The Orbit Problem in Zero and One Dimensions

Master's dissertation of Ventsislav K. Chonev

Supervised by Joël Ouaknine and James Worrell

Contents

Contents **2**

1 Introduction **3**
1.1 Verification of linear programs 3
1.2 Overview of this project 5

2 Mathematical Foundations **7**
2.1 Basic definitions and properties 7
2.2 Algebraic number fields 11
2.3 Representation of algebraic numbers 11
2.4 Monomorphisms 15
2.5 Magnitude bounds 16
2.6 Algebraic Number Power problem 18

3 The Zero-dimensional Orbit Problem **22**
3.1 Reduction 22
3.2 Solution 25

4 The One-dimensional Orbit Problem **29**
4.1 Reduction 29
4.2 Solution 31

5 Conclusion **36**

Bibliography **37**

Chapter 1

Introduction

1.1 Verification of linear programs

This project falls within the scope of Computer-Aided Verification. The field is concerned with developing methods to analyse software and hardware systems. This direction of research is motivated by the idea that as systems grow in scale and complexity, it becomes impossible to prove manually that they satisfy desirable properties. Manual inspection is expensive, because it requires human input, but is simultaneously error-prone. Testing is widely applied in the industry, but is incomplete and can only be used to show the existence of errors, never to prove their absence. Thus, the need for automated methods for verification is evident.

Verifying programs is inherently difficult. One of the earliest results of computability theory is the undecidability of the Halting Problem. A more general result, Rice's Theorem, states that all non-trivial properties of Turing Machines are undecidable. However, an important point is that these results place no restrictions on the program to be verified, allowing any Turing Machine as input. Another point is that the verification procedure is required to terminate and produce a correct answer. Therefore, modern research in software verification lifts these restrictions and focuses on two main directions:

1. Partial verification procedures, which are allowed to return no answer.
2. Complete verification procedures for restricted classes of programs.

One restricted class of programs is *linear programs*, which consist of unnested while loops defined in terms of linear or affine guards and assignments, such as

$$\textbf{while } (Bx > 0) \textbf{ do } x := Ax + c \textbf{ end}$$

where A, B are matrices of constants, c is a vector of constants and x is a vector of program variables.

Within this class, there is a wide variety of subclasses that may be considered, depending on the domain of the relevant variables (natural, integer, rational, algebraic, real, complex), the type of constraints in the loop guard (equalities, disequalities, strict and non-strict inequalities), the type of assignments (linear or affine), the dimension of the relevant state space. For linear programs written in matrix form, one may also consider special cases of matrices (stochastic, symmetric, diagonalisable, unitary).

Linear programs have elicited a considerable amount of interest in recent years. Questions about them are related to questions about Markov chains, quantum automata, probabilistic automata, linear recurrent sequences, and probabilistic model checking problems. Studying them is of practical value because they occur frequently in practice. A tool called *Terminator*, developed at Microsoft Research, Cambridge, uses theoretical results about the verification of linear programs to automatically prove termination and liveness properties of Windows device drivers.

This reduced computational model allows a variety of properties to be decided. Algorithms for such problems typically rely on polynomial-time reductions and various mathematical fields such as computational linear algebra, algebraic number theory and Galois theory. We will now sketch some examples of recent work on such problems.

Tiwari [11] showed that it is decidable whether the program

$$\textbf{while } (Bx > b) \textbf{ do } x := Ax + c \textbf{ end}$$

terminates on all initial values of x in $\mathbb{R}^n$. This work was later generalised by Braverman [12], who showed that termination is decidable over $\mathbb{Q}$ in the general case and over $\mathbb{Z}$ in the homogeneous case $c = 0$.

The form of the termination problem considered in [11, 12] has an implicit universal quantification over a large set of possible initial states. Alternatively, one can ask about termination from a subset of the state space, such as a lower-dimensional subspace, a semi-linear set, or even a single point. The latter is related to *Skolem's problem*: determine whether there exists a positive integer n such that $u^T A^n v = 0$, for given vectors u, v and a square matrix A. Observe that Skolem's problem is equivalent to the question of whether the program

$$\textbf{while } (u^T x \neq 0) \textbf{ do } x := Ax \textbf{ end}$$

halts on initial value $x := v$. Another equivalent formulation of the problem, though slightly less obvious, is whether a given linear recurrent sequence (x_n) has a *zero*: an index n such that $x_n = 0$.

The solvability of Skolem's problem is open. Algorithms are known for it in the case of matrices of small dimension, or equivalently, recurrences of low depth. A paper by Vereshchagin [13] gives a solution for recurrences of depth at most 4. Recently, an attempt was made by Halava et al. [8] to show decidability for recurrences of depth 5. Ongoing efforts are focusing on finding algorithms to solve further special cases. The general problem is known to be NP-hard (Blondel and Portier, [14]), but since it is not at all clear whether it is decidable in general, there is a large gap between the known upper and lower complexity bounds.

A related problem is the *Positivity problem*, which asks whether a linear recurrence is always non-negative. It has been solved for recurrences of depth 2 by Halava et al. [15]. The question of its decidability in the general case is open, but a reduction from Skolem to Positivity is known.

Much effort has been devoted to partial decision procedures for termination through the synthesis of ranking functions. This approach consists of finding a function f from the state space of the program into some well-founded set. This function must have the property that if x is a valuation of the variables which satisfies the loop guard, and the loop body transforms it into x', then $f(x') < f(x)$. A complete method for finding *linear* ranking functions was described by Podelski and Rybalchenko [16].

1.2 Overview of this project

This project is concerned with the *Orbit problem:*

> Given vectors $x, y \in \mathbb{Q}^n$ and a matrix $A \in \mathbb{Q}^{n \times n}$,
> does there exist $m \in \mathbb{N}$ such that $A^m x = y$?

This is equivalent to asking whether the linear program

$$\textbf{while } (x \neq y) \textbf{ do } x := Ax \textbf{ end}$$

halts. The problem was shown to be decidable, and in fact decidable in polynomial time, in a seminal paper of Kannan and Lipton [1]. In their conclusion, the authors discuss a related problem, the *generalised Orbit problem*:

> Given a vector $x \in \mathbb{Q}^n$, a matrix $A \in \mathbb{Q}^{n \times n}$ and a vector subspace S specified by a basis, does there exist $m \in \mathbb{N}$ such that $A^m x \in S$?

The Orbit problem concerns reachability from point to point in a linear system, whereas the generalised Orbit problem replaces the target point with a vector space. Thus, in a sense, the former is a zero-dimensional version of the latter. The authors conjecture that the generalised Orbit problem is decidable, and that its one-dimensional version is decidable in polynomial time. They suggest that future work should focus on showing decidability, and possibly low complexity, for versions of the problem of small dimension, in order to gain insight into the general case. Since the paper was published in 1986, no progress has been recorded in the literature regarding the decidability of these problems.

The contribution of this dissertation is to prove that the one-dimensional Orbit problem is decidable in polynomial time, as Kannan and Lipton conjectured. Additionally, we give a simpler proof of the polynomial-time decidability of the zero-dimensional version. Our proof eliminates the need to reason about the cumbersome Jordan canonical form and is much less involved.

Chapter 2 provides the necessary mathematical background and proves the results needed later. Chapter 3 investigates the zero-dimensional Orbit problem and proves that it is polynomial-time decidable. Chapter 4 shows the same for the one-dimensional version. Finally, Chapter 5 concludes this dissertation and gives suggestions for future work.

Chapter 2

Mathematical Foundations

A subset of the complex numbers which is of special mathematical interest is the set of algebraic numbers, that is, the roots of polynomials with rational coefficients. They frequently arise in matrix and polynomial problems over $\mathbb{Q}$. A problem concerning a rational matrix may often be reduced to questions regarding its eigenvalues, which arise as the roots of its characteristic polynomial.

In our solution to the one-dimensional Orbit problem, we will exploit such a connection between matrices and their eigenvalues. Sections 2.1 to 2.4 contain the textbook results which are necessary for our proof. The main theme is to show that operations on algebraic numbers can be carried out efficiently. For a broader view of algebraic number theory, we suggest Cohen [4] and Stewart and Tall [5].

Sections 2.5 and 2.6 discuss a decision problem on algebraic numbers and two of its special cases. The general problem is known to be decidable. Kannan and Lipton proved a polynomial-time bound for the first special case and used it to prove the zero-dimensional Orbit problem decidable in polynomial time. We will prove a polynomial-time bound for the second special case, and use it in our result for the one-dimensional Orbit problem. This complexity analysis is the subject of Section 2.6 and makes use of bounds given in Section 2.5.

