On Matrix Powering in Low Dimensions

Esther Galby[1], Joël Ouaknine[2], and James Worrell[2]

1 **École Normale Supérieure de Rennes, France**
2 **Departement of Computer Science, Oxford University, UK**

Abstract

We investigate the ***Matrix Powering Positivity Problem***, PosMatPow: given an $m \times m$ square integer matrix M, a linear function $f : \mathbb{Z}^{m \times m} \to \mathbb{Z}$ with integer coefficients, and a positive integer n (encoded in binary), determine whether $f(M^n) \geq 0$. We show that for fixed dimensions m of 2 and 3, this problem is decidable in polynomial time.

1998 ACM Subject Classification F.2.1 Numerical Algorithms and Problems

Keywords and phrases matrix powering, complexity, Baker's theorem

1 Introduction

An important theme in theoretical computer science is the complexity of performing calculations on large (often exponential) but succinctly presented structures. Notable examples include the analysis of various problems on succinctly represented graphs [10, 20], as well as the study of PosSLP [2], the problem of determining whether an arithmetic circuit, with addition, multiplication, and subtraction gates, evaluates to a positive integer. Allender *et al.* show that a substantial fragment of modern numerical analysis reduces in polynomial time to PosSLP, as do several other well-known questions such as the Sum-of-Square-Roots Problem, which itself is instrumental, among others, in solving the Euclidean Travelling Salesman Problem. Very recently, an interesting 'challenge' (spiritually attributed to Dyson) was proposed by Lipton: find an efficient algorithm that, given an integer n, determines whether the reversal of 2^n as a decimal number is a power of 5 [13].[1]

In all the above examples, the central issue is that the objects in question, while succinctly presented, are fundamentally of exponential size. In the case of PosSLP, for instance, it is trivial to construct an arithmetic circuit whose integer output is doubly exponential in the size of the circuit, i.e., requiring an exponential number of bits. In light of this observation, one might conjecture that the existence of polynomial-time algorithms for performing non-trivial calculations on such entities is generally doomed.

Perhaps surprisingly, polynomial-time algorithms *can* occasionally be found. For example, by exploiting deep results in analytic number theory, Hirvensalo *et al.* showed in [12] that the most significant digit in base 3 of expressions such as 2^n and the nth Fibonacci number could be computed in time polynomial in the size of n, something which at first sight is far from obvious.

In this paper, we are concerned with the ***Matrix Powering Positivity Problem***, PosMatPow: given an $m \times m$ square integer matrix M, a linear function $f : \mathbb{Z}^{m \times m} \to \mathbb{Z}$ with integer coefficients, and a positive integer n, determine whether $f(M^n) \geq 0$. Note that in general, the entries of M^n have exponentially many bits (in the size of n), even if M

[1] Here by "efficient" one requires a running time polynomial in the size, or bit length, of the representation of n.

is encoded in unary, so a naive calculation based (for example) on iterated squaring would necessarily require exponential time.[2]

Problems involving powers of matrices appear in a wide range of contexts. For 'small' (i.e., unary-encoded) powers, the complexity of powering has been thoroughly investigated in [15], and shown to lie in TC^0 for any fixed dimension. The complexity of PosMatPow (for 'large'—i.e., binary-encoded—powers) is instrumental in determining the overall complexity of the main algorithms presented in [18, 17] to decide positivity of linear recurrence sequences of low order. Allender *et al.* study the closely related problem of BitMatPow in [1] in which one seeks to determine the value of a specified bit in a large power of a given matrix. Already in dimension 2—and thus *a fortiori* in higher dimensions as well—Allender *et al.* provide some evidence that BitMatPow cannot be solved in polynomial time, although it is known to lie in the Counting Hierarchy CH.

Since the publication of Allender *et al.*'s seminal work [2], determining the complexity of PosSLP has become a problem a major importance; Etessami and Yannakakis, for example, show in [8] that the fundamental problem of finding mixed strategy profiles close to exact Nash equilibria in three-person games is PosSLP-hard. PosSLP lies in CH [2] but is not believed to belong to NP and much less to P. Unfortunately, no non-trivial lower bounds for it are known at present.

PosMatPow can be shown to reduce in polynomial time to both PosSLP and BitMatPow (increasing for the latter the dimension by 3). Thus in addition to upper complexity bounds, lower bounds for PosMatPow would be of significant interest.

The main result of this paper is that in dimensions 2 and 3, PosMatPow can be solved in polynomial time.[3] This upper bound is achieved by attacking the problem via spectral techniques, and making use of sophisticated tools from algebraic number theory, transcendence theory, and numerical analysis. We leave as a challenging open problem the complexity of PosMatPow in higher dimensions.

2 Preliminaries

We review some of the mathematical apparatus used throughout this paper. Since our approach is predicated on spectral techniques, the efficient manipulation of algebraic numbers is of central importance. The reader may however wish to skip this section on a first reading and proceed directly to Sections 3 and 4 in which the main algorithms are presented.

2.1 Algebraic Numbers and Baker's Theorem

For $p \in \mathbb{Z}[x]$ a univariate polynomial with integer coefficients, let us denote by $||p||$ the bit length of its representation as a list of coefficients encoded in binary. Note that the degree of p is at most $||p||$, and the ***height*** of p—i.e., the maximum magnitude of its coefficients—is at most $2^{||p||}$.

