Feasibility of Circuit Polynomials without Purple Swans

Weixun Deng*
deng15521037237@tamu.edu
Texas A&M University
USA

Alperen Ergür†
alperen.ergur@utsa.edu
University of Texas at San Antonio
USA

Grigoris Paouris‡
grigoris@tamu.edu
Texas A&M University
USA

J. Maurice Rojas§
rojas@tamu.edu
Texas A&M University
USA

ABSTRACT

Suppose f is a polynomial in n variables with degree d, exactly $n+k$ monomial terms, coefficients in $\{\pm 1, \ldots, \pm H\}$ for some $H \in \mathbb{N}$, and Newton polytope of positive volume. Testing real feasibility of such an f is a fundamental task whose bit-complexity remains a mystery, even in the first non-trivial case $k = 2$: The fastest algorithms so far have deterministic bit-complexity $(n \log(dH))^{O(n)}$. We prove a significant speed-up that holds for all but a small collection of inputs in the $k = 2$ case: Bit complexity $(n \log(dH))^{O(1)}$ for all but a $O\left(\frac{1}{2^n H}\right)$-fraction of the f above, for any fixed support. Our result follows by combining a connection to diophantine approximation with a more recent anti-concentration result. In particular, we show that for random inputs, Baker's famous theorem on linear forms in logarithms can be significantly sharpened. We also consider extensions beyond feasibility such as counting connected components and systems of circuit polynomials.

CCS CONCEPTS

• **Mathematics of computing** → **Nonlinear equations**; **Number-theoretic computations**.

KEYWORDS

circuit, sparse polynomial, isotopy type, positive, feasibility, connected component, Baker's Theorem, linear forms in logarithms, polynomial-time, discrete Gaussian, random

ACM Reference Format:
Weixun Deng, Alperen Ergür, Grigoris Paouris‡, and J. Maurice Rojas§. 2024. Feasibility of Circuit Polynomials without Purple Swans. In *International Symposium on Symbolic and Algebraic Computation (ISSAC '24), July 16–19, 2024, Raleigh, NC, USA.* ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3666000.3669716

* ‡ § Partially supported by NSF grant CCF-1900881.
†Partially supported by NSF grant CCF-2110075.

ISSAC '24, July 16–19, 2024, Raleigh, NC, USA

ACM ISBN 979-8-4007-0696-7/24/07
https://doi.org/10.1145/3666000.3669716

1 INTRODUCTION AND MAIN RESULTS

Counting the number of connected components (a.k.a. *pieces*) for the positive zero set, $Z_+(f)$, of a Laurent polynomial f, as a function of its monomial term structure, is a fundamental problem from real algebraic geometry that is still far from completely understood. This is unfortunate, because many real zero sets occuring in practice come from highly structured polynomials, and one of the most basic structures to consider is monomial term structure.

Definition 1.1. *Suppose $f \in \mathbb{Z}\left[x_1^{\pm 1}, \ldots, x_n^{\pm 1}\right]$ is of the form $f(x) = \sum_{i=1}^{n+k} c_i x^{a_i}$, where $c_i \in \{\pm 1, \ldots, \pm H\}$, $a_i \in \{-d, \ldots, d\}^{n\times 1}$, and $x^{a_i} := x_1^{a_{1,i}} \cdots x_n^{a_{n,i}}$ for all i. Assuming in addition that $A := \{a_1, \ldots, a_{n+k}\}$ has cardinality $n+k$, we call such an f an* n-variate $(n+k)$-nomial of type (A, d, H), *and say that f is* supported *on A. If we also have that the matrix $\widehat{\mathcal{A}} := \begin{bmatrix} 1 & \cdots & 1 \\ a_1 & \cdots & a_{n+k} \end{bmatrix} \in \mathbb{Z}^{(n+1)\times(n+k)}$ has rank $n+1$ then we call f an* honest *n-variate $(n+k)$-nomial.* ⋄

Remark 1.2. *The geometric restriction on the exponent vectors (via the rank of $\widehat{\mathcal{A}}$ above) makes the parameter n meaningful: Without this restriction, one could find a simple change of variables to reduce to a smaller n while still preserving $n+k$, e.g., the positive roots of $1 - xy + x^{100}y^{100}$ can be determined from the positive roots of $1 - u + u^{100}$ by substituting $u = xy$. Note also that the rank restriction on $\widehat{\mathcal{A}}$ forces $k \geq 1$.* ⋄

For the special case $n = 1$, Descartes' Rule tells us that the number of pieces (for the positive zero set of a univariate $(k+1)$-nomial) is at most k, and this bound is tight thanks to the explicit family of examples $(x_1 - 1)(x_1 - 2)\cdots(x_1 - k)$. However, for $n = 2$, it isn't even known if the number of pieces admits an upper bound of the form $k^{O(1)}$: The best upper bound is still exponential in k^2 [Kho91, BS09], and no family of examples evincing even $\Omega(k^2)$ pieces is known. However, recent *probabilistic* results in real fewnomial theory [BET-C19] suggest that, *on average* (for many natural coefficient distributions), the number of pieces should be $O(k^2)$.

Remark 1.3. *All O-, Ω-, and o-constants in our results are effective and absolute, i.e., they are actual* constants *that can be made explicit, albeit with some effort. Also, for an f as in Definition 1.1, we define the* size *of f to be*

$$\sum_{i=1}^{n+k} \left(\left\lceil \log_2(2 + |c_i|) \right\rceil + \sum_{j=1}^{n} \left\lceil \log_2(2 + |a_{i,j}|) \right\rceil \right),$$

i.e., the sum of the bit sizes of the coefficients and the exponents of f. So in our setting, polynomial-time algorithms have bit complexity

$((n+k)\log(dH))^{O(1)}$, *as opposed to many basic algorithms in classical computational algebra that have complexity* $(dn\log H)^{O(n)}$. ⋄

For the *algorithmic* question of actually counting the number of pieces for a given $Z_+(f)$, our knowledge is even sparser: A polynomial-time algorithm is known only for the cases $(n,k)\in\{(1,1),(1,2)\}$ [BRS09, Bih11]. There is also the folkloric fact that the case $k=1$ and n arbitrary never results in $Z_+(f)$ having more than 1 piece, and (for $k=1$) deciding between 0 and 1 pieces is doable in time $O(n)$ (see, e.g., Lemma 2.14 in Section 2.1 below). On the other hand, we have **NP**-hardness (for deciding non-emptiness, a.k.a. *feasibility*) if we fix any $\varepsilon>0$, let $n\longrightarrow\infty$, and take $k=n^\varepsilon$ [BRS09]. What happens between $k=1$ and $k=n^\varepsilon$ is still a mystery.

