists of only two equations:

$$\begin{aligned}\lambda^n &= a_1\varphi_1 + a_2\varphi_2 + a_3\varphi_3\\ \overline{\lambda}^n &= a_1\varphi_4 + a_2\varphi_5 + a_3\varphi_6\end{aligned}$$

Again, depending on the dimension of

$$span\left\{\begin{pmatrix}\varphi_1\\ \varphi_4\end{pmatrix}, \begin{pmatrix}\varphi_2\\ \varphi_5\end{pmatrix}, \begin{pmatrix}\varphi_3\\ \varphi_6\end{pmatrix}\right\}$$

this is either a trivially positive instance, or a lower-dimensional non-degenerate one.

Case III. All equivalence classes of $\sim$ are self-conjugate. As above, we perform a case analysis on $r = n \bmod L$ and pick out 4-tuples of equations. We rely on the fact that stems of classes are distinct real algebraic numbers to ensure that the resulting Skolem instances are non-degenerate and give us the desired bound on n.

6. CONCLUSION

We have shown that the higher-dimensional Orbit Problem is decidable in polynomial time when the target space has dimension one. We have also shown membership in

$\mathbf{NP}^{\mathbf{EqSLP}}$ in the two- and three-dimensional cases. It is known [Schönhage 1979] that $\mathbf{EqSLP} \subseteq \mathbf{coRP}$, so membership in $\mathbf{NP}^{\mathbf{RP}}$ follows immediately.

To obtain these results, we have exploited the connection between Skolem's Problem and the Orbit Problem in fixed dimension. We have shown that each Orbit Problem instance is equivalent to a Master System of equations, from which arises the need to find the zeros of exponential polynomials $\sum_{i=0}^{t} \alpha_i^n p_i(n)$ corresponding to linear recurrence sequences. In quantifying and strengthening the known bounds on n, we have also derived as a by-product a $\mathbf{PTIME}$ upper complexity bound for Skolem's Problem of order two, and an $\mathbf{NP}^{\mathbf{RP}}$ bound for orders three and four.

It is interesting to note that at first glance, the Orbit Problem in fixed dimension appears more difficult that Skolem's Problem for a fixed order, due to its unbounded matrix. On the contrary, the presence of more eigenvalues allows us more freedom when choosing equations from which to obtain a bound on the exponent n. Repetitions in the roots of the minimal polynomial of the matrix simplify matters greatly by bringing n into the base position. A Skolem instance offers no such freedom.

As a further step in research, it would be of interest to consider reachability to different types of targets in a linear system. We are currently working on a version of the Orbit Problem where the target is a convex polytope, that is, an intersection of halfspaces. Whilst the connection to Skolem's Problem is still present, requiring that $A^n x$ lie on one side of a hyperplane also draws a connection to the *Positivity Problem:* given a linear recurrence sequence $u_n := x^T A^n y$, determine whether the set $\{n : u_n \geq 0\}$ is non-empty. A related area of further research is reachability from sets more complex than a single point, as in [Ben-Amram et al. 2012; Braverman 2006; Tiwari 2004]. Finally, a version of the Orbit Problem with continuous time can be formulated [Hainry 2008] and could, we believe, benefit from the study of the continuous Skolem Problem [Bell et al. 2010].

ELECTRONIC APPENDIX

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

Received March 2014; revised ; accepted

A. ALGEBRAIC NUMBERS AND OPERATIONS ON THEM

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 and is denoted by $f_\alpha(x)$. The *degree* of $\alpha \in \mathbb{A}$ is defined as the degree of its minimal polynomial and is denoted by n_α. The *height* of α, denoted by H_α, is defined as the maximum absolute value of the coefficients of the integer polynomial cf_α, where c is the least common multiple of the denominators of the coefficients of f_α. The roots of $f_\alpha(x)$ (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

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

where a, b are respectively the free term and the leading coefficient of $f_\alpha(x)$. It follows that $\mathcal{N}_{abs}(\alpha) \in \mathbb{Q}$. An *algebraic integer* is an algebraic number α such that $f_\alpha \in \mathbb{Z}[x]$. The set of algebraic integers, denoted $\mathcal{O}_\mathbb{A}$, is a ring under the usual addition and multiplication.