A complex number α is ***algebraic*** if it is the root of a univariate polynomial with integer coefficients. The ***defining polynomial*** of α, denoted p_α, is the unique polynomial of least degree, and whose coefficients do not have common factors, which vanishes at α. The ***degree*** and ***height*** of α are respectively those of p_α.

[2] Note that we are working in the standard bit-model of complexity theory, rather than the unit-cost arithmetic model in which PosMatPow (and PosSLP) would trivially be in polynomial time.

[3] In dimension 3, this result requires that the base matrix M be encoded in unary. The exponent n and linear function f are, however, always encoded in binary.

A standard representation[4] for algebraic numbers is to encode α as a tuple comprising its defining polynomial together with rational approximations of its real and imaginary parts of sufficient precision to distinguish α from the other roots of p_α. More precisely, α can be represented by $(p_\alpha, a, b, r) \in \mathbb{Z}[x] \times \mathbb{Q}^3$ provided that α is the unique root of p_α inside the circle in $\mathbb{C}$ of radius r centred at $a + bi$. A separation bound due to Mignotte [16] asserts that for roots $\alpha \neq \beta$ of a polynomial $p \in \mathbb{Z}[x]$, we have

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

where d and H are respectively the degree and height of p. Thus if r is required to be less than a quarter of the root-separation bound, the representation is well-defined and allows for equality checking. Given a polynomial $p \in \mathbb{Z}[x]$, it is well-known how to compute standard representations of each of its roots in time polynomial in $||p||$ [19, 7, 4]. Thus given α an algebraic number for which we have (or wish to compute) a standard representation, we write $||\alpha||$ to denote the bit length of this representation. From now on, when referring to computations on algebraic numbers, we always implicitly refer to their standard representations.

Given algebraic numbers α and β, one can test whether $\alpha = \beta$ as well as membership in $\mathbb{R}$ in polynomial time. One can also compute $\alpha+\beta$, $\alpha\beta$, $1/\alpha$ (for non-zero α), $\overline{\alpha}$, $|\alpha|$, $\mathrm{Re}(\alpha)$, and $\mathrm{Im}(\alpha)$, all of which are algebraic, in polynomial time. Moreover, if $\alpha \in \mathbb{R}$, deciding whether $\alpha > 0$ can also be done in polynomial time. Efficient algorithms for all these tasks can be found in [7, 4].

We will also need the following result.

▶ **Proposition 1.** Given algebraic numbers α and β, together with an integer $n \geq 0$, one can decide whether $\alpha^n = \beta$ in time polynomial in both $||\alpha|| + ||\beta||$ and $||n|| = \lceil \log_2 n \rceil$.

Proposition 1 can be proved directly using elementary algebraic number theory. Alternatively, it is an immediate consequence of the following lemma:

▶ **Lemma 1.** *Let α and β be non-zero complex algebraic numbers, and consider the free abelian group L under addition given by $L = \{(u, v) \in \mathbb{Z}^2 : \alpha^u \beta^v = 1\}$. L has a basis whose vectors are polynomially bounded in $||\alpha|| + ||\beta||$. Moreover, such a basis can be computed in time polynomial in $||\alpha|| + ||\beta||$.*

Note in the above that the bound is on the *magnitude* of the vectors in the basis (rather than the bit length of their representation), which follows from a deep result of Masser [14]. For a proof of Lem. 1, see also [11, 6].

Proposition 1 now easily follows: given α and β, compute a basis B for the corresponding free abelian group L, and decide whether $(n, -1) \in L = \mathrm{span}(B)$, which can be done in polynomial time. For example, if L has rank 2, i.e., $B = \{(u_1, v_1), (u_2, v_2)\}$ for some integers u_1, v_1, u_2, and v_2, the problem is equivalent to determining whether there exist integers x and y such that $x(u_1, v_1) + y(u_2, v_2) = (n, -1)$. Since the u_i's and v_i's have magnitude polynomial in $||\alpha|| + ||\beta||$, the size of this problem instance is logarithmic (hence *a fortiori* polynomial) in $||\alpha|| + ||\beta||$ and polynomial in $||n||$. Since solving linear equations over the integers can be carried out in polynomial time, the desired result follows.

[4] Note that this representation is not unique.

We also record the following bounds, which are immediately derived from classical analytic results on polynomials (see, e.g., [21]). For α a non-zero algebraic number of height H, we have

$$\frac{1}{H+1} < \alpha < H+1\,. \tag{2}$$

If E and F are two fields such that $F \subseteq E$, we say that E is an *extension* of F and the ***degree of*** E ***over*** F, denoted $[E:F]$, is defined to be the dimension of E considered as a vector space over F. The degree is multiplicative: if E is an extension of F and F is itself an extension of L, then E is an extension of L of degree $[E:L] = [E:F][F:L]$.

A ***number field*** is an extension of $\mathbb{Q}$ of finite degree. In particular, given any algebraic numbers $\alpha_1, \ldots, \alpha_k$, $\mathbb{Q}(\alpha_1, \ldots, \alpha_k)$ is the number field comprising all complex numbers that are equal to some polynomial in $\alpha_1, \ldots, \alpha_k$ with rational coefficients.