In particular, the fastest algorithms for just deciding if there are any pieces at all (for $k=2$ and n arbitrary) have deterministic bit-complexity $(n\log(dH))^{O(n)}$ [BRS09]. (See also [BPR06, BR14] for much more powerful and general algorithms which, unfortunately, are no faster in the $k=2$ case.) So we prove the following significant speed-up for the case $k=2$:

Theorem 1.4. *Following the notation above, for a fraction of* $1-O\left(\frac{1}{2^nH}\right)$ *of honest n-variate* $(n+2)$*-nomials of type* (A,d,H)*, we can decide whether* $Z_+(f)$ *is empty in deterministic time* $O(n^{3.373}\log^3(ndH))$.

Theorem 1.4 is proved in Section 3. There, we will also see that there are in fact A for which the fraction of inputs on which we can go faster is 1 (i.e., all inputs supported on A). The first step of the proof is to reduce our feasibility question to a diophantine problem: Determining the sign of an integer linear combination of logarithms of integers. This reduction is straightforward, after we review a special case of the $\mathcal{A}$-discriminant [GKZ94] — the *circuit* case — in Section 2. So the main difficulty is understanding the subtle behavior of linear forms in logarithms.

Remark 1.5. *The circuit discriminant has also been used recently in other computational settings such as chemical reaction networks [PKC22], entropy computation and signomial programming [CS16], and polynomial optimization [PRT09, MSdW19]. For instance, if one wants to check if the minimum of a (real) n-variate* $(n+1)$*-nomial over* $\mathbb{R}^n_+$ *is greater than, equal to, or less than an input rational number, then this query reduces (in the most difficult case) to deciding the sign of a circuit discriminant.* ⋄

By combining with a more refined classification of isotopy types from [BDPRRR24], we can even *count* the number of pieces of $Z_+(f)$ within the same time bound as Theorem 1.4, and this addendum is currently being finalized. However, the underlying probabilistic technique is the same for both results, so we cover it now.

Remark 1.6. *There is a very natural sibling to the problem of counting connected components of a hypersurface in* $\mathbb{R}^n_+$ *defined by a single circuit polynomial: Counting the number of isolated roots in* $\mathbb{R}^n_+$ *of an* $n\times n$ system *of circuit polynomials. This problem was recently given a deterministic speed-up in [Roj24], but the complexity bound there is still exponential in n, similar to the hypersurface case. However, via our techniques here, it appears that we can also speed up root counting for circuit* systems *to time polynomial in n for "most" inputs. There still remain a number of technical hurdles, one of which is the need to extend probabilistic estimates for linear combinations of logarithms of random rational numbers to random* real algebraic *numbers. So we lay the groundwork for the latter problem in Section 1.3 below.* ⋄

1.1 Probabilistic Bounds on Linear Forms in Logarithms

A landmark 1966 result in transcendental number theory due to Baker can be coarsely summarized as follows [Bak77]: Let $\Lambda(b,\xi):=\sum_{i=1}^m b_i\log\xi_i$ be an integer linear combination of logarithms of rational numbers. Let H denote the maximal absolute value among the integers that appear in the numerators and denominators of the ξ_i, and let $B:=\max_i|b_i|$. A special case of *Baker's Theorem on Linear Forms in Logarithms* [Bak77, BW93, Mat00, Nes03] then implies

$$\Lambda(b,\xi)\neq 0\Rightarrow\log|\Lambda(b,\xi)|>-O(\log H)^m\log B. \tag{1}$$

Baker's Theorem was remarkably difficult to prove, and the special case $m=2$ already implies the solution to Hilbert's Seventh Problem (proving that a^b is transcendental for $a\notin\{0,1\}$ algebraic and b algebraic and irrational).

Baker won a Fields Medal in 1970 for his lower bound, and later his bound also proved useful for many other important problems in number theory, e.g., computing explicit upper bounds for the size of integer points on curves of genus 1 (see, e.g., [Sch92]). More recently, Baker's lower bound has also found use in the design and analysis of algorithms for real algebraic geometry and parsing (see, e.g., [BRS09, BHPR11, BSY14, Roj24]).

It has been conjectured that Baker's lower bound is far from sharp: Lang and Waldschmidt used a simple heuristic argument to motivate a conjecture that the optimal bound should be $-O(m\log(HB))$ [Lan78, Pg. 213]. However, there appears to have been no progress whatsoever, for close to half a century, on their conjecture. Our algorithmic goals happen to naturally motivate a probabilistic approach to this bound: Can we prove a sharper version of Baker's lower bound, for *most* inputs instead, and thereby prove that earlier algorithms for real feasibility can be sped up most of the time?

Remark 1.7. *The title of our paper was inspired by the title of [AL17], which studies algorithms that are fast on average, outside of a small region of inputs. In contrast, we study a setting where* worst-case *complexity is low for all inputs outside of a small region.* ⋄

1.2 Two Models of Discrete Randomness

We start by stating an important consequence of Corollary 1.4 from an elegant paper of Rudelson and Vershynin [RV15]. We let $\langle u,v\rangle$ denote the standard dot product of $u,v\in\mathbb{R}^n$.

Lemma 1.8. *Consider a random vector* $X=(X_1,\ldots,X_m)$ *where the* X_i *are independent random variables. Let* $p,t>0$ *be parameters such that* $\sup_{z\in\mathbb{R}}\mathbb{P}\{|X_i-z|\leq t\}\leq p$ *for all* $i\in\{1,\ldots,m\}$*. Then for all* $b\in\mathbb{R}^m$ *we have*

$$\sup_{z\in\mathbb{R}}\mathbb{P}\{|\langle X,b\rangle-z|\leq t|b|_2\}\leq\sqrt{2}p. \qquad \blacksquare$$

Using the special case $d=1$ of [RV15, Cor. 1.4] and rescaling to allow an inner product with an arbitrary $b\in\mathbb{R}^m$, we obtain an upper bound of $O(p)$ for the right-hand side of the inequality above. Applying [LPP16] then immediately yields the refinement to $\sqrt{2}p$.

