On the Complexity of the Orbit Problem

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

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

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. 2016. On the complexity of the orbit problem. J. ACM 63, 3, Article 23 (June 2016), 18 pages.
DOI: <http://dx.doi.org/10.1145/2857050>

1. INTRODUCTION

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

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

The decidability of this problem remained open for over 10 years, until it was shown to be decidable in polynomial time by 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 $\mathbf{A} \in \mathbb{Q}^{m \times m}$, a vector $\mathbf{x} \in \mathbb{Q}^m$, and a subspace $\mathcal{V}$ of $\mathbb{Q}^m$, decide whether there exists a non-negative integer n such that $\mathbf{A}^n \mathbf{x} \in \mathcal{V}$.

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

Authors' addresses: V. Chonev, IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria; J. Ouaknine and J. Worrell, Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD, UK.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2016 ACM 0004-5411/2016/06-ART23 \$15.00

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

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

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

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

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

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

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

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

Then, with target space

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

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

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

The key observation is that the eigenvalues of $\mathbf{A}$ 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 $\mathbf{A}^n \mathbf{x} \in \mathcal{V}$. Computable bounds on such an n are obtained utilising the work of Mignotte et al. [1984] and Vereshchagin [1985], quantifying and strengthening some of the bounds given for the Skolem Problem. In the case of a one-dimensional target subspace $\mathcal{V}$, the resulting bound is polynomial in the size of the problem representation, allowing for all exponents n up to the bound to be checked directly and yielding a polynomial-time algorithm. Unfortunately, when $\mathcal{V}$ has dimension two or three, the bounds on n are exponential in the size of the input, leading to an $\mathbf{NPRP}^{\text{guess-and-check}}$ procedure, in which an $\mathbf{RP}$ oracle is used to check whether $\mathbf{A}^n \mathbf{x} \in \mathcal{V}$ for a guessed value of n . Finally, the case in which the eigenvalues of $\mathbf{A}$ 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 $\mathbf{A}^n \mathbf{x} \in \mathcal{V}$. Noting that there are at most exponentially many such residue classes, we can directly incorporate this case analysis into an $\mathbf{NPRP}$ algorithm.

1.1. Related Work

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

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

2. OUTLINE OF ARTICLE

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

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

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

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

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

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

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

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

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

²Indeed, the eigenvalues of $\mathbf{A}^L$ are exactly λ_i^L where λ_i are the eigenvalues of $\mathbf{A}$. If for any two distinct such eigenvalues, say, $\lambda_i^L \neq \lambda_j^L$, we have $(\lambda_i^L / \lambda_j^L)^t = 1$, 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 $\lambda_i^L = \lambda_j^L$.

instances is positive. Unfortunately, this reduction to the non-degenerate case carries an exponential overhead, as L is, in general, exponentially large in the size of the input data.

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

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

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

3. REDUCTION

3.1. Matrix Power Problem

Suppose we are given a matrix $A \in \mathbb{Q}^{m \times m}$, a vector $\mathbf{x} \in \mathbb{Q}^m$ and a target vector space $\mathcal{V} \subseteq \mathbb{Q}^m$ specified by a basis of rational vectors $\mathbf{y}_1, \dots, \mathbf{y}_k$. We wish to decide whether there exists $n \in \mathbb{N}$ such that $A^n \mathbf{x} \in \mathcal{V}$.

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 $v = \max\{m \mid \mathbf{x}, A\mathbf{x}, \dots, A^m \mathbf{x} \text{ are linearly independent}\}$, $B = \{\mathbf{x}, A\mathbf{x}, \dots, A^v \mathbf{x}\}$, $\mathcal{U} = \text{span}(B)$ and $D = [\mathbf{x} \ A\mathbf{x} \ \dots \ A^v \mathbf{x}]$. It is clear that $\mathcal{U}$ is invariant under the linear transformation A , so consider the restriction of A to $\mathcal{U}$. Suppose $\mathbf{b} = (b_0, \dots, b_v)^T$ are

the coordinates of $\mathbf{A}^{v+1}\mathbf{x}$ with respect to B , that is, $\mathbf{A}^{v+1}\mathbf{x} = \mathbf{D}\mathbf{b}$. The restriction of $\mathbf{A}$ to $\mathcal{U}$ with respect to the basis B is described by the matrix

$$\mathbf{M} = \begin{bmatrix} 0 & 0 & \dots & 0 & b_0 \\ 1 & 0 & \dots & 0 & b_1 \\ 0 & 1 & \dots & 0 & b_2 \\ & & \ddots & & \vdots \\ 0 & 0 & \dots & 1 & b_v \end{bmatrix}.$$