The *canonical representation* of an algebraic number α is its minimal polynomial $f_\alpha(x)$, along with a numerical approximation of $Re(\alpha)$ and $Im(\alpha)$ of sufficient precision

to distinguish α from its Galois conjugates. More precisely, we represent α by the tuple

$$(f_\alpha, x, y, R) \in (\mathbb{Q}[x] \times \mathbb{Q} \times \mathbb{Q} \times \mathbb{Q})$$

meaning that α is the unique root of f_α 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 $\alpha_i \neq \alpha_j$ of a polynomial $p(x)$,

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

where n and H are the degree and height of p, respectively. Thus, if R is restricted to be less than a quarter of the root separation bound, the representation is well-defined and allows for equality checking. Observe that given f_α, the remaining data necessary to describe α is polynomial in the length of the input. It is known how to obtain polynomially many bits of the roots of any $p \in \mathbb{Q}[x]$ in polynomial time [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 $\ln|\alpha|$ is polynomially large. Notice also that $1/\ln|\alpha|$ 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 $\sum_{i=0}^{n} \|p_i\|$ where n is the degree of the polynomial and p_i are its coefficients.

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. The resultant of $f_\alpha(x - y)$ and $f_\beta(y)$, 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 $n_\alpha n_\beta$, while its height is bounded by $H_{\alpha+\beta} \leq H_\alpha^{n_\alpha} H_\beta^{n_\beta}$ [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 n_α, 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 A be the companion matrix of f_α. Then $p(\alpha)$ is an eigenvalue of $p(A)$. We can calculate $p(A)$ using only polynomial space. Then from the formula

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

it is evident that the coefficients of the characteristic polynomial of $p(A)$ 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.

B. 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 $[\mathbb{K} : \mathbb{Q}]$. The primitive element theorem states that 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 $n_\theta = [\mathbb{K} : \mathbb{Q}]$. The proof of the primitive element theorem is constructive and shows how to obtain a primitive element for $\mathbb{K} = \mathbb{Q}(\alpha_1, \ldots, \alpha_k)$ given $\alpha_1, \ldots, \alpha_k$. There exist exactly n_θ monomorphisms from $\mathbb{K}$ into $\mathbb{C}$, given by $\theta \to \theta_i$, where θ_i are the Galois conjugates of θ. If $\alpha \in \mathbb{K}$, then $n_\alpha | n_\theta$. Moreover, if $\sigma_1, \ldots, \sigma_{n_\theta}$ are the monomorphisms from $\mathbb{K}$ into $\mathbb{C}$ then $\sigma_1(\alpha), \ldots, \sigma_{n_\theta}(\alpha)$ are exactly the Galois conjugates of α, each repeated n_θ/n_α times. The *norm of* α *relative to* $\mathbb{K}$ is defined as

$$\mathcal{N}_{\mathbb{K}/\mathbb{Q}}(\alpha) = \prod_{i=1}^{n_\theta} \sigma_i(\alpha) = (\mathcal{N}_{abs}(\alpha))^{n_\theta/n_\alpha}$$

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 $\{0\}$. 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_i 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 occurs frequently in what follows, so we state it explicitly here.

LEMMA B.1. *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 $v_P(\alpha) \neq 0$. □

C. BAKER'S THEOREM AND VAN DER POORTEN'S THEOREM

THEOREM C.1. *[Wüstholz and Baker 1993] Let $\alpha_1, \ldots, \alpha_m$ be algebraic numbers other than 0 or 1, and let $b_1, \ldots, b_m$ be rational integers. Write*

$$\Lambda = b_1 \ln \alpha_1 + \ldots + b_m \ln \alpha_m$$

Let $A_1, \ldots, A_m, B \geq e$ be real numbers such that, for each $j \in \{1, \ldots, m\}$, A_j is an upper bound for the height of α_j, and B is an upper bound for $|b_j|$. Let d be the degree of the extension field $\mathbb{Q}(\alpha_1, \ldots, \alpha_m)$ over $\mathbb{Q}$. If $\Lambda \neq 0$, then

$$\ln |\Lambda| > -(16md)^{2(m+2)} \ln(A_1) \ldots \ln(A_m) \ln(B)$$

LEMMA C.2. *Let $\lambda, b \in \mathbb{C}$, 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 $a, \chi \in \mathbb{R}$ such that $0 < \chi < 1$ and $|\phi(n)| \leq a\chi^n$. There exists an effective bound m such that if*

$$\lambda^n = \phi(n) + b \tag{12}$$

then $n < m$. Moreover, if $\lambda, b \in \mathbb{A}$ and $a, \chi \in \mathbb{Q}$ are given as input, then m is at most exponential in the length of the input $L = \|\lambda\| + \|b\| + \|a\| + \|\chi\|$.