1.2.1 Random Integers with Controlled Bit-Size.

Proposition 1.9. *Let $H \in \mathbb{N}$ and consider a uniformly random α chosen from $\{\pm 1, \ldots, \pm H\}$. Then, for any $z \in \mathbb{R}$ and any $\varepsilon \in \left(\frac{1}{H}, \frac{1}{4}\right)$ we have:*

$$\mathbb{P}\{|\log|\alpha| - z| \leq \varepsilon\} \leq 9\varepsilon.$$

Proof: Let us use $|S|$ for the cardinality of a set S. Note that for any interval $[s,t]$ we have that

$\mathbb{P}\{\log|\alpha| \in [s,t]\} = \mathbb{P}\{|\alpha| \in [e^s, e^t]\}$, which is in turn bounded from above by $\frac{2\left|[e^s,e^t]\cap[1,\ldots,H]\right|}{2H}$. So if $e^{z-\varepsilon} > H$ then we have $\mathbb{P}\{|\log|\alpha| - z| \leq \varepsilon\} = 0$. If $e^{z-\varepsilon} \leq H$ then we have

$$\mathbb{P}\{|\log|\alpha| - z| \leq \varepsilon\} \leq \frac{2\left[e^{z+\varepsilon} - e^{z-\varepsilon} + 1\right]}{2H} = \frac{1}{H} + 2\left(e^{2\varepsilon} - 1\right)\frac{e^{z-\varepsilon}}{2H}.$$

For any $\varepsilon < \frac{1}{4}$ Taylor's expansion yields $e^{2\varepsilon} - 1 \leq \frac{1}{1-2\varepsilon} - 1 = \frac{2\varepsilon}{1-2\varepsilon} \leq 4\varepsilon$. So we have

$$\mathbb{P}\{|\log\alpha - z| \leq \varepsilon\} \leq \frac{1}{H} + 8\varepsilon\left(\frac{e^{z-\varepsilon}}{2H}\right) \leq 9\varepsilon. \qquad \blacksquare$$

Remark 1.10. *The probability distribution in Proposition 1.9 was chosen for simplicity. Similar estimates can be easily derived for more general distributions, e.g., the uniform distribution on $\{x-H, \ldots, -1, 1, x, \ldots, x+H\}$ for $0 \leq x \leq H$.* $\diamond$

Combining Lemma 1.8 and Proposition 1.9 gives the following probabilistic estimate on integer linear sums of logs.

Corollary 1.11. *Fix any $b := (b_1, \ldots, b_m) \in (\mathbb{Z} \setminus \{0\})^m$ and let $\alpha = (\alpha_1, \ldots, \alpha_m)$ be a uniformly random vector in $\{\pm 1, \ldots, \pm H\}^m$. Then for any $z \in \mathbb{R}$ and $\varepsilon \in \left(\frac{1}{H}, \frac{1}{4}\right)$ we have:*

$$\mathbb{P}\{|b_1\log|\alpha_1| + \cdots + b_m\log|\alpha_m| - z| \leq \varepsilon|b|_2\} \leq 9\sqrt{2}\varepsilon$$

Since $b_i^2 \geq 1$ for all i, this also yields

$$\mathbb{P}\{|b_1\log|\alpha_1| + \cdots + b_m\log|\alpha_m| - z| \leq \varepsilon\sqrt{m}\} \leq 9\sqrt{2}\varepsilon. \qquad \blacksquare$$

1.2.2 Discrete Gaussians.

A distribution that is commonly used in discrepancy theory and integer programming applications is the discrete Gaussian (see, e.g., [ADRSD15]).

Here we will consider discrete Gaussian centered at an arbitrary integer $a \in \mathbb{Z}$ with standard deviation H. More precisely, for any $x \in \mathbb{Z}$ we define $p(x) := e^{-\frac{(x-a)^2}{2H^2}}$ and set $Q := \sum_{x\in\mathbb{Z}} e^{-\frac{(x-a)^2}{2H^2}}$. Then the discrete Gaussian centered at a with standard deviation H is the random variable X that takes integer values with the following weights: $\mathbb{P}(X = x) := \frac{p(x)}{Q}$.

First we make a quick computation:

$$1 + Q = 2\sum_{y=0}^{\infty} e^{-\frac{y^2}{2H^2}} \geq 2\int_0^{\infty} e^{-\frac{y^2}{2H^2}}\,dy = \int_{-\infty}^{\infty} e^{-\frac{y^2}{2H^2}}\,dy = H\sqrt{2\pi}$$

So $Q > H\sqrt{2\pi} - 1$, and for any $H > 2$, this also yields $Q > 2H$.

Lemma 1.12. *Let a be an arbitrary integer and let X be the discrete Gaussian centered at a with standard deviation H where $H > |a|$. Then, for any real number z and any $\frac{1}{4} > \varepsilon > \frac{1}{H}$ we have*

$$\mathbb{P}\{|\log|X| - z| \leq \varepsilon\} \leq 9\varepsilon.$$

Proof: Note that for any interval $[s,t]$ we have

$$\mathbb{P}\{\log|X| \in [s,t]\} = \mathbb{P}\{|X| \in [e^s, e^t]\} \leq \frac{2 \cdot \sum_{y\in[e^s,e^t]\cap\mathbb{Z}} e^{-\frac{(y-a)^2}{2H^2}}}{Q}$$

So we have

$$\mathbb{P}\{|\log|X| - z| \leq \varepsilon\} \leq \frac{\sum_{y\in[e^{z-\varepsilon},e^{z+\varepsilon}]\cap\mathbb{Z}} e^{-\frac{(y-a)^2}{2H^2}}}{H}$$

(Note that in both events above, the event $X = 0$ would imply $\log|X| = -\infty$, so the latter event is excluded.)

For any $\varepsilon < \frac{1}{4}$ Taylor's expansion yields $e^{2\varepsilon} - 1 \leq \frac{1}{1-2\varepsilon} - 1 = \frac{2\varepsilon}{1-2\varepsilon} \leq 4\varepsilon$. So we have $e^{z+\varepsilon} - e^{z-\varepsilon} \leq 4\varepsilon\, e^{z-\varepsilon}$. If $H \geq a \geq e^{z-\varepsilon}$ then we have

$$\mathbb{P}\{|\log|X| - z| \leq \varepsilon\} \leq \frac{4\varepsilon e^{z-\varepsilon} + 1}{H} \leq 9\varepsilon$$

If $a < e^{z-\varepsilon}$ then let $d := \min_{y\in[e^{z-\varepsilon},e^{z+\varepsilon}]\cap\mathbb{Z}} |y - a|$. Note that $a + d > e^{z-\varepsilon}$, and thus we have $e^{z+\varepsilon} - e^{z-\varepsilon} \leq 4\varepsilon(a+d)$. So we have

$$\mathbb{P}\{|\log|X| - z| \leq \varepsilon\} \leq \frac{(4\varepsilon(a+d)+1)e^{-\frac{d^2}{2H^2}}}{H} \leq \varepsilon + 4\varepsilon\frac{(a+d)e^{-\frac{d^2}{2H^2}}}{H}$$

If $d < H$ then we are done. Otherwise, let $\delta = \frac{d}{H} \geq 1$ and note that

$$\mathbb{P}\{|\log|X| - z| \leq \varepsilon\} \leq \left(1 + 8\delta e^{-\frac{\delta^2}{2}}\right)\varepsilon \leq 9\varepsilon \qquad \blacksquare$$

Combining Lemmata 1.8 and 1.12 immediately gives the following probabilistic estimate on integer linear sums of logs.