2.1 Basic definitions and properties

We say that $\alpha \in \mathbb{C}$ is an *algebraic number* just if there exists a non-zero polynomial $p \in \mathbb{Q}[x]$ such that $p(\alpha) = 0$. Further, if p may be chosen to be in $\mathbb{Z}[x]$ and monic, then α is an *algebraic integer*. We define the *minimal*

polynomial of α, denoted $f_\alpha(x)$, to be the monic polynomial in $\mathbb{Q}[x]$ of least degree which vanishes at α. The *degree* and *height* of α are defined as the degree and height[1] of f_α.

Proposition 1. *The minimal polynomial $f_\alpha(x)$ of α is unique and irreducible over $\mathbb{Q}$. Further, any polynomial which vanishes at α must be a multiple of $f_\alpha(x)$.*

Proof. If $f(x)$ and $g(x)$ are minimal polynomials of α, then $h(x) = f(x) - g(x)$ is of lower degree and vanishes at α, so $h(x)$ is the zero polynomial. This proves uniqueness.

If $f_\alpha(x) = u(x)\,v(x)$ for some $u, v \in \mathbb{Q}[x]$ of non-zero degree, then $u(\alpha) = 0$ or $v(\alpha) = 0$ must hold. But the degrees of u, v are strictly smaller than the degree of f_α, which contradicts the minimality of f_α. This proves irreducibility.

Suppose $p(\alpha) = 0$ for some $p \in \mathbb{Q}[x]$. If u and v are the quotient and the remainder of the polynomial division of $p(x)$ by $f_\alpha(x)$, then

$$0 = p(\alpha) = u(\alpha)\,f_\alpha(\alpha) + v(\alpha) = v(\alpha)$$

But $\deg(v) < \deg(f_\alpha)$, so v must be identically zero. Therefore, f_α divides p.

■

The roots of $f_\alpha(x)$ (including α) are called the *Galois conjugates* of α. By Proposition 1, if α satisfies some polynomial equation $p(x) = 0$, then f_α divides p, so the Galois conjugates of α must also satisfy the equation. Therefore, we say that Galois conjugates are *algebraically indistinguishable* from each other.

Proposition 2. *The polynomial f_α has no repeated roots.*

Proof. Suppose α is a repeated root of $f_\alpha(x)$. Then $f_\alpha(x) = (x-\alpha)^2\,u(x)$ for some $u \in \mathbb{C}[x]$. The derivative of f_α is

$$f'_\alpha(x) = (x-\alpha)\,(2u(x) + (x-\alpha)\,u'(x))$$

Therefore, $f'_\alpha(\alpha) = 0$, which contradicts the minimality of f_α because $f'_\alpha(x) \in \mathbb{Q}[x]$ and $\deg(f'_\alpha) < \deg(f_\alpha)$.

[1] The height of a polynomial is the maximum absolute value of its coefficients.

■

An algebraic number α is an *n-th root of unity* just if it is a root of $x^n - 1$. The least $n \in \mathbb{N}^+$ such that α is an n-th root of unity is called the *order* of α. An n-th root of unity is called *primitive* just if its order is n. The n-th roots of unity are exactly $\exp(2\pi it/n)$, for $t \in \{1, \ldots, n\}$. Moreover, $\exp(2\pi it/n)$ has order n if and only if $\gcd(t, n) = 1$. If α is an n-th root of unity, then so are all integer powers of α. Further, if α is primitive, then $\alpha^1, \ldots, \alpha^n$ are exactly the n-th roots of unity in some order. The *n-th cyclotomic polynomial* is defined as

$$C_n(x) = \prod_{\substack{1 \leq t \leq n \\ \gcd(t,n)=1}} (x - \exp(2\pi it/n))$$

It is easy to prove by induction that $C_n \in \mathbb{Z}[x]$ and that it is minimal for the primitive n-th roots of unity. Its degree is $\varphi(n)$, where φ is Euler's totient function:

$$\varphi(n) = p_1^{k_1 - 1} p_2^{k_2 - 1} \ldots p_s^{k_s - 1} (p_1 - 1)(p_2 - 1) \ldots (p_s - 1)$$

where $p_1^{k_1} \ldots p_s^{k_s}$ is the decomposition of n into primes. It satisfies the inequality

$$\sqrt{n} \leq \varphi(n) \leq n - 1$$

Therefore, if α is algebraic with degree d, either α is not a root of unity, or it is a root of unity of order at most d^2.

Given monic polynomials $f, g \in \mathbb{Q}[x]$ with roots $\{\alpha_i\}, \{\beta_j\}$, their *resultant* is defined as

$$\mathcal{R}(f, g) = \prod_{i=1}^{\deg(f)} \prod_{j=1}^{\deg(g)} (\alpha_i - \beta_j)$$

Viète's Laws may be used to show that $\mathcal{R}(f, g) \in \mathbb{Q}$. The resultant is computable in polynomial time, see section 3.3 of [4].

Proposition 3. *The algebraic numbers form a field under addition and multiplication.*

Proof. It suffices to show closure under addition, subtraction, multiplication and division.

Let α, β be algebraic numbers and $g(x)$ be the resultant of $f_\alpha(x-y)$ and $f_\beta(y)$, interpreted as polynomials in y. One of the roots of $f_\alpha(x-y)$ is $y = x - \alpha$. Therefore, $\alpha + \beta$ is a root of $g(x)$. Similarly,

1. $\alpha - \beta$ is a root of $\mathcal{R}(f_\alpha(x+y), f_\beta(y))$,
2. $\alpha\beta$ is a root of $\mathcal{R}\left(y^{\deg(\alpha)} f_\alpha(x/y), f_\beta(y)\right)$,
3. α/β is a root of $\mathcal{R}(f_\alpha(xy), f_\beta(y))$,

■

The *absolute norm* of an algebraic number α, denoted $\mathcal{N}(\alpha)$, is the product of its Galois conjugates, including itself:

$$\mathcal{N}(\alpha) = \prod_{\beta\,:\,f_\alpha(\beta)=0} \beta$$

By Viète's Laws,

$$\mathcal{N}(\alpha) = (-1)^{\deg(\alpha)} a_0$$

where a_0 is the constant term of f_α.

The notion of minimal polynomial extends to matrices. For $A \in \mathbb{Q}^{n\times n}$, the *minimal polynomial* of A, denoted $f_A(x)$, is the monic polynomial in $\mathbb{Q}[x]$ of least degree such that $f_A(A)$ is the zero matrix. If $p \in \mathbb{Q}[x]$, $p(A) = 0$ if and only if f_A divides p. A well-known result is the following:

Theorem. *(Cayley-Hamilton) The minimal polynomial f_A of A divides the characteristic polynomial of A.*

Thus, the minimal polynomial is of degree at most the size of the matrix, and its roots coincide with the eigenvalues of A, up to multiplicity.

2.2 Algebraic number fields

Let K and L be fields, and K be a subfield of L. Then L is called a *field extension* of K, denoted $L : K$. In this case, L may be seen as a vector space over K, where vector addition is addition in L, and scalar multiplication takes $\lambda \in K$, $v \in L$ and returns $\lambda v \in L$. The dimension of this vector space is called the *degree* or *dimension* of the field extension, denoted $[L : K]$.

A *number field* is a field extension L of $\mathbb{Q}$ such that $[L : \mathbb{Q}]$ is finite. For $\alpha_1, \ldots, \alpha_n \in \mathbb{C}$, define $\mathbb{Q}(\alpha_1, \ldots, \alpha_n)$ to be the smallest field extension of $\mathbb{Q}$ containing $\alpha_1, \ldots, \alpha_n$.

Proposition 4. *If α is algebraic with degree d, then $[\mathbb{Q}(\alpha) : \mathbb{Q}] = d$ and*

$$\mathcal{B} = \left\{\alpha^0, \ldots, \alpha^{d-1}\right\}$$

is a basis for $\mathbb{Q}(\alpha)$.

Proof. If $c_0\alpha^0 + \cdots + c_{d-1}\alpha^{d-1} = 0$ for some coefficients $c_i \in \mathbb{Q}$, then the polynomial $c_0x^0 + \cdots + c_{d-1}x^{d-1}$ vanishes at α and is of smaller degree than f_α. Therefore, $c_0 = \cdots = c_{d-1} = 0$, proving linear independence.