Let $p \in \mathbb{Z}[x]$ be a quadratic polynomial with roots α and β. Then $[\mathbb{Q}(\alpha) : \mathbb{Q}] \leq 2$ and $\beta \in \mathbb{Q}(\alpha)$. On the other hand, if p is a cubic polynomial with roots α, β, and γ, then $[\mathbb{Q}(\alpha, \beta) : \mathbb{Q}] \leq 6$ and $\gamma \in \mathbb{Q}(\alpha, \beta)$. For K a number field with $\lambda, \overline{\lambda} \in K$, we have $[K(|\lambda|) : K] \leq 2$ since $|\lambda|^2 = \lambda\overline{\lambda}$. And also $[K(\nu) : \mathbb{Q}] \leq [K : \mathbb{Q}][\mathbb{Q}(\nu) : \mathbb{Q}]$ for any number field K and algebraic number ν.

Finally, we give a version of Baker's deep theorem on linear forms in logarithms. The particular statement we have chosen is a sharp formulation due to Baker and Wüstholz [3]. In what follows, log refers to the principal value of the complex logarithm function given by $\log z = \log|z| + i \arg z$, where $-\pi < \arg z \leq \pi$.

▶ **Theorem 2** (Baker and Wüstholz). *Let $\alpha_1, \ldots, \alpha_m \in \mathbb{C}$ be algebraic numbers different from 0 or 1, and let $b_1, \ldots, b_m \in \mathbb{Z}$ be integers. Write*

$$\Lambda = b_1 \log \alpha_1 + \ldots + b_m \log \alpha_m\,.$$

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

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

2.2 Matrix Powers and Linear Recurrence Sequences

We recall some basic facts about linear algebra and linear recurrence sequences. An excellent reference on the latter is [9].

Let $M \in \mathbb{Z}^{m \times m}$ be a square integer matrix of dimension m. In this paper, we work with a *binary* encoding of M in two dimensions, and with a *unary* encoding of M in three dimensions. Both encodings are denoted $||M||$, relying on context for disambiguation.

In two dimensions (with binary encoding), we note that eigenvalues of M have degree at most 2 and height at most $2^{2||M||}$.

In three dimensions (with unary encoding), eigenvalues of M have degree at most 3 and height at most $||M||^3$.

Let $f : \mathbb{Z}^{m \times m} \to \mathbb{Z}$ be a linear function with integer coefficients: $f(x_1, \ldots, x_{m^2}) = b_1 x_1 + \ldots + b_{m^2} x_{m^2}$ for integers $b_1, \ldots, b_{m^2}$. Since m-dimensional square integer matrices can be viewed as m^2-tuples of integers, we shall assume a fixed order for entries and freely apply such functions to square matrices. We always encode f as a list of its coefficients in binary, and denote the size of this encoding by $||f||$.

Let $p(x) = x^m - a_1 x^{m-1} - \ldots - a_m$ be the characteristic polynomial of M. For any $k \geq 0$, let $u_k = f(M^k)$. Then the sequence $\langle u_k \rangle_{k=0}^{\infty}$ is an integer linear recurrence sequence (LRS) obeying the recurrence

$$u_{k+m} = a_1 u_{k+m-1} + \ldots + a_m u_k \,. \tag{3}$$

Indeed, since $M^m - a_1 M^{m-1} - \ldots - a_m I = \mathbf{0}$ by the Cayley-Hamilton theorem, multiplying this equation by M^k, applying f on both sides and invoking linearity yields Eq. (3).

The characteristic polynomial of this LRS is p, hence the characteristic roots are the eigenvalues of M. Let us write

$$\mathrm{spec}(M) = \{\rho_1, \ldots, \rho_\ell, \lambda_1, \overline{\lambda_1}, \ldots, \lambda_p, \overline{\lambda_p}\} \,,$$

where each $\rho_i \in \mathbb{R}$ and each $\lambda_j \in \mathbb{C} \setminus \mathbb{R}$. There are now univariate polynomials $A_1, \ldots, A_\ell$ and $C_1, \ldots, C_p$ such that, for all $n \geq 0$,

$$u_k = \sum_{i=1}^{\ell} A_i(k) \rho_i^k + \sum_{j=1}^{p} \left(C_j(k) \lambda_j^k + \overline{C_j}(k) \overline{\lambda_j}^k \right) \,. \tag{4}$$

This expression is referred to as the ***exponential polynomial*** solution of $\langle u_k \rangle_{k=0}^{\infty}$. The polynomials A_i have real algebraic coefficients and the polynomials C_j have complex algebraic coefficients. The degree of each of these polynomials is at most one less than the multiplicity of the corresponding eigenvalue; thus in particular, these polynomials are identically constant when M has no repeated eigenvalue. For fixed m, all coefficients appearing in these polynomials can be computed in time polynomial in $||\langle f, M \rangle||$ (whether M is encoded in binary or unary),[5] since they can be obtained by solving a system of linear equations involving the m constants $u_0, \ldots, u_{m-1}$. As a result, these coefficients all belong to $\mathbb{Q}(\mathrm{spec}(M))$, and their height (*qua* algebraic numbers) is bounded above by $2^{\mathcal{O}(||\langle f, M \rangle||)}$ (again regardless of whether M is encoded in binary or unary).

2.3 Approximation Algorithms for Transcendental Functions

Finally, we recall some classical numerical algorithms which are later invoked to efficiently compute approximations of transcendental functions applied to algebraic numbers.