Corollary 1.13. *Let $b = (b_1, \ldots, b_m) \in \mathbb{Z}^m$ be an integer vector with $b_i \neq 0$. Let $\alpha = (\alpha_1, \ldots, \alpha_m)$ be a random vector where the α_i are independent discrete Gaussian random variables centered at integers x_i with variances H_i^2 where $H_i > |x_i|$. Let $H := \min_{1\leq i\leq m} H_i$, then for any $z \in \mathbb{R}$ and $\varepsilon \in \left(\frac{1}{H}, \frac{1}{4}\right)$ we have*

$$\mathbb{P}\{|b_1\log|\alpha_1| + \cdots + b_m\log|\alpha_m| - z| \leq \varepsilon|b|_2\} = O(\varepsilon).$$

Since $b_i^2 \geq 1$ for all i, this also yields

$$\mathbb{P}\{|b_1\log|\alpha_1| + \cdots + b_m\log|\alpha_m| - z| \leq \varepsilon\sqrt{m}\} = O(\varepsilon). \qquad \blacksquare$$

Remark 1.14. *We emphasize, per Remark 1.3, that the O-constants above are truly constants, i.e., there is no dependence on any additional parameters.* $\diamond$

1.3 Random Algebraic Integers

As noted in Remark 1.6, it appears that our approach to new probabilistic speed-ups can be extended to counting isolated roots in $\mathbb{R}_+^n$ of $n \times n$ *systems* of circuit polynomials (all with the same support). Technically, counting pieces in our setting here is accomplished by using signs of linear combinations of logarithms of rational numbers to decide which discriminant chamber contains f. (This is explained further in Section 2.) To count real solutions of circuit systems instead, there is a reduction (using *Gale Dual form* [BS07]) to counting real roots of a linear combination of logarithms of *degree one polynomials*. While the latter problem appears to be purely transcendental, one can reduce it (as analyzed in [Roj24]) to computing several signs of linear forms of logarithms of *real algebraic numbers*. So we can extend our approach to systems provided we

have sufficiently strong extensions of Corollaries 1.11 and 1.13 to real algebraic numbers. (There are additional steps to this program, but we leave the details for future work.) This motivates Corollary 1.21 below as a first step toward the harder problem of counting real roots of circuit systems.

Suppose we are given $\alpha \in \mathbb{C}$ that generates a degree d field extension $\mathbb{Q}(\alpha)$ of $\mathbb{Q}$. To generate random algebraic integers from the number field $\mathbb{Q}(\alpha)$ that have (absolute logarithmic) height at most h, let us first recall some basics (see, e.g., [BG06] for further background).

Definition 1.15. *Let α be an algebraic number with minimal polynomial $\gamma_0 + \cdots + \gamma_d x^d \in \mathbb{Z}[x]$ satisfying $\gcd(\gamma_0, \ldots, \gamma_d) = 1$, and let $\alpha_1, \ldots, \alpha_d$ be all the conjugates of α. Then the* absolute multiplicative height *of α, denoted by $H(\alpha)$ is defined as*

$$H(\alpha) := \left(|\gamma_d| \prod_{i=1}^{d} \max\{1, \alpha_i\} \right)^{\frac{1}{d}}$$

The absolute logarithmic height *is defined as $h(\alpha) := \log H(\alpha)$.* $\diamond$

Henceforth, we will mean *absolute logarithmic* height when we speak of the height of an algebraic number.

Lemma 1.16. *Let $u = \xi_0 + \cdots + \xi_{d-1}\alpha^{d-1} \in \mathbb{Q}[\alpha]$ where $\mathbb{Q}(\alpha)$ is a degree d extension of $\mathbb{Q}$. Then*

$$h(u) \leq d \max_i h(\xi_i) + \frac{d(d-1)}{2} h(\alpha) + \log d.$$

Proof: By [BG06, Sec. 1.5.14, Pg. 18] the height of a product of algebraic numbers is the sum of their respective heights. So $h(\xi_i \alpha^i) = h(\xi_i) + ih(\alpha)$. By the standard upper bound on the height of a sum (see, e.g., [BG06, Sec. 1.5.16, Pg. 19]), we then obtain the stated bound. ■

Now we consider the following model of randomness: We generate $u = \xi_0 + \cdots + \xi_{d-1}\alpha^{d-1}$ by assuming that ξ_i is an independent discrete Gaussian random variable with mean zero and variance H_i^2 where $H_i \geq 1$. Let

$$\eta := \max\{\log H_0, \ldots, \log H_{d-1}, h(\alpha)\}$$

Lemma 1.17. *With the notation above, we have:*

$$\mathbb{P}\left(h(u) \leq 5d^2\eta + 2\log d \right) \geq \frac{1}{2} \tag{2}$$

Proof. Using Lemma 1.16, and Markov's inequality, we have for any $t > 0$ that

$$\begin{aligned} \mathbb{P}\left(h(u) \geq t\right) &\leq \mathbb{P}\left(d \max_{0 \leq i \leq d-1} |\log|\xi_i|| + \frac{d(d-1)}{2} h(\alpha) + \log d \geq t \right) \\ &\leq \frac{\frac{d(d-1)}{2} h(\alpha) + \log d + d\,\mathbb{E} \max_{0 \leq i \leq d-1} |\log|\xi_i||}{t} \\ &\leq \frac{2d^2\eta + \frac{1}{2}d(d-1)\eta + \log d}{t} \end{aligned}$$

where we used the fact that

$$\begin{aligned} \mathbb{E} \max_{0 \leq i \leq d-1} |\log|\xi_i|| &\leq \sum_{i=0}^{d-1} E|\log|\xi_i|| \\ &\leq \sum_{i=0}^{d-1} \log|\{E|\xi_i|\}| \leq d\eta + \log(2)d \end{aligned}$$

(since $\mathbb{E}|\xi_i| \leq \left(\mathbb{E}|\xi_i|^2\right)^{\frac{1}{2}} \leq H_i$). By choosing $t := 5d^2\eta + 2\log d$ we get the result. ■

So if we condition on x satisfying

$$h(u) \leq 5d^2\eta + 2\log d \tag{3}$$

then our randomness model is a uniform sample among the algebraic numbers $u = \xi_0 + \cdots + \xi_{d-1}\alpha^{d-1}$ that satisfy Inequality (3). We call this model of randomness $\Gamma(c_0, \ldots, c_{d-1}, H_0, \ldots, H_{d-1}, \alpha)$. We will need the following lemma (see [ADRSD15, Lemma 2.9] or [MP12, Lemma 2.8]).