It is easy to check that $\mathbf{D}\mathbf{M} = \mathbf{A}\mathbf{D}$. Thus, if some vector $\mathbf{z}$ has coordinates $\mathbf{z}'$ with respect to B , so $\mathbf{z} = \mathbf{D}\mathbf{z}'$, then $\mathbf{A}\mathbf{z}$ has coordinates $\mathbf{M}\mathbf{z}'$ with respect to B , so $\mathbf{A}\mathbf{z} = \mathbf{D}\mathbf{M}\mathbf{z}'$. By induction, for all $n \in \mathbb{N}$, $\mathbf{A}^n\mathbf{x} = \mathbf{D}\mathbf{M}^n\mathbf{x}'$, where $\mathbf{x}' = (1, 0, \dots, 0)^T$. Next we calculate a basis $\{\mathbf{w}_1, \dots, \mathbf{w}_t\}$ for $\mathcal{W} \stackrel{\text{def}}{=} \mathcal{U} \cap \mathcal{V}$. Then, if $\mathbf{w}_i$ are such that $\mathbf{w}_i = \mathbf{D}\mathbf{w}'_i$ for all i , then we have

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

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

Define the matrices $\mathbf{T}_1, \dots, \mathbf{T}_t$ by

$$\mathbf{T}_i = [\mathbf{w}'_i \quad \mathbf{M}\mathbf{w}'_i \quad \dots \quad \mathbf{M}^v\mathbf{w}'_i].$$

We will show that $\mathbf{M}^n\mathbf{x}' \in \text{span}\{\mathbf{w}'_1, \dots, \mathbf{w}'_t\}$ if and only if $\mathbf{M}^n \in \text{span}\{\mathbf{T}_1, \dots, \mathbf{T}_t\}$. If for some coefficients κ_i we have

$$\mathbf{M}^n = \sum_{i=1}^t \kappa_i \mathbf{T}_i,$$

then, considering the first column of both sides, we have

$$\mathbf{M}^n\mathbf{x}' = \sum_{i=1}^t \kappa_i \mathbf{w}'_i.$$

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

Thus, we have reduced the Orbit Problem to the *Matrix Power Problem*: determining whether some power of a given matrix lies inside a given vector space of matrices.

Now we will perform a further reduction step. It is clear that within the space $\mathcal{T} \stackrel{\text{def}}{=} \text{span}\{\mathbf{T}_1, \dots, \mathbf{T}_t\}$ it suffices to consider only matrices of the shape $P(\mathbf{M})$ where $P \in \mathbb{Q}[x]$.

We find a basis for the space $\mathcal{P} \stackrel{\text{def}}{=} \{P(\mathbf{M}) \mid P \in \mathbb{Q}[x]\}$ and then a basis $\{P_1(\mathbf{M}), \dots, P_s(\mathbf{M})\}$ for $\mathcal{P} \cap \mathcal{T}$. Then $\mathbf{M}^n \in \mathcal{T} \iff \mathbf{M}^n \in \mathcal{P} \cap \mathcal{T}$. We call the problem of determining, given $\mathbf{M}$ and $P_1, \dots, P_s$, whether there exists $n \in \mathbb{N}$ such that $\mathbf{M}^n \in \text{span}\{P_1(\mathbf{M}), \dots, P_s(\mathbf{M})\}$, the *polynomial version* of the matrix power problem. Observe that $\dim(\mathcal{V}) \geq \dim(\mathcal{T}) \geq \dim(\mathcal{T} \cap \mathcal{P})$, so the dimension of the target vector space does not grow during the described reductions. All described operations may be performed in polynomial time using standard techniques from linear algebra.

3.2. Master System of Equations

Suppose now we have an instance $(\mathbf{A}, P_1, \dots, P_s)$ of the polynomial version of the matrix power problem. Calculate the minimal polynomial of $\mathbf{A}$ and obtain canonical

representations of its roots $\alpha_1, \dots, \alpha_k$, that is, the eigenvalues of $\mathbf{A}$. This may be done in polynomial time, see Section A.1. Throughout, for an eigenvalue α_i , we will denote by $\text{mul}(\alpha_i)$ the multiplicity of α_i as a root of the minimal polynomial of the matrix.