Moreover, α^d is linearly dependent on $\mathcal{B}$, evidenced by $f_\alpha(\alpha) = 0$. Therefore, for all i, $\alpha^i \in span(\mathcal{B})$, so $span(\mathcal{B}) = \mathbb{Q}(\alpha)$. This proves that $\mathcal{B}$ is a basis and $[\mathbb{Q}(\alpha) : \mathbb{Q}] = |\mathcal{B}| = d$.

■

In general, if K is a number field, then $K = \mathbb{Q}(\alpha_1, \ldots, \alpha_n)$ for finitely many algebraic numbers $\alpha_1, \ldots, \alpha_n$. In fact, a stronger property holds:

Theorem. *(Primitive Element) If K is a number field, then $K = \mathbb{Q}(\theta)$ for some algebraic number θ.*

Proof. See Theorem 2.2 of [5].

■

2.3 Representation of algebraic numbers

A canonical way of specifying an algebraic number α is to refer to it by its minimal polynomial f_α. We include a numerical approximation of $Re(\alpha)$ and

$Im(\alpha)$ of sufficient accuracy to distinguish α from its Galois conjugates. More precisely, we represent α as a tuple

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

with the meaning 'α is the unique root of f_α which is inside the circle centred at (x, y) in the complex plane with radius R'. In order to make this well-defined, we invoke a *root separation* bound. One such bound, due to Mignotte [6], states that for roots $\alpha_i \neq \alpha_j$ of a polynomial $p(x)$,

$$|\alpha_i - \alpha_j| > \frac{\sqrt{6}}{n^{\frac{n+1}{2}} H^{n-1}}$$

where $n = \deg(p)$ and H is the height of p. If we restrict R to be less than half of the root separation bound for f_α, the disk (x, y, R) is guaranteed to include at most one root of f_α, making the representation unambiguous. However, we will restrict R to be less than *a quarter* of the root separation bound, in order to allow equality checking.

Observe that if we adopt the convention that $R = 2^{-t}$ for some $t \in \mathbb{N}$, then this requirement becomes

$$t > 2 + \frac{n+1}{2} \log_2 n + (n-1) \log_2 H - \log_2 \sqrt{6}$$

which is bounded by a polynomial in n and $\log_2 H$. Therefore, given f_α, the number of bits required to describe α is polynomial in the size of the input. Pan [7] gives an algorithm to obtain polynomially many bits of the roots of a given polynomial in polynomial time. Thus, the crux of identifying an algebraic number will be finding its minimal polynomial. The rest is just numerical approximation, which we assume we can always do.

Proposition 5. *Given canonical representations $\alpha = (f_\alpha, x_\alpha, y_\alpha, R_\alpha)$ and $\beta = (f_\beta, x_\beta, y_\beta, R_\beta)$, canonical representations of $\alpha + \beta$, $\alpha - \beta$, $\alpha\beta$ and α/β are computable in polynomial time. Moreover, it is decidable in polynomial time whether $\alpha \in \mathbb{N}$, $\alpha \in \mathbb{Z}$, $\alpha \in \mathbb{Q}$, $\alpha = \beta$.*

Proof. Clearly, $\alpha \in \mathbb{N}$ iff $f_\alpha(x) = x - t$ for some $t \in \mathbb{N}$, which can be decided by inspection. Similarly for $\alpha \in \mathbb{Z}$ and $\alpha \in \mathbb{Q}$.

To decide $\alpha = \beta$, first examine f_α and f_β. If they are distinct, then clearly $\alpha \neq \beta$. Otherwise, determine if the disks $(x_\alpha, y_\alpha, R_\alpha)$, $(x_\beta, y_\beta, R_\beta)$ have a common point by calculating

$$d^2 = |(x_\alpha + iy_\alpha) - (x_\beta + iy_\beta)|^2$$

and checking $d^2 \leq (R_\alpha + R_\beta)^2$. If not, then $\alpha \neq \beta$. Otherwise, observe that the distance from any point in one disk to any point in the other is at most $2R_\alpha + 2R_\beta$, which is less than the root separation bound. Therefore, $\alpha = \beta$.

For the arithmetic operations, we use the resultant method. We compute the relevant resultant (see Proposition 3), factor it into irreducible polynomials [3] and use the numerical approximations of α and β to determine which factor is the minimal polynomial of the resulting algebraic number.

■

Proposition 6. *Given a canonical description $\alpha = (f_\alpha, x_\alpha, y_\alpha, R_\alpha)$, it is decidable in polynomial time whether α is a root of unity. Further, if it is, its order n is computable in polynomial time, along with the integer $t \in \{1, \ldots, n\}$ such that $\alpha = \exp(2\pi it/n)$.*

Proof. Let $d = \deg(\alpha)$. It suffices to check whether f_α divides some polynomial $x^j - 1$ for $j = 1, \ldots, d^2$. If not, then α is not a root of unity. Otherwise, n is the least j such that f_α divides $x^j - 1$. Then we use numerical approximation to determine which n-th root of unity is equal to α.

■

When the intermediate results of an algorithm are confined to some number field $\mathbb{Q}(\gamma)$, it is convenient to represent algebraic numbers by their coordinates with respect to a basis for $\mathbb{Q}(\gamma)$. If $d = \deg(\gamma)$, the tuple

$$(a_0, a_1, \ldots, a_{d-1}) \in \mathbb{Q}^d$$

is called the *standard representation of $\sum_{i=0}^{d-1} a_i \gamma^i$ with respect to $\mathbb{Q}(\gamma)$*. By linear independence, each algebraic number in $\mathbb{Q}(\gamma)$ has a unique standard representation.

Proposition 7. *Given a canonical representation of γ and standard representations of $\alpha = (a_0, \ldots, a_{d-1})$ and $\beta = (b_0, \ldots, b_{d-1})$ with respect to $\mathbb{Q}(\gamma)$, the standard representations of $\alpha + \beta$, $\alpha - \beta$, $\alpha\beta$ and α/β are computable in polynomial time. Moreover, it is decidable in polynomial time whether $\alpha \in \mathbb{N}$, $\alpha \in \mathbb{Z}$, $\alpha \in \mathbb{Q}$, $\alpha = \beta$.*

Proof. The decidability questions are immediate from the uniqueness of the representation: $\alpha = \beta$ iff their tuples are equal, $\alpha \in \mathbb{Q}$ iff there are only zeroes in components $1, \ldots, d-1$ of the description of α, and so on. Addition and subtraction are also trivial.

For multiplication, we inductively precompute the standard representations of $\gamma^d, \gamma^{d+1}, \ldots, \gamma^{2d-2}$. In the base case, we obtain the coordinates of $\gamma^d = (u_0, u_1, \ldots, u_{d-1})$ directly from the identity $f_\gamma(\gamma) = 0$. If for some i we have $\gamma^i = (t_0, \ldots, t_{d-1})$, then

$$\gamma^{i+1} = \left(\sum_{j=0}^{d-1} t_j \gamma^j \right) \gamma = t_{d-1}\gamma^d + \sum_{j=1}^{d-1} t_{j-1}\gamma^j$$

so the coordinates of γ^{i+1} are:

$$t_{d-1}(u_0, u_1, \ldots, u_{d-1}) + (0, t_0, \ldots, t_{d-2})$$

Having done this precomputation, we expand the product

$$\alpha\beta = \left(\sum_{j=0}^{d-1} a_j \gamma^j \right) \left(\sum_{j=0}^{d-1} b_j \gamma^j \right)$$

and substitute the coordinates of $\gamma^0, \ldots, \gamma^{2d-2}$.

To divide, we use the extended Euclidean algorithm on the polynomial $b(x)$ with coefficients $b_0, \ldots, b_{d-1}$ and f_γ. As b and f_γ are coprime, this gives $u, v \in \mathbb{Q}[x]$ such that

$$u(x)\, b(x) + v(x)\, f_\gamma(x) = 1$$

Then we have

$$\frac{\alpha}{\beta} = \frac{a(\gamma)}{b(\gamma)} = a(\gamma)u(\gamma)$$

whose coordinates may be calculated by multiplication.