Given a real number t and a positive integer m, we say that $q \in \mathbb{Q}$ is an ***m-bit approximation of*** t if $|t - q| < 2^{-m}$. We also sometimes refer to the calculation of such a q as "computing m bits of t", even though strictly speaking this form of words is not perfectly accurate.

▶ **Proposition 2.**

1. There exists an algorithm which takes as input a real algebraic number $\rho > 0$, together with a positive integer m, and returns an m-bit approximation of $\log \rho$ in time polynomial in both $||\rho||$ and m.
2. There exists an algorithm which takes as input two non-zero real algebraic numbers a and b, together with a positive integer m, and returns an m-bit approximation of $\arctan b/a$ in time polynomial in both $||a|| + ||b||$ and m.
3. There exists an algorithm which takes as input a positive integer m and returns an m-bit approximation of π in time polynomial in m.

[5] We write $||\langle f, M \rangle||$ to denote the size of the joint encoding of f and M.

Proposition 2 follows from classical approximation results for transcendental functions due to Brent [5], together with the fact that we can compute approximations of algebraic numbers with polynomially many bits in polynomial time (see, e.g., [19]). Below we sketch the process for (1) of Prop. 2, relying on the following result.

▶ **Theorem 3** (Brent). *For any fixed real numbers $0 < a < b$, there exists an algorithm which, given an integer $p \geq 0$, evaluates $\log x$ in time $\mathcal{O}(p \log^2 p \log\log p)$, with relative error at most $\mathcal{O}(2^{-p})$, uniformly for all $x \in [a, b]$.*

Let ρ and m be as in Prop. 2 (1), and denote the height of ρ by H, recalling that $H \leq 2^{||\rho||}$. By Eq. (2), we have $\rho > 1/(H+1)$, and for simplicity assume that $\rho < 1$; the alternative can be handled in a similar manner as what follows. We now aim to select a positive integer k with certain properties, to be listed in the remainder of this proof; we will then choose a specific value for k later on in such a way as to discharge all our assumptions.

The first requirement is that k be at most polynomial in $||\rho||$ and m. Next, compute $r \in \mathbb{Q}$ with $1/(2H) < r \leq \rho$ such that $\rho - r < 2^{-k}$; this can be achieved in time polynomial in k by computing polynomially many bits of ρ. Fix the interval $[a, b] = [2, 4]$ in Thm. 3 and find $j \in \mathbb{N}$ such that $2^j r \in [2, 4]$. Thanks to our lower bound on r such j is at most polynomial in $||\rho||$ and can be obtained in polynomial time.

We now invoke Thm. 3 to compute $u \in \mathbb{Q}$ such that $\dfrac{|u - \log 2^j r|}{\log 2^j r} < \dfrac{1}{2^{m+3}}$ in time polynomial in m. Since $\log 2^j r \leq \log 4 < 2$, we have $|v - \log r| < 2^{-m-2}$, where $v = u - j \log 2$. By invoking Thm. 3 once more, we can compute $v' \in \mathbb{Q}$ in time polynomial in m such that $|v' - v| < 2^{-m-2}$, whence $|v' - \log r| < 2^{-m-1}$.

Since the derivative of $\log x$ at point r is r^{-1}, we conclude that $|\log \rho - \log r| < r^{-1} 2^{-k}$. If we make the additional requirement on k that $r^{-1} 2^{-k} < 2^{-m-1}$, we can combine with our previous inequality to obtain $|v' - \log \rho| < 2^{-m}$, yielding an m-bit approximation of $\log \rho$ as required.

Finally, it remains to show that $k \in \mathbb{N}$ can be chosen so as to meet our various assumptions, a straightforward task which we leave to the reader.

3 The Two-Dimensional Matrix Powering Positivity Problem

The main result of this section is the following:

▶ **Theorem 4.** *In two dimensions,* PosMatPow *(with full binary encoding) is decidable in polynomial time.*

Proof. Consider an instance of the two-dimensional Matrix Powering Positivity Problem, comprising a linear function $f : \mathbb{Z}^4 \to \mathbb{Z}$, a 2×2 integer matrix M, and an integer $n \geq 0$. Assume that all this data is encoded in binary and denote by $||\langle f, M, n \rangle||$ the size of the instance. We wish to decide in polynomial time whether $f(M^n) \geq 0$. To this end, we consider the sequence $u_k = f(M^k)$ and study the exponential polynomial solution (Eq. 4) in which two cases arise: either (i) both eigenvalues of M are real (including the possibility of a single repeated real eigenvalue), or (ii) both eigenvalues are complex conjugates. In the latter, it is worth pointing out that the sign of the sequence will forever fluctuate.

In Case (i), let $\rho_1, \rho_2 \in \mathbb{R}$ be the eigenvalues of M. We distinguish two subcases depending on the multiplicity of ρ_1. If ρ_1 is repeated (i.e., $\rho_1 = \rho_2$), then for all $k \geq 0$,

$$u_k = (ak + b)\rho_1^k$$

where a and b are two real algebraic constants; recall moreover from Sec. 2.2 that ρ_1, a, and b can all be computed in polynomial time, and have representations of size linear in $||\langle f, M\rangle||$.