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

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

if and only if

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

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

Thus, in order to decide whether there exists an exponent n and coefficients κ_i such that $\mathbf{A}^n = \sum_{i=1}^s \kappa_i P_i(\mathbf{A})$, it is sufficient to solve the system of Equations (1) where the unknowns are $n \in \mathbb{N}$ and $\kappa_1, \dots, \kappa_s \in \mathbb{C}$. Each eigenvalue α_i contributes $\text{mul}(\alpha_i)$ equations that specify that $P(x)$ and its first $\text{mul}(\alpha_i) - 1$ 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 $0 \leq j < \text{mul}(\alpha_i)$. We will also denote by $Eq(\alpha_i)$ the set of equations contributed by α_i to the system:

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

For example, if the minimal polynomial of $\mathbf{A}$ has roots $\alpha_1, \alpha_2, \alpha_3$ with multiplicities $\text{mul}(\alpha_i) = i$ and the target space is $\text{span}\{P_1(\mathbf{A}), P_2(\mathbf{A})\}$, then the system contains six equations:

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

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

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

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

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

4. ONE-DIMENSIONAL TARGET SPACE

Suppose we are given a one-dimensional matrix power problem instance $(\mathbf{A}, P)$ and wish to decide whether $\mathbf{A}^n \in \text{span}\{P(\mathbf{A})\}$ for some n . We have constructed a system of equations in the exponent n and the coefficient κ as in Equation (1). For example, if the roots of the minimal polynomial of $\mathbf{A}$ are $\alpha_1, \alpha_2, \alpha_3$ with multiplicities $\text{mul}(\alpha_j) = j$,

then the system is

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

We solve the first conjunct as in Case 2. If there is no solution, then we are done. Otherwise, the solution is some congruence $n \equiv t_1 \pmod{t_2}$. 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

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

For each such Equation (4), we calculate the right-hand side in polynomial time, using the methods outlined in Section A.1 and check whether it is in $\mathbb{N}$. If not, then the system has no solution. Otherwise, Equation (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

$$n \equiv t_1 \pmod{t_2}$$

$$n = n_0$$

$$\kappa = \alpha_1^n / P(\alpha_1).$$

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

5. TWO-DIMENSIONAL TARGET SPACE

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

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

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

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

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

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

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

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

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

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

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

$$C_i = \{\bar{\alpha} \mid \alpha \in C_i\}.$$

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

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

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

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

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

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

Let

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

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

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

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

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

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

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

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

$$\begin{bmatrix} \alpha^n \\ \vdots \\ \alpha^n \end{bmatrix} = \mathbf{B} \begin{bmatrix} \kappa_1 \\ \kappa_2 \end{bmatrix}, \quad (5)$$

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

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

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

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

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

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

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

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

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

$$\begin{aligned} A_1\alpha^n + A_2n\alpha^{n-1} + A_3n(n-1)\alpha^{n-2} &= 0 \\ \iff A_1\alpha^2 + A_2n\alpha + A_3n(n-1) &= 0. \end{aligned}$$

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

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

We divide the first equation by the second to obtain

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

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

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

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

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

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

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

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

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

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

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

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

$$C_2 = \{\beta\omega'_1, \dots, \beta\omega'_l\}.$$

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

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

$$\begin{aligned}\alpha^n &= \kappa_1 c_1 + \kappa_2 c_2 \\ \beta^n &= \kappa_1 c_3 + \kappa_2 c_4.\end{aligned}$$

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

$$\begin{aligned}\alpha^n &= \kappa_1 c_1 \\ \beta^n &= \kappa_1 c_3\end{aligned}$$

and, hence,

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

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

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

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

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

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

Then, dividing, we obtain

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

which limits n to at most one, exponentially large candidate value $\alpha c_3/c_1$. Otherwise, if the vectors on the right-hand side of Equation (8) are linearly independent over $\mathbb{A}$, then 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 that is exponential in the size of the input follows from Lemma F.4. This again relies on the fact that α/β cannot be a root of unity.