Lemma 1.18. *Let $x_1, \ldots, x_d$ be independent discrete Gaussian variables with mean 0 and standard deviation 1, and let $\gamma_1, \ldots, \gamma_d \in \mathbb{R}$. Then for every $t > 0$ we have*

$$\mathbb{P}\left(\left| \sum_{i=1}^{d} \gamma_i x_i \right| \geq t \left(\sum_{i=1}^{d} \gamma_i^2 \right)^{\frac{1}{2}} \right) \leq 2e^{-\frac{t^2}{2}}. \tag{4}$$

Lemma 1.19. *Let $\alpha > 0$ fixed real and ξ_i be independent discrete Gaussian random variables with mean 0 and standard deviation $H > 2$. Let $u := \sum_{i=0}^{d-1} \xi_i \alpha^i$. Then, for any $\varepsilon \in (0, \frac{1}{2}]$ and any z we have*

$$\mathbb{P}\left(|\log|u| - z| \leq \varepsilon\right) = \sqrt{2}\left(8\sqrt{2} \cdot \varepsilon \sqrt{\log 1/\varepsilon} + \frac{1}{2H} \right). \tag{5}$$

Proof: Let $B := H\left(\sum_{i=0}^{d-1} a^{2i}\right)^{\frac{1}{2}}$. We have that

$$\mathbb{P}\left(|\log|u| - z| \leq \varepsilon\right) = \mathbb{P}\left(e^{z-\varepsilon} \leq |u| \leq e^{z+\varepsilon}\right). \tag{6}$$

We consider cases.

Case 1: Assume that $e^{z-\varepsilon} \geq B\sqrt{2\log 1/\varepsilon}$. Then (4) implies

$$\mathbb{P}\left(|u| \geq e^{z-\varepsilon}\right) \leq \mathbb{P}\left(|u| \geq B\sqrt{2\log 1/\varepsilon}\right) \leq 2\varepsilon. \tag{7}$$

Case 2: Assume that $e^{z-\varepsilon} \leq B\sqrt{2\log 1/\varepsilon}$. First observe that for every $\delta > 0$, $w \in \mathbb{R}$, and $0 \leq i \leq d-1$,

$$\mathbb{P}\left(|\xi_i - w| \leq \delta\right) \leq \frac{\delta + 1}{2H}.$$

(as in the proof of Lemma 1.12). We apply Lemma 1.8 and then, for any interval I of width $\frac{2B\delta}{H}$, we have

$$\mathbb{P}\left(|u| \in I\right) \leq \sqrt{2}\frac{\delta + 1}{2H}.$$

We choose $\delta = \frac{2\varepsilon e^{z-\varepsilon} H}{B}$ and we notice that the interval $[e^{z-\varepsilon}, e^{z-\varepsilon}]$ has length at most $2\varepsilon e^{z-\varepsilon} = 4\frac{\delta B}{H}$ So in this case we have that

$$\mathbb{P}\left(e^{z-\varepsilon} \leq |u| \leq e^{z+\varepsilon}\right) \leq \sqrt{2}\frac{2\delta + 1}{2H} \leq \sqrt{2}\frac{8\sqrt{2} \cdot \varepsilon\sqrt{\log 1/\varepsilon} \cdot H + 1}{2H}$$

Combining the above we complete the proof. ■

Combining our last two lemmata with the definition of conditional probability gives us the following.

Lemma 1.20. *Let α be an algebraic number with $|\alpha| \geq 1$, and degree d over $\mathbb{Q}$. Let X be a random variable distributed according to $\Gamma(c_0, \ldots, c_{d-1}, H_0, \ldots, H_{d-1}, \alpha)$. Then for any $z \in \mathbb{R}$ and $\varepsilon \in \left(\frac{1}{H}, \frac{1}{4}\right)$ we have $\mathbb{P}\{|\log|X| - z| \leq \varepsilon\} = O\left(\varepsilon\sqrt{\ln(1/\varepsilon)}\right)$.* ■

Corollary 1.21. *Fix $b = (b_1, \ldots, b_m) \in (\mathbb{Z} \setminus \{0\})^m$ and let $(u_1, \ldots, u_m)$ be a random vector where each u_i is an independent random variable distributed according to $\Gamma(c_0, \ldots, c_{d-1}, H_0, \ldots, H_{d-1}, u)$. Then for any $z \in \mathbb{R}$ and $\varepsilon \in \left(\frac{1}{H}, \frac{1}{4}\right)$ we have*

$$\mathbb{P}\{|b_1 \log|u_1| + \cdots + b_m \log|u_m| - z| \leq \varepsilon |b|_2\} = O(\varepsilon\sqrt{\ln(1/\varepsilon)}).$$

Since $b_i^2 \geq 1$ for all i this also yields

$$\mathbb{P}\{|b_1 \log|u_1| + \cdots + b_m \log|u_m| - z| \leq \varepsilon\sqrt{m}\} = O\left(\varepsilon\sqrt{\ln(1/\varepsilon)}\right). \blacksquare$$

2 WHICH SIDE ARE YOU ON?: CIRCUIT DISCRIMINANTS AND THEIR SIGNS

Let us first recall a rational function of absolute values that is related to a particular class of $\mathcal{A}$-discriminant polynomials.

Definition 2.1. *Suppose $A = \{a_1, \ldots, a_{m+2}\} \subset \mathbb{Z}^n$ is such that $\widehat{\mathcal{A}} := \begin{bmatrix} 1 & \cdots & 1 \\ a_1 & \cdots & a_{m+2} \end{bmatrix} \in \mathbb{Z}^{(n+1)\times(m+2)}$ has distinct columns and rank $m+1$ for some $m \leq n$. Let $b \in \mathbb{Z}^{(m+2)\times 1}$ be any generator of the right $\mathbb{Z}$-nullspace of $\widehat{\mathcal{A}}$. We then call A a* non-degenerate *circuit if and only if b has no zero coordinates (and a* degenerate *circuit otherwise). Also, for any non-degenerate circuit $A \subset \mathbb{R}^n$ of cardinality $m+2$, and any nonzero real $c_1, \ldots, c_{m+2}$, we define $\Xi_A(c_1, \ldots, c_{m+2}) := \left(\prod_{i=1}^{m+2} |c_i/b_i|^{b_i}\right) - 1$.*