Assuming that $a \neq 0$ (the treatment being straightforward otherwise), it is easy to see that u_k has the same sign as

$$\begin{array}{ll} b\rho_1^k, & \text{when } k < -b/a \\ a\rho_1^k, & \text{when } k > -b/a\,. \end{array}$$

Note that for $-b/a \in \mathbb{N}$, $u_{-b/a} = 0$. It therefore remains to compare n to $-b/a$. Since arithmetic and inequality testing on algebraic numbers can be performed in polynomial time, the desired result follows.

Now assume that $\rho_1 \neq \rho_2$. Then we can compute two real algebraic constants a and b such that for all $k \geq 0$,

$$u_k = a\rho_1^k + b\rho_2^k\,.$$

If any of ρ_1, ρ_2, a, or b is zero, the solution is immediate. Thus assume otherwise and consider the sequence

$$\frac{u_k}{\rho_1^k} = a + b\left(\frac{\rho_2}{\rho_1}\right)^k\,.$$

We first check whether u_n is zero for our given exponent n, or equivalently whether $(\rho_2/\rho_1)^n = -a/b$; by Prop. 1, this can be done in polynomial time. Otherwise, we have $u_n/\rho_1^n > 0$ iff

$$b\left(\frac{\rho_2}{\rho_1}\right)^n > -a\,. \tag{5}$$

Assume without loss of generality that the expressions on both sides of the inequality are positive (something which is readily checked). Then Eq. (5) holds iff

$$\log|b| + n\log|\rho_2| - n\log|\rho_1| > \log|a|\,.$$

In other words, the sign of the expression

$$\Lambda = \log|b| + n\log|\rho_2| - n\log|\rho_1| - \log|a|$$

determines that of u_n (modulo the sign of ρ_1^n). Note that $\rho_1, \rho_2, a, b \in \mathbb{Q}(\rho_1)$, so that the number field $\mathbb{Q}(|b|, |\rho_1|, |\rho_2|, |a|)$ has degree 2 over $\mathbb{Q}$. Moreover, we can easily compute an upper bound H on the heights of ρ_1, ρ_2, a, and b, such that $\log(H) = \mathcal{O}(||\langle f, M\rangle||)$. By Baker's theorem (Thm. 2), we then have

$$|\Lambda| > \exp\left(-(16\cdot 4\cdot 2)^{2(4+2)}(\log H)^4 \log n\right) = \frac{1}{n^{128^{12}(\log H)^4}} = \frac{1}{n^{||\langle f,M\rangle||^{\mathcal{O}(1)}}}\,.$$

Thus in order to determine the sign of Λ, it suffices to compute $||\langle f, M\rangle||^{\mathcal{O}(1)}\log_2 n = ||\langle f, M, n\rangle||^{\mathcal{O}(1)}$ bits of Λ, i.e., a polynomial number of bits in the size of our problem instance $\langle f, M, n\rangle$. By Prop. 2, $||\langle f, M, n\rangle||^{\mathcal{O}(1)}$-bit approximations of $\log|b|$, $\log|\rho_2|$, $\log|\rho_1|$, and $\log|a|$ can be obtained in polynomial time, whence the desired result follows.

We now turn to Case (ii), in which M has two complex eigenvalues λ and $\overline{\lambda}$. We have, for all $k \geq 0$, $u_k = c\lambda^k + \overline{c}\,\overline{\lambda}^k$, where c is a complex algebraic constant. Equivalently, letting $\theta = \arg\lambda$ and $\varphi = \arg c$, we can write

$$u_k = |c||\lambda|^k\cos(k\theta + \varphi)\,. \tag{6}$$

Note that $\cos(n\theta+\varphi)=0$ iff $(e^{i\theta})^n = e^{i(-\varphi\pm\pi/2)}$. Since $e^{i\theta}$ and $e^{i\varphi}$ are algebraic numbers with size linear in $||\langle f,M\rangle||$, by Prop. 1 the latter can be checked in time polynomial in $||\langle f,M,n\rangle||$.

Let us therefore assume that $u_n \neq 0$ for our given exponent n. We aim to bound (in absolute value) the expression $n\theta+\varphi$ away from $\pm\pi/2$ (modulo 2π). To this end, write

$$\Gamma = \arg e^{i(n\theta+\varphi)} = n\theta+\varphi-2m\pi\,,$$

where m is the unique integer such that $-\pi<\Gamma\leq\pi$. The delicate situation is now if Γ is 'close' to $\pm\pi/2$. If that is not the case (for instance, if $||\Gamma|-\pi/2|>0.1$, say), then one can readily compute the sign of $\cos(n\theta+\varphi)$, and therefore that of u_n, in polynomial time. On the other hand, if Γ is close to $\pm\pi/2$ (for instance, if $||\Gamma|-\pi/2|<0.5$, say), then one can readily determine the sign of Γ. Assume that we are in the latter situation, and without loss of generality suppose that $\Gamma>0$. Write

$$\Lambda = \frac{\pi}{2}-\Gamma = \frac{1}{i}\left(n\log\frac{\lambda}{|\lambda|}+\log\frac{c}{|c|}+(1-4m)\log i\right)\,,$$

and let H be an upper bound for the heights of $\lambda/|\lambda|$ and $c/|c|$. Note that the degree of $\mathbb{Q}(\lambda,|\lambda|,c,|c|,i)$ over $\mathbb{Q}$ is at most 16. Since $|1-4m|\leq 2n+1$, if follows from Baker's theorem that

$$|\Lambda| > \exp\left(-768^{10}(\log H)^2\log(2n+1)\right) = \frac{1}{(2n+1)^{768^{10}(\log H)^2}} = \frac{1}{n^{||\langle f,M\rangle||^{\mathcal{O}(1)}}}\,, \tag{7}$$

since $\log(H)=\mathcal{O}(||\langle f,M\rangle||)$.