PROOF. The left-hand side of (12) 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 (12) will always be off the unit circle. This happens when

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

The difficult case is when b is on the unit circle. We will use Baker's theorem to derive a bound on n. Consider the angle Λ between λ^n and b. This angle can be zero for at most one value of n, because λ is not a root of unity. Otherwise, we have

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

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 $d = \max\{n_\lambda, n_b\}$ is a bound on the degrees. Then by Baker's theorem, we have

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

which is equivalent to

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

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. For n large enough, the right-hand side will forever be smaller in magnitude than the left-hand side.

We will now quantify the bound on n. Let $p_1 = (48d)^{10}\ln^2 H$ and $p_2 = 2$. Observe that if λ and b are canonically represented algebraic numbers, then p_1, p_2 are polynomials in the size of the input. Then (12) cannot hold if

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

which is equivalent to

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

Define $p_3 = \ln(2) + \ln(a) + p_1\ln(p_2)$ and $p_4 = \max\{p_3, p_1\}$ (also polynomials in the size of the input). Then it suffices to have

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

which is guaranteed by

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

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

LEMMA C.3. *Suppose $\lambda_1, \lambda_2, a, b, c \in \mathbb{C}$ 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 $w, \chi \in \mathbb{R}$, $\chi \in (0,1)$. Then there exists a computable bound m such that if*

$$a\lambda_1^n = b\lambda_2^n + c + \phi(n) \tag{13}$$

then $n < m$. Moreover, if $\lambda_1, \lambda_2, a, b, c \in \mathbb{A}$ and $w, \chi \in \mathbb{Q}$ are given, then m is at most exponentially large in the length of the input $\|\lambda_1\| + \|\lambda_2\| + \|a\| + \|b\| + \|c\| + \|w\| + \|\chi\|$.

PROOF. Multiplying the equation by $\bar{c}/|c||a|$ 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

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

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 L_1 be the horizontal line through z_1 and L_2 the horizontal line through z_2. Let $L_1 \cap \mathcal{O}_1 = \{x_1, z_1\}$, $L_1 \cap \mathcal{O}_2 = \{y_1, z_1\}$, $L_2 \cap \mathcal{O}_1 = \{x_2, z_2\}$ and $L_2 \cap \mathcal{O}_2 = \{y_2, z_2\}$. It is trivial that $z_2 = \overline{z_1}$, $x_2 = \overline{x_1}$, $y_2 = \overline{y_1}$.

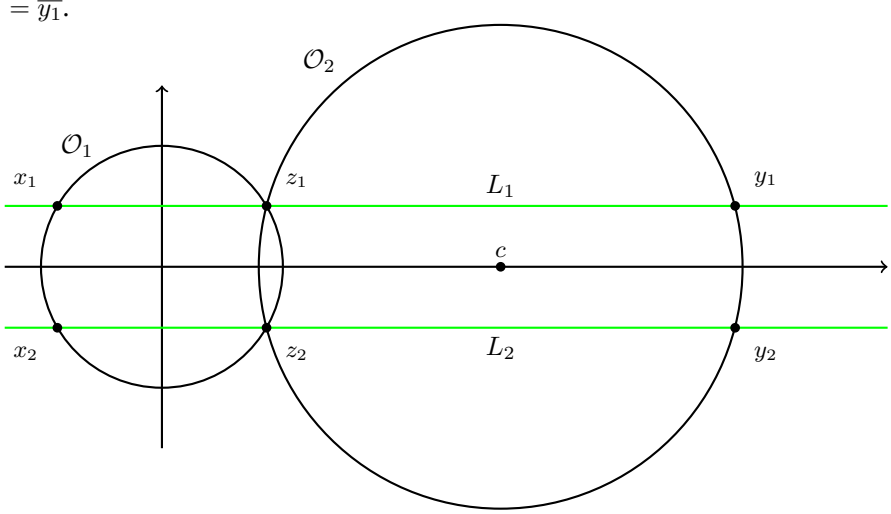

We first argue that for n large enough, (13) can hold only if for some intersection point z_i, $Re(z_i)$ lies between $Re(f(n))$ and $Re(g(n))$, or $Im(z_i)$ lies between $Im(f(n))$ and $Im(g(n))$. This can only be violated in two symmetric situations: either $f(n)$ is on the arc z_1z_2 of $\mathcal{O}_1$ 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 $f(n)$ 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 equality cannot hold for n large enough because $\phi(n)$ tends to zero exponentially quickly. This is the case when