In our setting, the c_i will always be the coefficients of a polynomial f supported on the circuit A. So we will often abuse notation by writing $\Xi_A(f)$ instead of $\Xi_A(c_1, \ldots, c_{m+2})$, assuming $f(x) = c_1 x^{a_1} + \cdots + c_{m+2}x^{a_{m+2}}$. When restricted to a suitable orthant in $\mathbb{R}^{m+2}$, our Ξ_A is a monomial multiple of the *$\mathcal{A}$-discriminant polynomial* $\Delta_{\mathcal{A}}$ from [GKZ94, Ch. 9]. From the development of [GKZ94, Ch. 9] (restricted to $\mathbb{R}$) we have the following summary of the key properties of Ξ_A that we'll need:

Theorem 2.2. *Suppose $A = \{a_1, \ldots, a_{m+2}\} \subset \mathbb{Z}^n$ is a non-degenerate circuit of cardinality $m+2$, $f \in \mathbb{R}\left[x_1^{\pm 1}, \ldots, x_n^{\pm 1}\right]$ is supported on A, and $f(x) = c_1 x^{a_1} + \cdots + c_{m+2}x^{a_{m+2}}$. Then $Z_+(f)$ has a singularity if and only if $\Xi_A(f) = 0$ and $\operatorname{sign}(b_1 c_1) = \cdots = \operatorname{sign}(b_{m+2}c_{m+2})$. In particular, when $m = n$, such a $Z_+(f)$ has at most 1 singular point.* $\blacksquare$

Example 2.3. *$A = \{0, 2, 7\} \subset \mathbb{Z}^1$ is a non-degenerate circuit, and we see that a suitable $b \in \mathbb{Z}^3$ is $b = (5, -7, 2)^\top$ (taking $(\cdot)^\top$ to mean transpose). Theorem 2.2 then tells us that $f(x) := c_1 + c_2x^2 + c_3x^7$ has a degenerate positive root if and only if $[[c_1, c_3 > 0 > c_2$ or $c_1, c_3 < 0 < c_2]$ and $\left|\frac{c_1}{5}\right|^5 \left|\frac{c_2}{-7}\right|^{-7} \left|\frac{c_3}{2}\right|^2 = 1]$. Note that the last equality is equivalent to $5\log|c_1| - 7\log|c_2| + 2\log|c_3| = 5\log(5) - 7\log(7) + 2\log(2)$. Note also that $\Xi_A(5 - 7x^2 + 2x^7) = 0$ here, and the unique degenerate root of $5 - 7x^2 + 2x^7$ is 1.* $\diamond$

Example 2.4. *$A = \{(0,0), (2,2), (7,7)\} \subset \mathbb{Z}^2$ is also a non-degenerate circuit of cardinality 3 and the same $b \in \mathbb{Z}^3$ from Example 2.3 works for this example as well. We then get exactly the same criteria for $c_1 + c_2x_1^2x_2^2 + c_3x_1^7x_2^7$ to have a degenerate root as in Example 2.3. However, $5 - 7x_1^2x_2^2 + 2x_1^7x_2^7$ has infinitely many degenerate roots in $\mathbb{R}_+^2$: They are all of the form $(x_1, x_2) = (r, 1/r)$ for $r \in \mathbb{R}_+$.* $\diamond$

We let Conv$A$ denote the *convex hull* of A, i.e., the smallest convex set containing A.

Theorem 2.5. *[BRS09, Thm. 2.17] Following the notation of Theorem 2.2, $Z_+(f)$ is empty if and only if exactly one of the following two conditions holds:*

1. *All the c_i have the same sign.*
2. *Conv$A$ is a simplex, $\operatorname{sign}(b_1c_1) = \cdots = \operatorname{sign}(b_{m+2}c_{m+2})$, and $(\Xi_A(f) + 1)^{\operatorname{sign}(b_j)} < 1$ where j is the unique index with $\operatorname{sign}(b_ib_j) < 0$ for all $i \neq j$.*

Furthermore, $Z_+(f)$ consists of a single point if and only if all the following conditions hold: $m = n$, Conv$A$ is a simplex, $\operatorname{sign}(b_1c_1) = \cdots = \operatorname{sign}(b_{m+2}c_{m+2})$, and $\Xi_A(f) = 0$. $\blacksquare$

Note in particular that when Conv$A$ is *not* a simplex, checking emptiness for $Z_+(f)$ reduces to merely checking whether all the coefficients of f have the same sign or not. It is easily checked that Conv$A$ is a simplex if and only if the b-vector has exactly one nonzero coordinate differing in sign from all the other coordinates.

Remark 2.6. *Unravelling the characterization above, we see that unless all the c_i have the same sign, and* Conv$A$ *has a particular shape, we will need to compare a high-degree monomial in the c_i against 1 to know if $Z_+(f)$ is empty. The latter calculation is then clearly equivalent to computing the sign of $b_1 \log|c_1/b_1| + \cdots + b_{m+2}\log|c_{m+2}/b_{m+2}|$. This is our central reduction to linear forms in logarithms.* $\diamond$

Example 2.7. *Suppose $A \subset \mathbb{Z}^3$ consists of the columns of $\begin{bmatrix} 24 & 68 & -47 & 52 & 71 \\ -85 & -10 & -51 & 11 & 87 \\ -90 & 33 & 1 & 28 & 46 \end{bmatrix}$. Then* Conv$A$ *is a simplex, and Theorem 2.5 (along with a bit of Morse Theory [BDPRRR24]) tells us (assuming $c_1, c_2, c_3, c_5 > 0 > c_4$) that $Z_+\left(c_1x_1^{24}x_2^{-85}x_3^{-90} + c_2x_1^{68}x_2^{-10}x_3^{33} + c_3x_1^{-47}x_2^{-51}x_3 + c_4x_1^{52}x_2^{11}x_3^{28} + c_5x_1^{71}x_2^{87}x_3^{46}\right)$ is empty, a single point, or isotopic to a 2-sphere, according as $\sum_{i=1}^5 b_i \log|c_i|$ is less than, equal to, or greater than $\sum_{i=1}^5 b_i \log|b_i|$, where $b = (43403, 600796, 150818, -1138887, 343870)^\top$. This condition can clearly be handled reasonably via floating calculation on a computer — provided sufficient accuracy is used for the underlying logarithms.* $\diamond$