Thanks to Eq. (6) and the definition of Γ, the quantities u_n, $\cos(n\theta+\varphi)$, and $\cos(\Gamma)$ have the same sign. Since we are assuming that $0<\Gamma\leq\pi$, it follows that $\cos(\Gamma)$ and Λ have the same sign as well. By Eq. (7), the sign of Λ (and therefore that of u_n) can be determined by computing $||\langle f,M\rangle||^{\mathcal{O}(1)}\log_2 n = ||\langle f,M,n\rangle||^{\mathcal{O}(1)}$ bits of Λ, which can be done in polynomial time thanks to Prop. 2, by noting that for any algebraic $\alpha\in\mathbb{C}\setminus\{i,-i\}$ of modulus 1, $\log\alpha$ can be obtained by computing $\arctan(\mathrm{Im}(\alpha)/\mathrm{Re}(\alpha))$.

This concludes the proof of Thm. 4. ◀

4 The Three-Dimensional Matrix Powering Positivity Problem

We now move to three dimensions. We are given a linear function $f:\mathbb{Z}^9\to\mathbb{Z}$, a 3×3 integer matrix M, and an integer $n\geq 0$. We assume that the base matrix M is encoded in unary, whereas the function f and exponent n are encoded in binary. Note in particular that this includes the important special case in which the base data $\langle f,M\rangle$ is fixed.

Our main result is as follows:

▶ **Theorem 5.** *In three dimensions,* PosMatPow *(with unary encoding of the base matrix and binary encoding of the linear function and of the exponent) is decidable in polynomial time.*

Proof. Let $\langle f,M,n\rangle$ be as above. We seek to determine whether $f(M^n)\geq 0$.

As before, write $u_k = f(M^k)$. Our strategy is to exhibit a bound N, of *magnitude* polynomial in $||\langle f,M\rangle||$, such that the sign of u_k is easily determined for $k\geq N$. Note on the other hand that, if $n<N$, then one can simply compute u_n outright in polynomial time.

We split our analysis into two main cases: (i) either M only has real eigenvalues, or (ii) two of M's eigenvalues are complex conjugates.

In Case (i), let ρ_1 be a real dominant eigenvalue of M. We focus on the hardest instance, in which ρ_1 has multiplicity 1; the other two alternatives are considerably simpler and can be handled similarly.

Let ρ_2 be a second real eigenvalue of M. Here, the critical case is when ρ_2 is repeated; the (easier) alternative can again be handled in similar fashion, and is therefore omitted.

We can thus write

$$u_k = a\rho_1^k + (bk + c)\rho_2^k \,, \tag{8}$$

where we assume that $a \neq 0$. Observe that since M is encoded in unary, one has an $||M||^3$ upper bound for the maximum height H of the eigenvalues ρ_1 and ρ_2. Likewise, the real algebraic numbers $1/a$, b, and c all have representations of size linear in $||\langle f, M\rangle||$, and therefore have magnitude at most $2^{\mathcal{O}(||\langle f,M\rangle||)}$.

Note that if ρ_2 and ρ_1 have the same modulus, then $\rho_2 = -\rho_1$ and the treatment is straightforward; we therefore assume that $|\rho_2| < |\rho_1|$. Since both ρ_1 and ρ_2 have degree at most 3, Mignotte's root-separation bound (Eq. 1) entails that $|\rho_2| = |\rho_1| - \delta$, where $\delta = \Omega(H^{-2})$.

Let $\gamma = \rho_2/\rho_1$. Equation (8) can be rewritten as

$$\frac{u_k}{\rho_1^k} = a + (bk + c)\gamma^k \,. \tag{9}$$

We also have

$$|\gamma| = \frac{|\rho_2|}{|\rho_1|} = \frac{|\rho_1| - \delta}{|\rho_1|} = 1 - \frac{\delta}{|\rho_1|} \,.$$

By Eq. (2), $|\rho_1| \leq H + 1$, from which we immediately conclude that $|\gamma| = 1 - \varepsilon$, where $1/\varepsilon = \mathcal{O}(H^3) = \mathcal{O}(||M||^9)$. We now aim to establish a bound N of magnitude polynomial in $||\langle f, M\rangle||$ such that, for $k \geq N$,

$$|a| > \left|(bk + c)\gamma^k\right| \,. \tag{10}$$

By Eq. (9), the sign of u_k is then automatically obtained for any k beyond N.

From the inequality $\log(1 - \varepsilon) < -\varepsilon$, we have $|\gamma| = 1 - \varepsilon < e^{-\varepsilon}$. In order for Eq. (10) to hold, it therefore suffices to have $|a| > |bk + c|e^{-k\varepsilon}$. Letting $x = k\varepsilon$, this translates to $e^x > |(bx/\varepsilon + c)/a|$. Thanks to our bounds on $|1/a|$, $|b|$, $|c|$, and $1/\varepsilon$, we can find $B = 2^{\mathcal{O}(||\langle f,M\rangle||)}$ such that $|(bx/\varepsilon + c)/a| \leq Bx$ for all $x \geq 1$. Now clearly the inequality $e^x > Bx$ holds provided that $x \geq 2\log B$, or equivalently that $k \geq (2\log B)/\varepsilon$. Letting $N = \lceil (2\log B)/\varepsilon \rceil = ||\langle f, M\rangle||^{\mathcal{O}(1)}$ and putting everything together, we see that Eq. (10) holds for $k \geq N$, as required.