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

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, so we can give lower bounds for them using Baker's theorem as in Lemma C.2:

$$|Re(z_i-f(n))|\geq(p_1n)^{-p_2}$$

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

$$|Re(g(n)-z_i)|=|b|(\cos(n\delta+\eta)-\cos(\theta))\geq(p_3n)^{-p_4}$$

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

$$|f(n)-g(n)|\geq2(p_5n)^{-p_6}$$

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 (13) cannot hold. In the manner of Lemma C.2, 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 cosine theorem 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 $A\leq1+c$, we have $A+B\leq2$, 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

$$\gamma+n\alpha+k_n2\pi\in[-\pi,\pi)$$

A lower bound on this angle follows from Baker's theorem:

$$|\gamma+n\alpha+k_n2\pi|\geq(p_7n)^{-p_8}$$

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 C.2, we can show that the bound on n is exponentially large in the input. □

THEOREM C.4. *[van der Poorten 1977] Let $\alpha_1, \ldots, \alpha_n$ be algebraic numbers of degree at most d belonging to a number field $\mathbb{K}$ and with heights at most $A_1, \ldots, A_n$. Let P be a prime ideal of $\mathbb{K}$ containing the rational prime p. If $\alpha_1^{b_1}\alpha_2^{b_2}\ldots\alpha_n^{b_n} \neq 1$ for rational integers $b_1, \ldots, b_n$ with absolute values at most $B \geq e^2$, then*

$$v_P(\alpha_1^{b_1}\ldots\alpha_n^{b_n} - 1) \leq (16(n+1)d)^{12(n+1)}(p^d/\ln(p))\Omega(\ln(B))^2$$

where $\Omega = \ln(A_1)\ldots\ln(A_n)$.

D. SKOLEM'S PROBLEM, ORDER 2

In this section, we consider the problem of whether a linear recurrence sequence u_n of order 2 contains 0 as an element. The characteristic equation of the recurrence may have one repeated root $\theta \neq 0$, or two distinct roots θ_1, θ_2, giving either

$$u_n = (A + Bn)\theta^n$$

or

$$u_n = A\theta_1^n + B\theta_2^n$$

Solving the problem in the former case is trivial. In the latter case, $u_n = 0$ if and only if $(\theta_1/\theta_2)^n = -B/A$, so this case is an instance of the *algebraic number power problem:* decide whether there exists $n \in \mathbb{N}$ such that

$$\alpha^n = \beta \tag{14}$$

for given $\alpha, \beta \in \mathbb{A}$. The algebraic number power problem is decidable [Halava et al. 2005]. Kannan and Lipton [Kannan and Lipton 1986] proved a polynomial bound on n when β has the form $p(\alpha)$ for a given $p \in \mathbb{Q}[x]$ and α is not a root of unity. We give a brief recapitulation of the decidability proof and extract a polynomial bound on n from it.

LEMMA D.1. *Suppose $\alpha, \beta \in \mathbb{A}$. If α is not a root of unity, then there exists a computable bound m such that if (14) holds, then $n < m$. Moreover, m is polynomial in the length of the input $\|\alpha\| + \|\beta\|$.*

PROOF. Let $\mathbb{K} = \mathbb{Q}(\alpha, \beta)$. If α is not an algebraic integer, then by Lemma B.1 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 α is an algebraic integer. It is not a root of unity, so by Kronecker's theorem [Kronecker 1875], α has a Galois conjugate $\sigma(\alpha)$ with magnitude strictly greater than 1. In fact, a significant strengthening of Kronecker's theorem, due to Blanksby and Montgomery [Blanksby and Montgomery 1971], guarantees the existence of a conjugate $\sigma(\alpha)$ such that

$$|\sigma(\alpha)| > 1 + \frac{1}{30 n_\alpha^2 \ln(6 n_\alpha)}$$

which implies

$$\frac{1}{\ln |\sigma(\alpha)|} < 60 n_\alpha^2 \ln(6 n_\alpha)$$

Then if $\alpha^n = \beta$, we have

$$n = \frac{\ln |\sigma(\beta)|}{\ln |\sigma(\alpha)|} < \ln |\sigma(\beta)| 60 n_\alpha^2 \ln(6 n_\alpha)$$