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

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

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

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

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

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

which gives

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

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

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

6. THREE-DIMENSIONAL TARGET SPACE

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

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

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

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

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

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

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

$$A_1\alpha^n + A_2\beta^n + A_3\gamma^n + A_4\delta^n = 0 \quad (10)$$

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

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

$$\begin{bmatrix} \alpha^n \\ \overline{\alpha}^n \\ n\alpha^{n-1} \\ n\overline{\alpha}^{n-1} \end{bmatrix} = \kappa_1 \begin{bmatrix} P_1(\alpha) \\ P_1(\overline{\alpha}) \\ P'_1(\alpha) \\ P'_1(\overline{\alpha}) \end{bmatrix} + \kappa_2 \begin{bmatrix} P_2(\alpha) \\ P_2(\overline{\alpha}) \\ P'_2(\alpha) \\ P'_2(\overline{\alpha}) \end{bmatrix} + \kappa_3 \begin{bmatrix} P_3(\alpha) \\ P_3(\overline{\alpha}) \\ P'_3(\alpha) \\ P'_3(\overline{\alpha}) \end{bmatrix}.$$

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

$$A_1\alpha^n + A_2\overline{\alpha}^n + A_3n\alpha^{n-1} + A_4n\overline{\alpha}^{n-1} = 0.$$

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

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

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

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

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

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

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

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

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

$$\begin{aligned}\lambda^n &= \kappa_1c_1 + \kappa_2c_2 + \kappa_3c_3 \\ \bar{\lambda}^n &= \kappa_1c_4 + \kappa_2c_5 + \kappa_3c_6 \\ \alpha^n &= \kappa_1c_7 + \kappa_2c_8 + \kappa_3c_9 \\ n\alpha^{n-1} &= \kappa_1c_{10} + \kappa_2c_{11} + \kappa_3c_{12}\end{aligned}$$

to obtain the non-degenerate instance

$$A_1\lambda^n + A_2\bar{\lambda}^n + A_3\alpha^n + A_4n\alpha^{n-1} = 0,$$

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

$$\begin{aligned}\lambda^n &= \kappa_1c_1 + \kappa_2c_2 + \kappa_3c_3 \\ \bar{\lambda}^n &= \kappa_1c_4 + \kappa_2c_5 + \kappa_3c_6 \\ \alpha^n &= \kappa_1c_7 + \kappa_2c_8 + \kappa_3c_9.\end{aligned}$$

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

$$\begin{aligned}\lambda^n &= \kappa_1c_1 + \kappa_2c_2 + \kappa_3c_3 \\ \bar{\lambda}^n &= \kappa_1c_4 + \kappa_2c_5 + \kappa_3c_6.\end{aligned}$$

Again, depending on the dimension of

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

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

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

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

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

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

ELECTRONIC APPENDIX

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

REFERENCES

- V. Arvind and T. Vijayaraghavan. 2011. The orbit problem is in the GapL hierarchy. *J. Comb. Optim.* 21, 1 (2011), 124–137.
- Alan Baker. 1975. *Transcendental Number Theory*. Cambridge University Press, Cambridge.
- A. Baker and G. Wüstholz. 1993. Logarithmic forms and group varieties. *J. Reine Angew. Math.* 442 (1993), 19–62.
- Amir M. Ben-Amram, Samir Genaim, and Abu Naser Masud. 2012. On the termination of integer loops. *ACM Trans. Program. Lang. Syst.* 34, 4 (2012), 16.
- Jean Berstel and Maurice Mignotte. 1976. Deux propriétés décidables des suites récurrentes linéaires. *Bull. Soc. Math. France* 104 (1976), 175–184.
- P. Blanksby and H. Montgomery. 1971. Algebraic integers near the unit circle. *Acta Arith.* (1971), 355–369.
- V. D. Blondel and N. Portier. 2002. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. *Linear Algebra and Its Applications* 351 (2002), 91–98.
- M. Braverman. 2006. Termination of integer linear programs. In *Proceedings of the 18th International Conference on Computer Aided Verification (CAV, LNCS 4144)*. Springer, Berlin, 372–385.
- Jin-yi Cai, Richard J. Lipton, and Yechezkel Zalcstein. 2000. The complexity of the ABC problem. *SIAM J. Comput.* 29, 6 (2000), 1878–1888.
- Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2013. The orbit problem in higher dimensions. In *STOC*. ACM, New York, NY, 941–950.
- Ventsislav Chonev, Joël Ouaknine, and James Worrell. 2015. The polyhedron-hitting problem. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'15)*. SIAM, Philadelphia, PA, 940–956. <http://dl.acm.org/citation.cfm?id=2722129.2722193>.
- H. Cohen. 1993. *A Course in Computational Algebraic Number Theory*. Springer, Berlin.
- Graham Everest, Alf van der Poorten, Thomas Ward, and Igor Shparlinski. 2003. *Recurrence Sequences*. American Mathematical Society, Washington DC.
- Emmanuel Hainry. 2008. Reachability in linear dynamical systems. In *Logic and Theory of Algorithms*. Springer, Berlin, 241–250.
- V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. 2005. *Skolem's Problem – On the Border between Decidability and Undecidability*. Technical Report 683. Turku Centre for Computer Science.
- G. Hansel. 1986. Une démonstration simple du théorème de Skolem-Mahler-Lech. *Theor. Comput. Sci.* 43 (1986), 91–98.
- Michael A. Harrison. 1969. *Lectures on Linear Sequential Machines*. Academic Press, New York, NY.
- R. Kannan and R. Lipton. 1986. Polynomial-time algorithm for the orbit problem. *J. ACM* 33, 4 (1986), 808–821.
- Ravindran Kannan and Richard J. Lipton. 1980. The orbit problem is decidable. In *Proceedings of the 12th Annual ACM Symposium on Theory of Computing (STOC)*. ACM, New York, NY, 252–261.
- L. Kronecker. 1875. Zwei sätze über gleichungen mit ganzzahligen koeffizienten. *J. Reine Angew. Math.* 53 (1875), 173–175.
- Christer Lech. 1953. A note on recurring series. *Arkiv för Matematik* 2 (1953), 417–421.