Remark 2.8. *We thus see that the sign of $\Xi_A(f)$ (or, equivalently, the sign of $\log(\Xi_A(f) + 1)$) appears to determine the isotopy type of $Z_+(f)$, at least in certain orthants of coefficient space. We call the connected components of the complement of the zero set of $\Xi_A(f)$, in the orthants of $(\mathbb{R} \setminus \{0\})^{n+2}$,* discriminant chambers. *One aspect of circuits that helps make computing the isotopy type of $Z_+(f)$ tractable (for f a circuit polynomial) is that every orthant of $\mathbb{R}^{n+2}$ contains at most 2 discriminant chambers. So, in the circuit case, the topological behavior of $Z_+(f)$ depends mainly on whether $f \in Z_{\mathbb{R}}(\Xi_A)$, or on which "side" of $Z_{\mathbb{R}}(\Xi_A)$ f lies.* $\diamond$

The degenerate circuit analogue of Theorem 2.5 is similar. In particular, recall that for any degenerate circuit $A = \{a_1, \ldots, a_{m+2}\} \subset \mathbb{R}^n$ of cardinality $m+2$, with corresponding right null vector b for $\widehat{\mathcal{A}}$, the subset $B := \{a_i \mid b_i \neq 0\}$ is a non-degenerate circuit. We also let $f_B(x) := \sum_{a_i \in B} c_i x^{a_i}$.

Theorem 2.9. *[BRS09, Thm. 2.18] Following the notation above, suppose A is a degenerate circuit, $f \in \mathbb{R}\left[x_1^{\pm 1}, \ldots, x_n^{\pm 1}\right]$ is supported on A, and $f(x) = c_1 x^{a_1} + \cdots + c_{m+2}x^{a_{m+2}}$. Then $Z_+(f)$ is empty if and only if at least one of the following conditions holds:*

1. *All the c_i have the same sign.*

(2) *(a)* $\mathrm{Conv}A$ *is a simplex, (b)* $\mathrm{sign}(b_i c_i)$ *is constant as* i *ranges over all indices with* $a_i \in B$*, (c)* $\mathrm{sign}(c_i c_j) < 0$ *for some* i *with* $a_i \notin B$*, and (d)* $(\Xi_B(f_B)+1)^{\mathrm{sign}(b_j)} \leq 1$ *where* j *is the unique index with* $b_j \neq 0$ *and* $\mathrm{sign}(b_i b_j) \leq 0$ *for all* $i \neq j$*.* ■

Example 2.10. *With* $A = \{(0,0),(1,0),(2,0),(0,1)\}$ *it is easily checked that* $b = (1,-2,1,0)^\top$ *is a suitable right nullvector for* $\widehat{\mathcal{A}}$*, and this* b *has a unique negative coordinate. So the* j *from Theorem 2.9 is* $j=2$*. Furthermore, for* $f(x_1,x_2) = 1 - cx_1 + x_1^2 + x_2$ *and* $c > 0$*, we see that* $f_B = 1 - cx_1 + x_1^2$*,* $\Xi_B(f_B)+1 = \frac{2}{c}$*,* $\mathrm{sign}(b_j) = -1$*, and thus* $Z_+(f)$ *is empty if and only if* $c \leq 2$*.* ⋄

So in the end, although the indexing is slightly more complicated for the degenerate circuit case, we can again reduce detecting emptiness of $Z_+(f)$ to checking the sign of an integer linear form in logarithms of integers.

Before moving on, we must also recall an explicit bound on the complexity of computing the sign of a linear form in logarithms. First, we recall the following paraphrase of a bound of Matveev [Mat00, Cor. 2.3], considerably strengthening earlier bounds of Baker and Wustholtz [BW93]. (See also [BMS06, Thm. 9.4].)

Theorem 2.11. *Suppose* K *is a degree* d *real algebraic extension of* $\mathbb{Q}$*,* $c_1,\ldots,c_m \in K\setminus\{0\}$*, and* $b_1,\ldots,b_m \in \mathbb{Z}\setminus\{0\}$*. Let* $B := \max\{|b_1|,\ldots,|b_m|\}$ *and* $\log H_i := \max\{dh(c_i), |\log c_i|, 0.16\}$ *for all* i*. Then* $\sum_{i=1}^m b_i \log c_i \neq 0$ *implies that* $\log\left|\sum_{i=1}^m b_i \log c_i\right|$ *is strictly greater than* $-1.4 \cdot m^{4.5} 30^{m+3} d^2 (1+\log d)(1+\log B)\prod_{i=1}^m \log H_i$*.* ■

We must also recall the following classical fact on approximating logarithms via Arithmetic-Geometric Iteration:

Theorem 2.12. *[Ber03, Sec. 5] Given any positive* $x \in \mathbb{Q}$ *of logarithmic height* h*, and* $\ell \in \mathbb{N}$ *with* $\ell \geq h$*, we can compute* $\lfloor \log_2 \max\{1, \log|x|\} \rfloor$ *and the* ℓ *most significant bits of* $\log x$ *in time* $O(\ell \log^2 \ell)$*.* ■

Taking $d = 1$, a consequence of the preceding two bounds derived in [Roj24, Proof of Lemma 4.2] is the following algorithmic complexity bound:

Corollary 2.13. *For any* $b = (b_1,\ldots,b_m) \in \mathbb{Z}^m$ *and* $\gamma = (\gamma_1,\ldots,\gamma_m) \in \mathbb{Q}^m$ *with* $B := \max_i |b_i|$ *and* $\log H := \max_i h(\gamma_i)$*, we can compute the sign of* $\Lambda(b, |\gamma|) = \sum_{i=1}^m b_i \log|\gamma_i|$ *in time*

$$O((31\log H)^m \log(B)\log^2(\log(B)\log H)).$$ ■

By combining Corollary 2.13 with Theorems 2.5 and 2.9, we immediately obtain an explicit (deterministic) complexity bound for detecting positive roots for circuit polynomials, i.e., the main results of [BRS09]. However, the resulting complexity bound is exponential in n. Our entire goal is to reduce this time bound to polynomial in n and, thanks to our probabilistic corollaries, we'll at least accomplish this for a large fraction of inputs. But first let us complete our background by reviewing real root detection for a single $(n+1)$-nomial.

2.1 A Brief Note on the Case $k = n+1$

We mentioned earlier that detecting real roots for an n-variate $(n+1)$-nomial is much easier than for an $(n+2)$-nomial. This is because of the following fact that can be found in earlier work of Reznick [Rez78]. For convenience, we provide a proof that works for real exponents as well.