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

The condition that α not be a root of unity is obviously necessary in Lemma D.1, because if β is also a root of unity, $\alpha^n = \beta$ could hold infinitely often. Indeed it is easy to exhibit linear recurrences of order 2 with infinitely many zeroes (for example, $u_1 = 0$, $u_2 = 1$, $u_{n+2} = u_n$), but by the Lemma, they all have two distinct roots whose ratio is a root of unity.

E. SKOLEM'S PROBLEM, ORDER 3

In this section we will focus on Skolem's Problem for linear recurrence sequences of order 3. The characteristic equation of such a sequence may have either three distinct roots α, β, γ, or one repeated real root α and one simple real root β, or one thrice repeated real root. Finding the zeroes of the latter type of linear recurrence is trivial, so we focus on the former two possibilities.

If the three roots are distinct, we are concerned with solving for $n \in \mathbb{N}$ equations of the form

$$A\alpha^n + B\beta^n + C\gamma^n = 0 \tag{15}$$

where $A, B, C, \alpha, \beta, \gamma \in \mathbb{A}$ are given and non-zero (if any of them is 0, then the sequence satisfies a recurrence relation of smaller order). Then (15) is equivalent to

$$\left(\frac{\beta}{\alpha}\right)^n = -\frac{C}{B}\left(\frac{\gamma}{\alpha}\right)^n - \frac{A}{B} \tag{16}$$

We will consider only non-degenerate sequences: the ratios of the roots α, β, γ are not roots of unity. Let also $|\alpha| \geq |\beta| \geq |\gamma|$. In Lemmas E.1, E.2, E.3 below, the length of the input is $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\| + \|\gamma\|$.

LEMMA E.1. *If $|\alpha| > |\beta|$, then there exists an effective bound m such that if equation (15) holds, then $n < m$. Moreover, m is at most exponential in the length of input.*

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

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

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 E.2. *If $|\alpha| = |\beta| > |\gamma|$, then there exists an effective bound m such that if equation (15) holds, then $n < m$. Moreover, m is at most exponential in the length of the input.*

PROOF. This is a direct application of Lemma C.2 to equation (16). □

LEMMA E.3. *If $|\alpha| = |\beta| = |\gamma|$, then there exist at most two values of n such that equation (15) holds. Moreover, they are at most exponential in the length of the input and are computable in polynomial time.*

PROOF. The left-hand side of (16) 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 (15) 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 recurrence sequences of order 3 with one repeated and one simple root. We are given $A, B, C, \alpha, \beta \in \mathbb{A}$, and the length of the input is $\|A\| + \|B\| + \|C\| + \|\alpha\| + \|\beta\|$. We wish to solve for n

$$(A + Bn)\alpha^n + C\beta^n = 0 \tag{17}$$

We will assume that B, C, α, β are all non-zero, otherwise the sequence on the left-hand side of (17) satisfies a linear recurrence of lower order.

LEMMA E.4. *There exists an effective bound m such that if (17) holds, then $n < m$. Moreover, m is at most exponential in the length of the input.*

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

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

we have

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

so

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

and (17) cannot hold. Now suppose $|\alpha| > |\beta|$ and rewrite (17) as

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

Equation (17) 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

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

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

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

to guarantee that (17) cannot hold. This is an exponential bound on n in the length of the input. □