We now turn to Case (ii), in which M has two complex conjugate eigenvalues λ and $\overline{\lambda}$, and one real eigenvalue ρ. For all $k \geq 0$, we have

$$u_k = a\rho^k + c\lambda^k + \overline{c}\overline{\lambda}^k \,, \tag{11}$$

for some algebraic constants $a \in \mathbb{R}$ and $c \in \mathbb{C}$. As before, ρ and λ have height bounded by $||M||^3$, whereas a and c have height bounded by $2^{\mathcal{O}(||\langle f,M\rangle||)}$. Moreover λ has degree at most 3 and $\rho, a, c \in \mathbb{Q}(\lambda, \overline{\lambda})$.

If $|\rho| > |\lambda|$, we can proceed straightforwardly through a growth argument akin to that invoked in Case (i) above, whereas if $|\rho| = |\lambda|$, the situation is very similar to Case (ii) of the two-dimensional instance of the problem, handled in the previous section, and can be dealt with in like fashion. We therefore focus on the situation in which $|\rho| < |\lambda|$.

Let $\theta = \arg\lambda$ and $\varphi = \arg c$. Equation (11) then becomes

$$u_k = a\rho^k + |c||\lambda|^k \cos(k\theta + \varphi)\,.$$

Writing $\gamma = \rho/|\lambda|$, we have

$$\frac{u_k}{|c||\lambda|^k} = \frac{a}{|c|}\gamma^k + \cos(k\theta + \varphi)\,, \tag{12}$$

and as before, $|\gamma| = 1 - \varepsilon$, where $1/\varepsilon = \mathcal{O}(||M||^9)$.

As in the two-dimensional case, we can check in polynomial time whether $\cos(n\theta+\varphi) = 0$, in which case the sign of u_n is readily determined. Otherwise, write $\Gamma = n\theta + \varphi - 2m\pi$, with $m \in \mathbb{Z}$ such that $-\pi < \Gamma \leq \pi$, and as before, without loss of generality, assume that Γ is 'close' to $\pi/2$, the other cases being either straightforward or handled similarly. Write $\Lambda = \pi/2 - \Gamma$, and note that

$$\cos(n\theta + \varphi) = \cos\Gamma = \sin\Lambda \qquad \text{and} \qquad |\sin\Lambda| > \frac{|\Lambda|}{2}\,. \tag{13}$$

We have

$$\Lambda = \frac{1}{i}\left(n\log\frac{\lambda}{|\lambda|} + \log\frac{c}{|c|} + (1-4m)\log i\right)\,.$$

Since λ has degree at most 3, the degree of $\mathbb{Q}(\lambda, \overline{\lambda}, |\lambda|, c, |c|, i)$ over $\mathbb{Q}$ is at most 48. Moreover, we can bound the height of $\lambda/|\lambda|$ and $c/|c|$ by some H with $\log H = \mathcal{O}(||\langle f, M\rangle||)$. Finally, we note that $|1 - 4m| \leq 2n + 1$. Applying Baker's theorem, we get

$$|\Lambda| > \exp\left(-2304^{10}(\log H)^2\log(2n+1)\right) = \frac{1}{(2n+1)^{2304^{10}(\log H)^2}} = \frac{1}{n^{||\langle f,M\rangle||^{\mathcal{O}(1)}}}\,. \tag{14}$$

It follows from Eqs. (13) and (14) that there is an absolute constant $T \in \mathbb{N}$ such that

$$|\cos(n\theta + \varphi)| > \frac{1}{n^{||\langle f,M\rangle||^T}}\,. \tag{15}$$

We now aim to establish a bound N of magnitude polynomial in $||\langle f, M\rangle||$ such that, if $n \geq N$, then

$$\left|\frac{a}{|c|}\gamma^n\right| < |\cos(n\theta + \varphi)|\,. \tag{16}$$

Thanks to Eq. (12), in that case the sign of u_n is the same as that of $\cos(n\theta + \varphi)$, and in turn the latter can be determined in polynomial time following the procedure outlined in Case (ii) of the two-dimensional instance of the problem, thanks to Eq. (14). On the other hand, if $n < N$, we simply note that u_n can then be computed outright in polynomial time.

By Eq. (15), and recalling that $|\gamma| = 1 - \varepsilon$, it is sufficient for Eq. (16) to hold to have

$$\frac{|a|}{|c|}(1-\varepsilon)^n < \frac{1}{n^{||\langle f,M\rangle||^T}}\,,$$

or equivalently (noting that $\log(1-\varepsilon) < 0$),

$$n > -\frac{||\langle f, M\rangle||^T}{\log(1-\varepsilon)}\log n - \frac{\log(|a|/|c|)}{\log(1-\varepsilon)}\,.$$

Multiplying the above equation by 2 and writing $2n = n + n$, it is then sufficient for both

$$n > -2\frac{||\langle f, M\rangle||^T}{\log(1-\varepsilon)} \log n \qquad \text{and} \tag{17}$$

$$n > -2\frac{\log(|a|/|c|)}{\log(1-\varepsilon)} \tag{18}$$

to hold.