■

2.4 Monomorphisms

A *monomorphism* from field a K to a field L is a function φ which preserves addition, multiplication and the units:

$$\begin{array}{ll} \varphi(x+y) = \varphi(x) + \varphi(y) & \text{for all } x, y \in K \\ \varphi(xy) = \varphi(x)\varphi(y) & \text{for all } x, y \in K \\ \varphi(1_K) = 1_L & \\ \varphi(0_K) = 0_L & \end{array}$$

A well-known result is the following:

Theorem. *If $K = \mathbb{Q}(\alpha)$ is a number field of degree n, then there exist exactly n distinct monomorphisms $\sigma_1, \ldots, \sigma_n$ from K into $\mathbb{C}$, defined by $\sigma_i(\alpha) = \alpha_i$, where $\alpha_1, \ldots, \alpha_n$ are the Galois conjugates of α.*

For example, take $K = \mathbb{Q}\left(i\sqrt{3}\right)$. There are exactly two monomorphisms from K to $\mathbb{C}$:

$$\sigma_1(\alpha) = i\sqrt{3}$$

$$\sigma_2(\alpha) = -i\sqrt{3}$$

For a proof of the theorem, see [5], page 38. Another well-known result is the Monomorphism Extension Theorem:

Theorem. *If K, L and H are fields, where $L : K$, and there exists a monomorphism σ from K to H, then there exists a monomorphism φ from L to H which agrees with σ.*

A good way of thinking about monomorphisms is as functions permuting the Galois conjugates. More precisely, we have the following:

Proposition 8. *If α is algebraic and σ is a monomorphism of $\mathbb{Q}(\alpha)$ into $\mathbb{C}$, then $\sigma(\alpha)$ is a Galois conjugate of α.*

Proof. We have $\sigma(p(y)) = p(\sigma(y))$ for all polynomials $p \in \mathbb{Q}[x]$ and all $y \in \mathbb{Q}(\alpha)$. Hence,

$$f_\alpha(\sigma(\alpha)) = \sigma(f_\alpha(\alpha)) = \sigma(0) = 0$$

which shows that $\sigma(\alpha)$ is a Galois conjugate of α.

■

This allows us to determine the Galois conjugates of a polynomial applied to an algebraic number:

Proposition 9. *If $\alpha = \alpha_1$ is algebraic, with Galois conjugates $\alpha_1, \ldots, \alpha_k$ and $q \in \mathbb{Q}[x]$, the Galois conjugates of $q(\alpha)$ are the distinct elements of the sequence $q(\alpha_1), \ldots, q(\alpha_k)$.*

Proof. Consider some α_i and its associated monomorphism σ_i from $\mathbb{Q}(\alpha)$ to $\mathbb{C}$. Then

$$q(\alpha_i) = q(\sigma_i(\alpha)) = \sigma_i(q(\alpha))$$

But by the previous proposition, this must be a Galois conjugate of $q(\alpha)$.

Now suppose θ is some Galois conjugate of $q(\alpha)$. There must be a monomorphism σ from $\mathbb{Q}(q(\alpha))$ to $\mathbb{C}$ which maps $q(\alpha)$ to θ. Moreover, $\mathbb{Q}(q(\alpha))$ is a subfield of $\mathbb{Q}(\alpha)$, so by the Monomorphism Extension Theorem, there must be a monomorphism φ from $\mathbb{Q}(\alpha)$ into $\mathbb{C}$ which agrees with σ. Therefore,

$$\theta = \sigma(q(\alpha)) = \varphi(q(\alpha)) = q(\varphi(\alpha))$$

But there are only k monomorphisms from $\mathbb{Q}(\alpha)$ to $\mathbb{C}$, one for each α_i, so φ must be among them and $\theta = q(\alpha_i)$ for some i.

■

Note that the sequence $q(\alpha_i)$ may contain duplicates. For example, take $q(x) = x^2$ and $\alpha = \alpha_1 = i$, $\alpha_2 = -i$.

2.5 Magnitude bounds

Now we focus on the two main results of this chapter. In this section, we quote bounds on magnitudes of algebraic numbers. In particular, we are concerned with upper and lower bounds on the magnitude of a polynomial evaluated at an algebraic number, in terms of the degrees and heights of the polynomial and the algebraic number.

While we have stated the bounds precisely for the sake of completeness, their exact detail is unimportant. Notice, however, that in all of them, heights appear only in the base position, whereas degrees appear both in the base and the exponent. Thus when we take logarithms, the resulting expressions will be bounded by polynomials in the degrees and the logarithms of heights.

Let n_t and H_t denote the degree and the height of t, where t can be a polynomial or an algebraic number.

Proposition 10. *If $p \in \mathbb{C}[x]$ and $p(\alpha) = 0$, then*

$$\frac{|u|}{H_p + |u|} < |\alpha| < \frac{H_p + |v|}{|v|}$$

where u and v are respectively the constant term and the leading coefficient of p. In particular, if $p(x) \in \mathbb{Q}[x]$ is the minimal polynomial of α, then $|\alpha| < H_p + 1$.

Proof. See page 30 of Shidlovskii [17].

■

Proposition 11. *(Blanksby and Montgomery) If α is an algebraic integer and not a root of unity, then there exists some Galois conjugate θ of α such that*

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

Proof. See [9].

■

It is easy to show that the Blanksby and Montgomery bound implies

$$\frac{1}{\log_2 |\theta|} \leq 60 n_\alpha^2 \ln(6 n_\alpha)$$

Proposition 12. *If α is an algebraic number, and $p \in \mathbb{Q}[x]$ such that $p(\alpha) \neq 0$, then*

$$|p(\alpha)| \geq \frac{1}{(3^{n_\alpha - 1} H_\alpha^{n_\alpha})^{n_p} H_p^{n_\alpha - 1}}$$

Proof. See page 31 of Shidlovskii [17].

■

Proposition 13. *If α is an algebraic number and $p \in \mathbb{Q}[x]$, then*

$$|\alpha| > 1 \Rightarrow |p(\alpha)| \leq (n_p + 1) H_p (H_\alpha + 1)^{n_p}$$

and

$$|\alpha| \leq 1 \Rightarrow |p(\alpha)| \leq (n_p + 1) H_p$$

Proof.

$$\begin{array}{rcll} |p(\alpha)| & = & \left|\sum_{i=0}^{n_p} p_i \alpha^i\right| & \\ & & & \text{(triangle inequality)} \\ & \leq & \sum_{i=0}^{n_p} |p_i| \, |\alpha|^i & \\ & & & (|p_i| \leq H_p) \\ & \leq & H_p \sum_{i=0}^{n_p} |\alpha|^i & \end{array}$$

If $|\alpha| > 1$, then the largest term in the summation is $|\alpha|^{n_p}$, which is bounded above by $(H_\alpha + 1)^{n_p}$ according to Proposition 10. If $|\alpha| \leq 1$, then the largest term is $|\alpha|^0 = 1$. This gives the desired inequalities.

■

2.6 Algebraic Number Power problem

We close this chapter by examining the following decision problem, which is related to the Orbit problem:

ALGEBRAIC NUMBER POWER
Given algebraic numbers α, β,
does there exist $m \in \mathbb{N}$ such that $\alpha^m = \beta$?

Reference [8] shows that ALGEBRAIC NUMBER POWER is decidable. We give a brief recapitulation of the proof. Then we consider two special cases of the problem and show that the bounds used for decidability in the general case are polynomial in the special cases.

Proposition 14. ALGEBRAIC NUMBER POWER *is decidable.*

Proof (Sketch). We consider three cases:

1. α is a root of unity.
2. α is an algebraic integer, but not a root of unity.
3. α is not an algebraic integer.

Case 1. If α is a primitive k-th root of unity, then the set of witnesses $W = \{m \mid \alpha^m = \beta\}$ is either empty, or one of the equivalence classes of $\mathbb{N}$ modulo k, depending on whether β is a k-th root of unity.

Case 2. If α is an algebraic integer but not a root of unity, α must have a Galois conjugate α_i such that $|\alpha_i| > 1$, by Blanksby and Montgomery's theorem. The monomorphism σ_i from $\mathbb{Q}(\alpha)$ into $\mathbb{C}$ defined by $\sigma_i(\alpha) = \alpha_i$ must extend to a monomorphism φ from $\mathbb{C}$ to $\mathbb{C}$ by the Monomorphism Extension Theorem. Applying it to $\alpha^m = \beta$ gives

$$\varphi(\alpha^m) = \alpha_i^m = \varphi(\beta)$$

Since $|\alpha_i| > 1$, m is bounded by

$$m \leq \frac{\log_2 |\varphi(\beta)|}{\log_2 |\alpha_i|}$$

Case 3. If α is not an algebraic integer, then a bound on m is obtained from ideal theory. A consequence of α not being an algebraic integer is that there exists a *prime ideal* P and an associated function v_P (the *p-adic valuation*) from algebraic numbers into $\mathbb{N}$, such that

- $v_P(xy) = v_P(x) + v_P(y)$ for all x, y
- $2^{v_P(x)} \leq |\mathcal{N}(x)|$ for all x
- $v_P(\alpha) \neq 0$

This is the only place in this report where we use ideals and p-adic valuation, so we take the existence of v_P for granted. For details, see reference [8].