- A. K. Lenstra, H. W. Lenstra, and L. Lovász. 1982. Factoring polynomials with rational coefficients. *Math. Ann.* 261 (1982), 515–534.
- B. Litow. 1997. A decision method for the rational sequence problem. In *Electronic Colloquium on Computational Complexity (ECCC)*, Vol. 4.
- K. Mahler. 1935. Eine arithmetische eigenschaft der Taylor-koeffizienten rationaler funktionen. *Proc. Akad. Wet. Amsterdam* 38 (1935), 51–60.
- K. Mahler and J. W. S. Cassels. 1956. On the Taylor coefficients of rational functions. *Math. Proc. Cambr. Philos. Soc.* 52 (1 1956), 39–48. Issue 01. DOI: <http://dx.doi.org/10.1017/S0305004100030966>
- M. Mignotte. 1982. Some useful bounds. *Comput. Algebr.* (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, 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. *Comput. Math. Appl.* 31, 12 (1996), 97–138.
- Arto Salomaa and Matti Soittola. 1978. *Automata—Theoretic Aspects of Formal Power Series*. Springer-Verlag, Berlin.
- Arnold Schönhage. 1979. On the power of random access machines. In *Automata, Languages and Programming*, Hermann Maurer (Ed.). Lecture Notes in Computer Science, Vol. 71. Springer, Berlin, 520–529.
- Th. Skolem. 1934. Ein verfahren zur behandlung gewisser exponentialer gleichungen und diophantischer gleichungen. *Skand. Mat. Kongr.* 8 (1934), 163–188.
- I. Stewart and D. Tall. 2002. *Algebraic Number Theory and Fermat’s Last Theorem* (3rd ed.). A. K. Peters.
- T. Tao. 2008. *Structure and Randomness: Pages from Year One of a Mathematical Blog*. American Mathematical Society, Washington, DC.
- Sergey Tarasov and Mikhail Vyalyi. 2011. Orbits of linear maps and regular languages. In *Computer Science – Theory and Applications*, Alexander Kulikov and Nikolay Vereshchagin (Eds.). Lecture Notes in Computer Science, Vol. 6651. Springer, Berlin, 305–316. DOI: http://dx.doi.org/10.1007/978-3-642-20712-9_24
- Ashish Tiwari. 2004. Termination of linear programs. In *Computer Aided Verification*. Springer, Berlin, 70–82.
- Alfred Jacobus van der Poorten. 1977. Linear forms in logarithms in the p-adic case. *Transcend. Theor. Adv. Appl.* (1977), 29–57.
- N. Vereshchagin. 1985. Occurrence of zero in a linear recursive sequence. *Math. Notes* 38 (1985), 609–615.
- Richard Zippel. 1997. Zero testing of algebraic functions. *Inform. Process. Lett.* 61, 2 (1997), 63–67.

Received April 2014; revised November 2015; accepted December 2015