For any $Q \geq 1$, one has $Q > 2\log Q$, thus $Q^2 > Q\log(Q^2)$. In other words, $x > Q\log x$ for $x = Q^2$. But by comparing derivatives at the point $x = Q^2$, we see that the inequality $x > Q\log x$ holds for all $x \geq Q^2$. Writing

$$Q = -2\frac{||\langle f, M\rangle||^T}{\log(1-\varepsilon)}\,,$$

we see that Eq. (17) holds provided $n \geq Q^2$. Since $1/\varepsilon = \mathcal{O}(||M||^9)$ and $|\log(1-\varepsilon)| > \varepsilon$, we immediately have $Q^2 = ||\langle f, M\rangle||^{\mathcal{O}(1)}$.

Next, let H be the maximum of the heights of a and c, noting that $\log H = \mathcal{O}(||\langle f, M\rangle||)$. By Eq. (2), $|a| < H + 1$ and $|c| > 1/(H+1)$, whence

$$Q' = \left|-2\frac{\log(|a|/|c|)}{\log(1-\varepsilon)}\right| = ||\langle f, M\rangle||^{\mathcal{O}(1)}\,.$$

It follows that by letting $N = \max\{\lceil Q^2\rceil, \lceil Q'\rceil\}$, both Eqs. (17) and (18) hold provided that $n \geq N$, as required.

This concludes the proof of Thm. 5. ◀

5 Concluding Remarks

It is worth noting that our results can be extended in a fairly minor way, by considering matrices M and linear functions f with *rational* entries and coefficients: indeed, the rational formulation of PosMatPow reduces straightforwardly to its integer counterpart at the cost of a polynomial blowup in size.

Further extensions however appear elusive under the present framework. In the three-dimensional case, for instance, encoding the base matrix in binary would not yield a sufficiently large spectral gap (difference in magnitude between the largest and second-largest eigenvalues) for our present approach to go through; more specifically, the value of N required so that Eq. (10) hold would then potentially be exponential, thereby not leading to a polynomial-time algorithm. In four dimensions or higher, the situation worsens: we do not know how to produce a polynomial-time algorithm even for *fixed* base data M and f. A critical case is encountered when there are four or more dominant complex eigenvalues, ostensibly precluding the use of Baker's theorem.

The reader will have noticed the presence of various 'galactic' constants appearing in the analysis of our algorithms, and perhaps conclude that the approach we have laid out is unlikely to be feasible in practice. It is worth noting, however, that our analysis merely serves to establish (large) polynomial-time upper bounds, without any expectation that such bounds need be tight. On the contrary, we conjecture that the proposed approach, under careful implementation and engineering, would prove quite efficient in practice. Substantiating this empirically might however be expected to require non-trivial efforts.

References

1 E. Allender, N. Balaji, and S. Datta. Low-depth uniform threshold circuits and the bit-complexity of straight-line programs. *Elec. Coll. on Comput. Complex.*, 177(1), 2013.

2 E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. *SIAM J. Comput.*, 38(5), 2009.

3 A. Baker and G. Wüstholz. Logarithmic forms and group varieties. *Jour. Reine Angew. Math.*, 442, 1993.

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

5 R. P. Brent. Fast multiple-precision evaluation of elementary functions. *J. ACM*, 23(2), 1976.

6 J.-Y. Cai, R. J. Lipton, and Y. Zalcstein. The complexity of the A B C problem. *SIAM J. Comput.*, 29(6), 2000.

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

8 K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points (extended abstract). In *Proceedings of FOCS*. IEEE Computer Society, 2007.

9 G. Everest, A. van der Poorten, I. Shparlinski, and T. Ward. *Recurrence Sequences*. American Mathematical Society, 2003.

10 H. Galperin and A. Wigderson. Succinct representations of graphs. *Inf. Control*, 56(3), 1983.

11 G. Ge. *Algorithms Related to Multiplicative Representations of Algebraic Numbers*. PhD thesis, U.C. Berkeley, 1993.

12 M. Hirvensalo, J. Karhumäki, and A. Rabinovich. Computing partial information out of intractable: Powers of algebraic numbers as an example. *Jour. Number Theory*, 130, 2010.

13 R. J. Lipton. A challenge from Dyson. *Blog entry*, September 2014. `http://rjlipton.wordpress.com/2014/09/09/a-challenge-from-dyson/`.

14 D. W. Masser. Linear relations on algebraic groups. In *New Advances in Transcendence Theory*. Cambridge University Press, 1988.

15 C. Mereghetti and B. Palano. Threshold circuits for iterated matrix product and powering. *Theoret. Informatics Appl.*, 34(1), 2000.

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

17 J. Ouaknine and J. Worrell. On the positivity problem for simple linear recurrence sequences. In *Proceedings of ICALP*, number 8573 in Springer LNCS, 2014.

18 J. Ouaknine and J. Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of SODA*. SIAM, 2014.

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

20 C. H. Papadimitriou and M. Yannakakis. A note on succinct representations of graphs. *Inf. Control*, 71(3), 1986.

21 Q. I. Rahman and G. Schmeisser. *Analytic Theory of Polynomials*. London Mathematical Society monographs. Oxford University Press, 2002.