Thus if $\alpha^m = \beta$, we have

$$v_P(\beta) = v_P(\alpha^m) = m v_P(\alpha)$$

which gives the bound

$$\log_2 |\mathcal{N}(\beta)| \geq v_P(\beta) = m v_P(\alpha) \geq m$$

In Case 2 and Case 3, we can decide $\exists m.\alpha^m = \beta$ by computing the powers of α up to the appropriate bound and comparing with β at each step.

■

Of particular interest for the Orbit problem are the following cases:

1. When $\beta = q(\alpha)$ for some $q \in \mathbb{Q}[x]$, which arises in Kannan and Lipton's argument for the zero-dimensional case.
2. When $\alpha = \alpha_i/\alpha_j$, $\beta = q(\alpha_i)/q(\alpha_j)$ for some algebraic numbers α_i, α_j and $q \in \mathbb{Q}[x]$, which arises in our solution for the one-dimensional case.

We will show that the bounds on m are polynomial in the length of the input.

Proposition 15. *Suppose we are given a canonical description of an algebraic α which is not a root of unity and a polynomial $q \in \mathbb{Q}[x]$. If $\alpha^m = q(\alpha)$ for some $m \in \mathbb{N}$, then m is bounded by a polynomial in n_α, n_q, $\log_2 H_\alpha$ and $\log_2 H_q$.*

Proof. If α is an algebraic integer, then we choose its Galois conjugate θ which satisfies the Blanksby and Montgomery bound. By algebraic indistinguishability, we have

$$\alpha^m = q(\alpha) \iff \theta^m = q(\theta)$$

Hence, the bound on m is

$$m \leq \frac{\log_2 |q(\theta)|}{\log_2 |\theta|}$$

Applying the upper bound on $|q(\theta)|$ from Proposition 13 and the Blanksby and Montgomery bound on $1/\log_2 |\theta|$ gives a polynomially-bounded expression.

If α is not an algebraic integer, the bound on m is

$$m \leq \log_2 |\mathcal{N}(q(\alpha))|$$

By Proposition 9, the Galois conjugates of $q(\alpha)$ are of the form $q(\alpha_i)$, where α_i are the Galois conjugates of α. Then $|\mathcal{N}(q(\alpha))|$ is a product of terms $|q(\alpha_i)|$. Each term can be bounded above by Proposition 13, and there are at most n_α such terms. Taking logarithms gives a polynomial bound.

■

Proposition 16. *Suppose we are given a canonical description of algebraic α, β such that α/β is not a root of unity, and a polynomial $q \in \mathbb{Q}[x]$. If*

$$\left(\frac{\alpha}{\beta}\right)^m = \frac{q(\alpha)}{q(\beta)}$$

for some $m \in \mathbb{N}$, then m is bounded by a polynomial in n_α, n_β, n_q, $\log_2 H_\alpha$, $\log_2 H_\beta$ and $\log_2 H_q$.

Proof.

If $\gamma = \alpha/\beta$ is an algebraic integer, then choose its Galois conjugate γ_k which satisfies the Blanksby and Montgomery bound. If φ is the extension to $\mathbb{C}$ of the monomorphism from $\mathbb{Q}(\gamma)$ to $\mathbb{C}$ which maps γ to γ_k, we need a polynomial upper bound on

$$m \leq \frac{\log_2 \left|\varphi\left(\frac{q(\alpha)}{q(\beta)}\right)\right|}{\log_2 |\gamma_k|} = \frac{\log_2 |q(\varphi(\alpha))| - \log_2 |q(\varphi(\beta))|}{\log_2 |\gamma_k|}$$

By Proposition 8, $\varphi(\alpha)$ is some Galois conjugate of α, and $\varphi(\beta)$ is some Galois conjugate of β. Propositions 12 and 13 give a polynomial bound on $\log_2 |q(\varphi(\alpha))| - \log_2 |q(\varphi(\beta))|$. Proposition 11 gives a polynomial bound on $1/\log_2 |\gamma_k|$ in terms of n_γ, which is at most $n_\alpha n_\beta$.

If γ is not an algebraic integer, then we need a polynomial upper bound on

$$m \leq \log_2 \left|\mathcal{N}\left(\frac{q(\alpha)}{q(\beta)}\right)\right| = \log_2 |\mathcal{N}(q(\alpha))| - \log_2 |\mathcal{N}(q(\beta))|$$

The Galois conjugates of $q(\alpha)$ are of the form $q(\alpha_i)$ where α_i are the Galois conjugates of α. There are at most n_α of them, so we can use Proposition 13 to obtain a polynomial upper bound on $\log_2 |\mathcal{N}(q(\alpha))|$. Similarly for $-\log_2 |\mathcal{N}(q(\beta))|$ using Proposition 12.

■

Chapter 3

The Zero-dimensional Orbit Problem

The zero-dimensional Orbit problem was originally proven decidable by Kannan and Lipton [1] in a two-step proof. First, the problem was reduced in polynomial time to a form of the Matrix Power problem. Then a connection between matrices and their eigenvalues was exploited to obtain a polynomial-time algorithm.

A central idea of the paper is to circumvent the difficulties associated with matrix problems and replace them with questions regarding algebraic numbers. However, near the end, the authors encounter a problematic case, and the argument reverts back to matrices. It makes use of the cumbersome Jordan canonical form, which detracts from the elegance of the original idea. Additionally, the reduction argument from the Orbit problem to the Matrix Power problem is set in the usual basis of unit vectors and overlooks a more convenient basis for the vector space of interest.

In this chapter, we give a simpler proof of Kannan and Lipton's result. We present our polynomial-time reduction in Section 3.1 and our solution to the reduced problem in Section 3.2.

3.1 Reduction

We begin with a polynomial-time reduction from Orbit to Matrix Power. The two problems are defined as follows:

ORBIT
Given vectors $x, y \in \mathbb{Q}^n$ and a matrix $A \in \mathbb{Q}^{n \times n}$, does there exist $m \in \mathbb{N}$ such that $A^m x = y$?

MATRIX POWER
Given matrices $A, D \in \mathbb{Q}^{n \times n}$, does there exist $m \in \mathbb{N}$ such that $A^m = D$?

Suppose we have an instance (A, x, y) of ORBIT. Define

$$\nu = \max \{m \mid x, Ax, \ldots, A^m x \text{ are linearly independent}\}$$

We compute ν in polynomial time using Gaussian elimination to check linear independence. Let $\mathcal{B} = \{x, Ax, \ldots, A^\nu x\}$ and $S = span(\mathcal{B})$. It is clear that $A^m x \in S$ for all m.

We use Gaussian elimination to check $y \in S$. If $y \notin S$, then (A, x, y) is a negative instance. Suppose otherwise. We will switch from the usual basis of unit vectors to $\mathcal{B}$.

We compute the coordinates $y_\mathcal{B} = [a_0, \ldots, a_\nu]^T$ of y and the coordinates $[b_0, \ldots, b_\nu]^T$ of $A^{\nu+1}x$ with respect to $\mathcal{B}$. Note that the coordinates of x are $x_\mathcal{B} = [1, 0, \ldots, 0]^T$.

Now suppose we have a point p with coordinates $[p_0, \ldots, p_\nu]^T$ with respect to $\mathcal{B}$. Then for Ap, we have:

$$\begin{aligned}
& Ap \\
= \; & A \textstyle\sum_{i=0}^{\nu} p_i \left(A^i x\right) \\
= \; & \left(\textstyle\sum_{i=1}^{\nu} p_{i-1} \left(A^i x\right)\right) + p_\nu \left(A^{\nu+1} x\right) \\
= \; & \textstyle\sum_{i=1}^{\nu} \left(p_{i-1} + p_\nu b_i\right) \left(A^i x\right) + p_\nu b_0 \left(A^0 x\right) \\
= \; & \begin{bmatrix} x & Ax & \ldots & A^\nu x \end{bmatrix} M \begin{bmatrix} p_0 & p_1 & \ldots & p_\nu \end{bmatrix}^T
\end{aligned}$$

where

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

Therefore, M describes the linear transformation of premultiplying by A with respect to $\mathcal{B}$. Thus,

$$A^m x - y = \begin{bmatrix} x & Ax & \dots & A^\nu x \end{bmatrix} (M^m x_\mathcal{B} - y_\mathcal{B})$$

which gives

$$A^m x = y \iff M^m x_\mathcal{B} = y_\mathcal{B}$$

Note that $x_\mathcal{B}, \dots, M^\nu x_\mathcal{B}$ are exactly the unit vectors of dimension $\nu + 1$, so

$$\begin{gathered} M^m x_\mathcal{B} = y_\mathcal{B} \\ \iff \\ M^m \begin{bmatrix} x_\mathcal{B} & Mx_\mathcal{B} & \dots & M^\nu x_\mathcal{B} \end{bmatrix} = \begin{bmatrix} y_\mathcal{B} & My_\mathcal{B} & \dots & M^\nu y_\mathcal{B} \end{bmatrix} \\ \iff \\ M^m = \begin{bmatrix} y_\mathcal{B} & My_\mathcal{B} & \dots & M^\nu y_\mathcal{B} \end{bmatrix} \end{gathered}$$

which gives an equivalent instance of Matrix Power.