F. SKOLEM'S PROBLEM, ORDER 4

In this section we will give lemmas which form a decidability proof for Skolem's Problem for order 4. Note that the problem is not known to be decidable for linear recurrence sequences of order 4 over $\mathbb{A}$, so we restrict ourselves to the situations which arise in the rational case. Assume algebraic numbers A, B, C, D and $\alpha, \beta, \gamma, \delta$ are given and the input has length $||I|| = \|A\| + \|B\| + \|C\| + \|D\| + \|\alpha\| + \|\beta\| + \|\gamma\| + \|\delta\|$. We wish to solve for n the following equations:

$$A\alpha^n + B\beta^n + C\gamma^n + D\delta^n = 0 \text{ (where } A, B, C, D \neq 0) \tag{18}$$

$$(A + Bn)\alpha^n + C\beta^n + D\gamma^n = 0 \text{ (where } B, C, D \neq 0) \tag{19}$$

$$(A + Bn)\alpha^n + (C + Dn)\beta^n = 0 \text{ (where } B, D \neq 0) \tag{20}$$

$$(A + Bn + Cn^2)\alpha^n + D\beta^n = 0 \text{ (where } C, D \neq 0) \tag{21}$$

$$(A + Bn + Cn^2 + Dn^3)\alpha^n = 0 \text{ (where } D \neq 0) \tag{22}$$

As before, we assume that the ratio of any distinct pair of $\alpha, \beta, \gamma, \delta$ is not a root of unity and that $\alpha, \beta, \gamma, \delta$ are all non-zero. Solving (22) is trivial. We can rearrange (21) 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 $|\alpha| < |\beta|$ or not. 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 equations (18)(19)(20) are more involved.

LEMMA F.1. *There exists an effective bound $m = 2^{||I||^{\mathcal{O}(1)}}$ such that if equation (20) holds, then $n < m$.*

PROOF. Rearrange (20) as

$$\lambda^n = -\frac{(C + Dn)}{(A + Bn)} \tag{23}$$

where $\lambda = \alpha/\beta$ is not a root of unity. The right-hand side of (23) tends to $-D/B$ as n tends to infinity.

If λ is an algebraic integer, then by Blanksby and Montgomery's theorem [Blanksby and Montgomery 1971], it has a Galois conjugate $\sigma(\lambda)$ such that

$$|\sigma(\lambda)| > 1 + \frac{1}{30n_\lambda^2 \ln(6n_\lambda)}$$

Assume the monomorphism σ has been applied to both sides of (23), so $|\lambda|$ is bounded away from 1 by an inverse polynomial in the size of the input. By the triangle inequality, if

$$n \geq \frac{|BC| + |AD| + |AB|}{|B|^2} \stackrel{def}{=} N_1 = 2^{||I||^{\mathcal{O}(1)}}$$

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 computable $N_2 \in ||I||^{\mathcal{O}(1)}$ such that if $n > N_2$, then $|\lambda^n| > |D/B| + 1$. Therefore, for $n > \max\{N_1, N_2\} = 2^{||I||^{\mathcal{O}(1)}}$, equation (23) cannot hold.

Second, suppose λ is not an algebraic integer. Then by Lemma B.1 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 (23) gives

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

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

$$n > \frac{|\sigma_i(BC)| + |\sigma_i(AD)| + |\sigma_i(AB)|}{|\sigma_i(B)|^2} \stackrel{def}{=} N_i = 2^{||I||^{\mathcal{O}(1)}}$$

then we have

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

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

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

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 (23) cannot hold. □

LEMMA F.2. *There exists an effective bound m such that if equation (19) holds, then $n < m = 2^{||I||^{\mathcal{O}(1)}}$.*

PROOF. First suppose $|\alpha| \geq |\beta|, |\gamma|$. Then the term $(A+Bn)\alpha^n$ is dominant. More precisely, rewrite (19) 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 (19) 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 (19) as

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

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 (24) impossible. The former inequality holds for $n > \ln|C/2D|/\ln|\gamma/\beta|$, 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 $r = \lceil -\ln(2)/\ln(\alpha/\beta) \rceil$, 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

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

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 (19) 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 C.2, because the right-hand side is a constant plus an exponentially decaying term, whereas the left-hand side is on unit circle. □

LEMMA F.3. *If $\alpha, \beta, \gamma, \delta$ do not all have the same magnitude, then there exists an effective bound m such that if equation (18) holds, then $n < m = 2^{||I||^{\mathcal{O}(1)}}$.*

PROOF. Let $|\alpha| \geq |\beta| \geq |\gamma| \geq |\delta|$. First, if $|\alpha| > |\beta|$, then $A\alpha^n$ is the dominant term in (18). 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

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

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 (18) as

$$\left(\frac{\beta}{\alpha}\right)^n = -\frac{A}{B} - \frac{C}{B}\left(\frac{\gamma}{\alpha}\right)^n - \frac{D}{B}\left(\frac{\delta}{\alpha}\right)^n \tag{25}$$

The left-hand side of (25) is on the unit circle, whereas the right is a constant plus exponentially decaying terms. An exponential bound on n follows from Lemma C.2.