We proceed with a reduction to the following problem:

> Matrix Power (Polynomial Version)
> Given a matrix $A \in \mathbb{Q}^{n \times n}$ and a polynomial $q \in \mathbb{Q}[x]$, does there exist $m \in \mathbb{N}$ such that $A^m = q(A)$?

Suppose we have an instance (A, D) of Matrix Power. We define $q_0, \dots, q_{n-1}$ to be unknowns ranging over $\mathbb{Q}$, and solve the linear system of n^2 equations $\sum_{i=0}^{n-1} q_i A^i = D$. If it has no solution, then we conclude that (A, D) is a negative instance of Matrix Power and terminate. Otherwise, we choose a solution $q_0, \dots, q_{n-1}$ and output (A, q) as the result of the reduction.

Proposition 17. *If (A, D) is a positive instance of* Matrix Power, *then the system $\sum_{i=0}^{n-1} q_i A^i = D$ has a solution.*

Proof. Suppose $A^m = D$ and let

$$x^m = f_A(x)\, u(x) + r(x)$$

where u and r are respectively the quotient and the remainder of the polynomial division of x^m by $f_A(x)$. Then

$$D = A^m = f_A(A)\,u(A) + r(A) = r(A)$$

Since $\deg(r) < \deg(f_A) \leq n$, r is a solution of the linear system.

■

Thus, if (A, D) is a positive instance of MATRIX POWER, the linear system is guaranteed to have a solution q. Any such solution yields a positive instance (A, q) of MATRIX POWER (POLYNOMIAL VERSION). If (A, D) is a negative instance and the system has no solution, we correctly determine that (A, D) is negative. Finally, if (A, D) is negative and the system has a solution q, then for all m, $A^m \neq D = q(A)$, so (A, q) is a negative instance of MATRIX POWER (POLYNOMIAL VERSION). This proves the correctness of the reduction.

3.2 Solution

Suppose we have an instance (A, q) of MATRIX POWER (POLYNOMIAL VERSION). To determine the existence of $m \in \mathbb{N}$ such that $A^m = q(A)$, we first calculate the minimal polynomial $f_A(x)$ of A. A conceptually simple procedure for computing f_A in polynomial time is to consider each possible degree $d \in \{1, \ldots, n\}$, determine if there is a polynomial of degree d which vanishes at A by solving the linear equation system $\sum_{i=0}^{d} a_i A^i = 0$ in $a_0, \ldots, a_d$, and take the least d for which such a polynomial is found. If needed, we scale the polynomial to make it monic. An efficient procedure which would be of more practical use is given in [2].

Having computed $f_A(x)$, we factor it [3] into irreducible polynomials in $\mathbb{Q}[x]$. Each factor is the minimal polynomial of its roots, so we can construct canonical representations of the roots of f_A in polynomial time. Let these roots be $\alpha_1, \ldots, \alpha_k$, with respective multiplicities $c_1, \ldots, c_k$. We construct a system of equations in m based on the roots and their multiplicities:

$\forall i \in \{1, \ldots, k\}$	α_i^m	$=$	$q(\alpha_i)$
	$m\alpha_i^{m-1}$	$=$	$q'(\alpha_i)$
	$m(m-1)\alpha_i^{m-2}$	$=$	$q''(\alpha_i)$
	$\vdots$	$\vdots$	$\vdots$
	$m(m-1)\ldots(m-c_i+2)\alpha_i^{m-c_i+1}$	$=$	$q^{(c_i-1)}(\alpha_i)$

Each root α_i contributes c_i equations to the system. Their left-hand sides are x^m and its first $c_i - 1$ derivatives evaluated at α_i. The respective right-hand sides are $q(x)$ and its first $c_i - 1$ derivatives evaluated at α_i. For example, if there are three distinct roots $\alpha_1, \alpha_2, \alpha_3$ with respective multiplicities $1, 2, 3$, the system is:

α_1^m	$=$	$q(\alpha_1)$
α_2^m	$=$	$q(\alpha_2)$
$m\alpha_2^{m-1}$	$=$	$q'(\alpha_2)$
α_3^m	$=$	$q(\alpha_3)$
$m\alpha_3^{m-1}$	$=$	$q'(\alpha_3)$
$m(m-1)\alpha_3^{m-2}$	$=$	$q''(\alpha_3)$

Let the system be S. We will now show that $A^m = q(A)$ if and only if m is a solution of S.

Proposition 18. *If $p(x) \in \mathbb{C}[x]$ and α is a root of p with multiplicity c, then α is a root of $p^{(0)}, p^{(1)}, \ldots, p^{(c-1)}$, but not $p^{(c)}$.*

Proof. Easy induction on c.

■

Proposition 19. *$A^m = q(A)$ if and only if m is a solution of S.*

Proof.

$$
\begin{aligned}
& A^m = q(A) \\
\iff & \\
& f_A \text{ divides } x^m - q(x) \\
\iff & \\
& \text{for all } i \in \{1, \ldots, k\},\ x^m - q(x) \text{ has root } \alpha_i \text{ with multiplicity at least } c_i \\
\iff & \\
& \text{for all } i \in \{1, \ldots, k\},\ \alpha_i \text{ is a root of } (x^m - q(x))^{(0)}, \ldots, (x^m - q(x))^{(c_i - 1)} \\
\iff & \\
& m \text{ solves } S
\end{aligned}
$$

■

Therefore, the system S is equivalent to the matrix equation $A^m = q(A)$. Now we will describe a procedure for solving it.

First, we perform some preliminary calculations. We check directly whether $m = 0$ is a witness for the problem instance. If so, we are done. Assume otherwise. Next, we check if some α_i is 0. If this is the case, the left-hand sides of the equations contributed by α_i are 0 for all m. We directly evaluate the right-hand sides. If one of them is non-zero, we can conclude the problem instance is negative. If all of them are 0, then the equations contributed by α_i are vacuously satisfied, so we discard them from the system. We may now assume that m ranges over $\mathbb{N}^+$ and that $\alpha_i \neq 0$ for all i.

Second, we consider three cases:

1. Some α_i is not a root of unity.
2. All α_i are roots of unity, unrepeated in f_A.
3. All α_i are roots of unity, at least one of which is repeated in f_A.

Case 1. We use the bound on m from Proposition 15. It is polynomial in the length of the input, so it suffices to compute A^m for all m up to the bound and compare with $q(A)$ at each step.

Case 2. The system S contains only equations of the form $\alpha^m = q(\alpha)$, with α a root of unity. Consider one such equation and let the order of α be s. For all values of m, the left-hand side is an s-th root of unity. The right-hand side $q(\alpha)$ is an s-th root of unity if and only if $q(\alpha)$ has the same coordinates in $\mathbb{Q}(\alpha)$ as one of $\alpha, \ldots, \alpha^s$. We check this using arithmetic in $\mathbb{Q}(\alpha)$.

If $q(\alpha)$ is not an s-th root of unity, then the problem instance is negative. Otherwise, the equation $\alpha^m = q(\alpha) = \alpha^r$ is equivalent to the congruence $m \equiv r \bmod s$, where r and s are known. We transform the entire system S in this way into an equivalent system of linear congruences, which we then solve in polynomial time [10].

Case 3. Let α be a root of unity that is repeated in f_A. Then S contains the equations:

$$\begin{aligned} \alpha^m &= q(\alpha) \\ m\alpha^{m-1} &= q'(\alpha) \end{aligned}$$

By the preliminary analysis, the left-hand sides of these two equations are non-zero. We check whether the right-hand sides are 0 using polynomial division of q and q' by f_α. If either right-hand side is 0, we conclude the instance is negative. Otherwise, we divide the two equations to obtain

$$m = \frac{q'(\alpha)}{q(\alpha)}\alpha$$

We directly calculate the right-hand side using arithmetic in $\mathbb{Q}(\alpha)$ and check if it is in $\mathbb{N}$. If not, then the instance is negative. Otherwise, we need only determine if this $m = m_0$ is a solution. We check that m_0 satisfies each equation $(\alpha^m)^{(t)} = q^{(t)}(\alpha)$ using arithmetic in $\mathbb{Q}(\alpha)$. For the left-hand side, we use $\alpha^u = \alpha^{u \bmod s}$, where s is the order of α.

This completes the proof that ORBIT is decidable in polynomial time.

Chapter 4

The One-dimensional Orbit Problem

The one-dimensional Orbit problem is defined as:

1D Orbit
Given vectors $x, y \in \mathbb{Q}^n$ and a matrix $A \in \mathbb{Q}^{n \times n}$,
do there exist $m \in \mathbb{N}$ and $k \in \mathbb{Q}$ such that $A^m x = ky$?

In this chapter, we show that 1D Orbit is decidable, and in fact decidable in polynomial time, as Kannan and Lipton conjectured. Our proof has a similar shape to that of the zero-dimensional version. First we reduce the problem to a one-dimensional version of Matrix Power and then use techniques from algebraic number theory to solve it.

4.1 Reduction

Define 1D Matrix Power to be the following decision problem:

1D Matrix Power
Given matrices $A, D \in \mathbb{Q}^{n \times n}$, do there
exist $m \in \mathbb{N}$ and $k \in \mathbb{Q}$ such that $A^m = kD$?

We will reduce 1D Orbit to 1D Matrix Power. The reduction is essentially the same as the one given in Section 3.1 for the zero-dimensional version. Given an instance (A, x, y) of 1D Orbit, we calculate

$$\nu = \max\{m \mid x, Ax, \ldots, A^m x \text{ are linearly independent}\}$$

and set

$$\mathcal{B} = \{x, Ax, \ldots, A^\nu x\}$$

We compute the coordinates $y_\mathcal{B}$ of y and the coordinates $[b_0, \ldots, b_\nu]^T$ of $A^{\nu+1}x$ with respect to $\mathcal{B}$, and set $x_\mathcal{B} = [1, 0, \ldots, 0]^T$. If $y \notin span(\mathcal{B})$, then the problem instance is negative and we are done. Otherwise, we have

$$A^m x = ky \iff M^m x_\mathcal{B} = k y_\mathcal{B} \iff M^m = k \left[\begin{array}{cccc} y_\mathcal{B} & My_\mathcal{B} & \ldots & M^\nu y_\mathcal{B} \end{array} \right]$$

where

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

This gives an equivalent instance of 1D Matrix Power. Next, we reduce further to the polynomial version of 1D Matrix Power, defined as the following decision problem:

1D Matrix Power (Polynomial Version)
Given a matrix $A \in \mathbb{Q}^{n \times n}$ and a polynomial $q \in \mathbb{Q}[x]$,
do there exist $m \in \mathbb{N}$ and $k \in \mathbb{Q}$ such that $A^m = kq(A)$?

Given an instance (A, D) of 1D Matrix Power, we solve the system $q(A) = D$ in the unknowns $q_0, \ldots, q_{n-1}$. If it has no solution, then we conclude that (A, D) is a negative instance of 1D Matrix Power. Otherwise, we choose a solution q and output (A, q). Then we have

$$A^m = kD \iff A^m = kq(A)$$

which completes the reduction.

4.2 Solution

Suppose we have an instance (A, q) of 1D MATRIX POWER (POLYNOMIAL VERSION). We compute the minimal polynomial f_A of A, canonical representations of its roots $\alpha_1, \ldots, \alpha_s$ and their respective multiplicities $c_1, \ldots, c_s$. Then we construct a system S of equations in m and k:

$\forall i \in \{1, \ldots, s\}$	α_i^m	$=$	$kq(\alpha_i)$
	$m\alpha_i^{m-1}$	$=$	$kq'(\alpha_i)$
	$m(m-1)\alpha_i^{m-2}$	$=$	$kq''(\alpha_i)$
	$\vdots$	$\vdots$	$\vdots$
	$m(m-1)\ldots(m-c_i+2)\alpha_i^{m-c_i+1}$	$=$	$kq^{(c_i-1)}(\alpha_i)$

The system is very similar to the one in the zero-dimensional case, but here we have an additional unknown k. Each root α_i contributes c_i equations to the system. The left-hand sides are x^m and its first $c_i - 1$ derivatives, evaluated at α_i. The respective right-hand sides are $kq(x)$ and its first $c_i - 1$ derivatives evaluated at α_i.

Proposition 20. *$A^m = kq(A)$ if and only if m and k satisfy the system S.*

Proof. The same as in Proposition 19, with the polynomial $q(x)$ replaced by $kq(x)$. We point out that the proof does not require $k \in \mathbb{Q}$. In particular, if we have some solution $(m, k) \in (\mathbb{N} \times \mathbb{C})$ of S, then $A^m = kq(A)$, which gives $k \in \mathbb{Q}$.

■

Now we will focus on solving S. We start with some preliminary analysis.

1. We check if $k = 0$ has a corresponding $m \in \mathbb{N}$ which satisfies the system, using the algorithm for the zero-dimensional case. If so, we are done. Assume otherwise.
2. Let c be the maximum of the multiplicities $c_1, \ldots, c_s$. For each $m = 0, \ldots, c-2$, we calculate A^m and check if it is a multiple of $q(A)$. If so, we are done. Otherwise, we can assume that $m \geq c - 1$.
3. We check if 0 is among the roots $\alpha_1, \ldots, \alpha_s$. If so, then the equations contributed by this root are all of the form $0 = kq^{(t)}(0)$. Since $k \neq 0$, this is equivalent to $q^{(t)}(0) = 0$, which we can easily check. If it holds, we dismiss the equation as vacuous. If not, then we are done. Now we can assume that $\alpha_i \neq 0$.

4. Finally, we check if the system contains some equation with right-hand side $kq^{(t)}(\alpha_i)$ equal to 0. This is done with a polynomial division of $q^{(t)}(x)$ by the minimal polynomial of α_i. If so, then we can conclude the system has no solution, because the previous steps guarantee the left-hand sides are all non-zero.

Now we can assume that both sides of all equations in S are non-zero. Next, we compute all quotients α_i/α_j, obtaining their canonical representations. We consider three cases:

1. Some quotient is not a root of unity.
2. All quotients are roots of unity, and all the roots of f_A are unrepeated.
3. All quotients are roots of unity, and some roots of f_A are repeated.

Case 1. Suppose some quotient α_i/α_j is not a root of unity. We have the equations

$$\alpha_i^m = kq(\alpha_i)$$

$$\alpha_j^m = kq(\alpha_j)$$

Hence,

$$\left(\frac{\alpha_i}{\alpha_j}\right)^m = \frac{q(\alpha_i)}{q(\alpha_j)}$$

so m is bounded by a polynomial in the length of the input, according to Proposition 16. We compute A^m for all values of m up to the bound, and at each step check if A^m is a multiple of $q(A)$.

Case 2. Now suppose all quotients are roots of unity, and there are no repeated roots in f_A. Then the system S contains only equations of the shape $\alpha_i^m = kq(\alpha_i)$. Then S is equivalent to

$$\bigwedge_{i<j}\left(\frac{\alpha_i}{\alpha_j}\right)^m = \frac{q(\alpha_i)}{q(\alpha_j)} \wedge k = \frac{\alpha_1^m}{q(\alpha_1)}$$

It is sufficient to determine whether there exists an m which solves

$$\bigwedge_{i<j} \left(\frac{\alpha_i}{\alpha_j}\right)^m = \frac{q(\alpha_i)}{q(\alpha_j)}$$

If so, then $k = \alpha_1^m / q(\alpha_1)$ is guaranteed to be rational by Proposition 20 and we need not check it.

Consider a single equation $(\alpha_i/\alpha_j)^m = q(\alpha_i)/q(\alpha_j)$. We need to determine if the right-hand side is a root of unity. This is easy if we have its canonical description, but it is not obvious whether this description has polynomial length. Evaluating q at α_i requires a polynomial number of additions and multiplications, and each operation results in a polynomial increase in the length of the representation. This gives an exponential upper bound on the length of the representation of the result, which is not strong enough for us. It could be that a canonical description of $q(\alpha_i)$ may be computed in polynomial time, but proving this will require more precise analysis.