Finally, if $|\alpha| = |\beta| = |\gamma| > |\delta|$, then an exponential bound on n follows from Lemma C.3 applied to equation (25). □

Thus, the only outstanding problem is to solve (18) when $|\alpha| = |\beta| = |\gamma| = |\delta|$. This is difficult for general $\alpha, \beta, \gamma, \delta$, so we will restrict ourselves to two sufficient special cases: when at least two of them are real, or when they are two pairs of complex conjugates. We will also assume that they are *algebraic integers*. This is sufficient for our application to the Orbit Problem, because any Orbit instance $\exists n. A^n x \in V$ may be reduced in polynomial time to an instance where A is an integer matrix, so that its eigenvalues are algebraic integers.

LEMMA F.4. *Let $\alpha, \beta, \gamma, \delta$ be algebraic integers with $|\alpha| = |\beta| = |\gamma| = |\delta|$. If $\alpha, \beta \in \mathbb{R}$ or if (α, β) and (γ, δ) are pairwise complex conjugates, there exists an effective bound m such that if equation (18) holds, then $n < m$. Moreover, m is exponential in the length of the input.*

PROOF. Let $\mathbb{K} = \mathbb{Q}(\alpha, \beta, \gamma, \delta, A, B, C, D)$. First suppose that $\alpha, \beta \in \mathbb{R}$. If $\alpha = \beta$, then we have a Skolem instance of order 3

$$\exists n. (A + B)\alpha^n + C\gamma^n + D\gamma^n = 0$$

An exponential bound on n follows from Lemma E.3. If $\alpha = -\beta$, then we consider even n and odd n separately, obtaining two Skolem instances of order 3 of the same type.

Now suppose $\beta = \overline{\alpha}$ 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 (18) leads to a Skolem instance of order 4 with roots whose magnitudes are not all the same. A bound on n follows from Lemma F.3.

Suppose then that α/β is not an algebraic integer. By the reasoning of Lemma B.1, 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 (18) as

$$A\alpha^n + B\beta^n = -C\gamma^n - D\delta^n \tag{26}$$

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 (26) is strictly smaller under v_P than the right-hand side, so (18) 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 (18) as

$$B\beta^n \left(\left(-\frac{C}{B} \right) \left(\frac{\gamma}{\beta} \right)^n - 1 \right) = A\alpha^n + D\delta^n \tag{27}$$

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 may apply Theorem C.4 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 C.4, we have

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

It is classical $\mathcal{N}(P) = p^f$ 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 the inequality from Theorem C.4 as

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

where E_1 is exponentially large in the input size and independent of n. Now we apply v_P to both sides of equation (27):

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

and

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

Equation (18) cannot hold if

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

which is implied by

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

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

$$E_2((\ln n)^2 + 1) < n$$

Since

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

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.

Paul C. Bell, Jean-Charles Delvenne, Raphaël M. Jungers, and Vincent D. Blondel. 2010. The continuous Skolem-Pisot problem. *Theoretical Computer Science* 411, 40-42 (2010), 3625–3634. http://www.sciencedirect.com/science/article/pii/S0304397510003397

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.

J. Berstel. 1974. *Deux Proprietes Decidables des Suites Redcurrentes Lineaires*. Centre de Calcule de L Esplanade -, U.E.R. Mathematique. http://books.google.co.uk/books?id=xFdrQwAACAAJ

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

V. Blondel and N. Portier. 2002. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. (2002).

Mark Braverman. 2006. Termination of integer linear programs. In *Computer aided verification*. 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.

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. *TUCS Technical Report* 683 (2005).

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 *STOC*. 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.

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

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.

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.

Sergey P. Tarasov and Mikhail N. Vyalyi. 2010. Orbits of linear maps and regular languages. *CoRR* abs/1011.1842 (2010).

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.

Gisbert Wüstholz and Alan Baker. 1993. Logarithmic forms and group varieties. *Journal für die reine und angewandte Mathematik* 442 (1993), 19–62. http://eudml.org/doc/153550

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