To avoid this difficulty, we find a primitive element θ, such that $\mathbb{Q}(\alpha_i, \alpha_j) = \mathbb{Q}(\theta)$. The proof of the Primitive Element theorem (page 37 in [5]) shows that such a primitive element is $\alpha_i + c\alpha_j$ for some small positive integer c which may be computed in polynomial time. Given θ, the next task is to compute the standard representations of α_i and α_j with respect to $\mathbb{Q}(\theta)$. This may be done using an algorithm for the *Field Membership* problem:

> Given canonical representations of α and θ,
> determine whether $\alpha \in \mathbb{Q}(\theta)$, and if so, find the
> standard representation of α with respect to $\mathbb{Q}(\theta)$.

Section 4.5 of [4] shows how to solve this problem in polynomial time. We calculate the coordinates of α_i and α_j in $\mathbb{Q}(\theta)$ and substitute them into the equation. We use operations within $\mathbb{Q}(\theta)$ to compute the standard representations $l(\theta)$ of α_i/α_j and $r(\theta)$ of $q(\alpha_i)/q(\alpha_j)$. Now we have the equation $l(\theta)^m = r(\theta)$.

If α_i/α_j has order d, then the powers $l(\theta)^1, \ldots, l(\theta)^d$ are exactly the d-th roots of unity. We compute their standard representations in polynomial time, and compare each with $r(\theta)$. If there is no match, then $q(\alpha_i)/q(\alpha_j)$ is not a d-th root of unity, and the system has no solution. Otherwise, if $r(\theta) = l(\theta)^t$, the equation $(\alpha_i/\alpha_j)^m = q(\alpha_i)/q(\alpha_j)$ is equivalent to the congruence $m \equiv t \bmod d$.

We process all the quotient equations in this manner, obtaining an equivalent system of congruences, which we can solve in polynomial time [10].

Case 3. Now suppose all quotients are roots of unity, and some roots of f_A have multiplicity greater than 1. We transform S into S' in the following way. First, we include in S' the quotients of all equations $\alpha_i^m = kq(\alpha_i)$ as in Case 2. Second, for each repeated root α_i of f_A, we take the quotient of its first and second equation, its second and third, and so on. Third, we include the equation $k = \alpha_1^m / q(\alpha_1)$.

For example, suppose f_A has roots $\alpha_1, \alpha_2, \alpha_3$ with respective multiplicities $1, 2, 3$. The original system S is:

α_1^m	$=$	$kq(\alpha_1)$
α_2^m	$=$	$kq(\alpha_2)$
$m\alpha_2^{m-1}$	$=$	$kq'(\alpha_2)$
α_3^m	$=$	$kq(\alpha_3)$
$m\alpha_3^{m-1}$	$=$	$kq'(\alpha_3)$
$m(m-1)\alpha_3^{m-2}$	$=$	$kq''(\alpha_3)$

We transform it into S':

$(\alpha_1/\alpha_2)^m$	$=$	$q(\alpha_1)/q(\alpha_2)$
$(\alpha_2/\alpha_3)^m$	$=$	$q(\alpha_2)/q(\alpha_3)$
$(\alpha_1/\alpha_3)^m$	$=$	$q(\alpha_1)/q(\alpha_3)$
α_2/m	$=$	$q(\alpha_2)/q'(\alpha_2)$
α_3/m	$=$	$q(\alpha_3)/q'(\alpha_3)$
$\alpha_3/(m-1)$	$=$	$q'(\alpha_3)/q''(\alpha_3)$
k	$=$	$\alpha_1^m/q(\alpha_1)$

It is easy to see that S' is equivalent to S. Now we will solve it.

First, we solve the set of equations where m appears only in the exponent as we did in Case 2 by computing a primitive element for each equation. This subsystem either has no solution, or is solved by all m satisfying some congruence. Assume the latter and suppose that the calculation returns a solution $m \equiv t_1 \bmod t_2$.

The remainder of S' is easier to solve. Each equation contributed by a repeated root α_i has the shape

$$\frac{\alpha_i}{m-t} = \frac{q^{(t)}(\alpha_i)}{q^{(t+1)}(\alpha_i)}$$

for a constant t, as we only took the quotients of successive α_i equations. This is equivalent to

$$m = t + \frac{q^{(t+1)}(\alpha_i)}{q^{(t)}(\alpha_i)}\alpha_i$$

For each such equation, we compute this expression using operations within $\mathbb{Q}(\alpha_i)$ and check if it is a positive integer. If not, then the system clearly has no solution. Otherwise, this equation points to a single candidate m. We do this for all equations where m appears outside the exponent. If they point to different values of m, the system has no solution. Otherwise, S' is equivalent to

$$\boxed{\begin{array}{rcl} k & = & \alpha_1^m / q(\alpha_1) \\ m & = & m_0 \\ m & \equiv & t_1 \bmod t_2 \end{array}}$$

We check whether the candidate m_0 obeys the congruence, and we are done.

This completes the proof that 1D ORBIT is decidable in polynomial time.

Chapter 5

Conclusion

The idea of devising a system of equations equivalent to the matrix equation is of central importance to our result. It allowed us to prove that the one-dimensional Orbit problem is decidable in polynomial time, as Kannan and Lipton conjectured in 1986. Additionally, it led to a concise argument for the zero-dimensional version.

Future work should focus on proving the two-dimensional Orbit problem decidable. A similar reduction leads to the problem of deciding whether there exist $m \in \mathbb{N}$, $u, v \in \mathbb{Q}$ such that

$$A^m = up(A) + vq(A)$$

where $p, q \in \mathbb{Q}[x]$. The idea of constructing an equivalent system based on the roots of f_A is also applicable. However, solving the system is more difficult, because the trick of considering quotients cannot be used directly to remove one of the unknowns from consideration.

A related problem is the *non-homogeneous one-dimensional* version, which asks to determine whether there exist $m \in \mathbb{N}$, $t \in \mathbb{Q}$ such that

$$A^m x = y + tz$$

That is, whether for some m, $A^m x$ is on a specified line which does *not* pass through the origin. This problem appears very similar to the two-dimensional Orbit problem, and it is possible that solving it will yield insight into the two-dimensional case.

Bibliography

[1] R. Kannan and R. Lipton. *Polynomial-Time Algorithm for the Orbit Problem.* Journal of the ACM, Vol **33**, No. 4, pp. 808–821 (1986).

[2] M. Neunhöffer and C. Praeger. *Computing Minimal Polynomials of Matrices.* LMS Journal of Computation and Mathematics, **11**, pp 252-279 (2008).

[3] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. *Factoring Polynomials with Rational Coefficients.* Mathematische Annalen **261** (4): pp 515–534 (1982).

[4] H. Cohen. *A Course in Computational Algebraic Number Theory.* Springer (1993).

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

[6] M. Mignotte. *Some Useful Bounds.* Computer Algebra, pp 259-263, Springer (1982).

[7] V. Pan. *Optimal and Nearly Optimal Algorithms for Approximating Polynomial Zeros.* Computers & Mathematics with Applications, **31** (12): pp 97-138 (1996).

[8] V. Halava, T. Harju, M. Hirvensalo, J. Karhumäki. *Skolem's Problem — On the Border Between Decidability and Undecidability.* TUCS Technical Report No 683 (2005).

[9] P. Blanksby and H. Montgomery. *Algebraic Integers Near the Unit Circle.* Acta Arith, pp 355-369 (1971).

[10] J. von zur Gathen and M. Sieveking. *Weitere zum Erfüllungsproblem polynomial äquivalente kombinatorische Aufgaben.* Komplexität von Entscheidungsproblemen, LNCS **43**, Springer-Verlag, pp 49-71 (1976).

[11] A. Tiwari. *Termination of Linear Programs.* CAV 2004. LNCS **3114**, pp. 70-82. (2004)

[12] M. Braverman. *Termination of Integer Linear Programs.* CAV 2006. LNCS, vol. **4144**. Springer-Verlag. (2004)

[13] N. Vereshchagin. *Occurrence of Zero in a Linear Recursive Sequence.* Math. Notes **38**, nos 1–2. (1985)

[14] V. D. Blondel and N. Portier. *The Presence of a Zero in an Integer Linear Recurrent Sequence is NP–hard to Decide.* Linear Algebra and Its Applications, **351-352**. (2002)

[15] V. Halava, T. Harju, M. Hirvensalo. *Positivity of Second Order Linear Recurrent Sequences.* Discrete Appl. Math. **154**. (2006)

[16] A. Podelski and A. Rybalchenko. *A Complete Method for the Synthesis of Linear Ranking Functions.* VMCAI 2004: Verification, Model Checking, and Abstract Interpretation. LNCS, Springer-Verlag. (2004)

[17] A. Shidlovskii. *Transcendental Numbers.* New York: de Gruyter studies in mathematics. (1989)