Lemma 2.14. *Suppose* $f \in \mathbb{R}\left[x_1^{\pm 1},\ldots,x_n^{\pm 1}\right]$ *can be written in the form* $f(x) = \sum_{i=1}^{n+1} c_i x^{a_i}$ *where* $A = \{a_1,\ldots,a_{n+1}\}$ *is the vertex set of an* n*-simplex, i.e., the rank of* $\widehat{\mathcal{A}}$ *is* $n+1$*. Then* $Z_+(f)$ *is empty if and only if all the* c_i *have the same sign.*

Proof: Substituting $x_i = e^{y_i}$ for all i, we see that $Z_+(f)$ is empty if and only if the real zero set, $Z_{\mathbb{R}}(g)$, of the exponential sum $g(y) := \sum_{j=1}^{n+1} c_j e^{a_j \cdot y}$ is empty. Since $Z_{\mathbb{R}}(g)$ is invariant under translation of A, we may assume a_1 is the origin.

Noting that the emptiness of $Z_{\mathbb{R}}(g)$ is invariant under invertible linear maps applied to the variables, we can substitute $y \mapsto My$, where we can consider y as a column vector, and let M be the inverse of the $n \times n$ matrix whose ith column is a_{i+1}. (M is invertible since the edge vectors of any vertex of a simplex are linearly independent.) So we may assume $g(y) = c_1 + c_2 e^{y_1} + \cdots + c_{n+1} e^{y_n}$. Finally, since $Z_{\mathbb{R}}(g)$ is invariant under nonzero scaling of g, and the emptiness of $Z_{\mathbb{R}}(g)$ is invariant under translation of the variables, we may assume $g(y) = \varepsilon_1 + \varepsilon_2 e^{y_1} + \cdots + \varepsilon_{n+1} e^{y_n}$ where $\varepsilon_i \in \{\pm 1\}$ has the same sign as c_i. Letting $u_i = e^{y_i}$ for all i, we are reduced to deciding the emptiness of $Z_+(\varepsilon_1 + \varepsilon_2 u_1 + \cdots + \varepsilon_{n+1} u_n)$. The latter zero set is clearly empty if and only if all the ε_i have the same sign. ■

3 THE PROOF OF THEOREM 1.4

First note that, out of the 2^{n+2} orthants of $(c_1,\ldots,c_{n+2}) \in (\mathbb{R}\setminus\{0\})^{n+2}$, exactly *two* of these orthants satisfy the condition

$$(\star) \qquad \mathrm{sign}(b_1 c_1) = \cdots = \mathrm{sign}(b_{n+2} c_{n+2}).$$

For those orthants *not* satisfying Condition $(\star)$, Theorems 2.5 and 2.9 tell us that checking $Z_+(f) \stackrel{?}{=} \emptyset$ is almost trivial: We merely need to check whether all the c_i have the same sign. Note also that the sign of $\Xi_A(f)$ (or, equivalently, $\log(\Xi_A(f)+1)$) is independent of the signs of the c_i. So the inputs where checking $Z_+(f) \stackrel{?}{=} \emptyset$ is harder are exactly the inputs where $\log(\Xi_A(f)+1)$ requires more accuracy to evaluate. So by Corollary 1.11, we obtain that we can decide $Z_+(f) \stackrel{?}{=} \emptyset$ easily on a fraction of $1 - O\left(\frac{1}{2^n H}\right)$ of our input f, since our underlying probability measure is uniform across all orthants.

So now we must precisely quantify what we mean by "more accuracy" and "easily": Corollary 1.11 tells us that in the two orthants satisfying Condition $(\star)$, with probability $1 - O(1/H)$, we have:

$$\left|\left(\sum_{i=1}^{n+2} b_i \log|c_i|\right) - \left(\sum_{i=1}^{n+2} b_i \log|b_i|\right)\right| \quad > \quad \frac{\sqrt{n+2}}{H} \tag{8}$$

if

$$\sum_{i=1}^{n+2} b_i \log|c_i| \neq \sum_{i=1}^{n+2} b_i \log|b_i|.$$

In other words, we now know that for most inputs in our two special orthants, "moderate" accuracy for each logarithm in the sums above will suffice to correctly determine which of $\sum_{i=1}^{n+2} b_i \log|c_i|$ or $\sum_{i=1}^{n+2} b_i \log|b_i|$ is bigger (or if they are equal). More precisely, simply let $B := \max_i |b_i|$ and let L_i and M_i be rational numbers satisfying $|L_i - \log|c_i|| < \frac{\sqrt{n+2}}{6nBH}$ and $|M_i - \log|b_i|| < \frac{\sqrt{n+2}}{6nBH}$. Then by the Triangle Inequality, the values of $\left(\sum_{i=1}^{n+2} b_i \log|c_i|\right) - \left(\sum_{i=1}^{n+2} b_i \log|b_i|\right)$ and $\left(\sum_{i=1}^{n+2} L_i\right) - \left(\sum_{i=1}^{n+2} M_i\right)$ differ by no more than $\frac{\sqrt{n+2}}{3H}$. In other

words, to decide whether $\Lambda(b,c)$ is negative, zero, or positive, we merely check whether $\left(\sum_{i=1}^{n+2} L_i\right) - \left(\sum_{i=1}^{n+2} M_i\right)$ is less than $-\frac{2\sqrt{n+2}}{3H}$, inside of the open interval $\left(-\frac{1\sqrt{n+2}}{3H}, \frac{1\sqrt{n+2}}{3H}\right)$, or greater than $\frac{2\sqrt{n+2}}{3H}$: These are the only possibilities that can occur on our $1 - O(1/H)$ fraction of inputs from our two orthants satisfying Condition ($\star$), thanks to Corollary 1.11.

To conclude, observe that Cramer's Rule (and Hadamard's Inequality for determinants) tells us that the height of b_i is $O(n\log(dn))$. So by Theorem 2.12, each $\log|c_i|$ and $\log|b_i|$ term can be approximated to our desired accuracy in time

$$O((n+\log(dH)+\log(H)+n\log(dn))\log^2(n\log(ndH))),$$

which is simply $O(n\log^3(ndH))$. So computing $L_1, M_1, \ldots, L_{n+2}, M_{n+2}$ takes time $O(n^2\log^3(ndH))$. The computation of b takes time $n^{3.373}\log^{1+o(1)}(nd)$ via fast integer linear algebra (see, e.g., [Roj24, Lemma 2.1]). So our overall time bound is $O(n^2\log^3(ndH)) + n^{3.373}\log^{1+o(1)}(nd) = O(n^{3.373}\log^3(ndH))$. ■

Example 3.1. *Suppose* $A \subset \mathbb{Z}^5$ *consists of the columns of*

$$\begin{bmatrix} -13 & 47 & -85 & -84 & 95 & 5 & 75 \\ 94 & -15 & 60 & 82 & -84 & 52 & -7 \\ 51 & -41 & 44 & -87 & -89 & -22 & 27 \\ 91 & 34 & -32 & -18 & 16 & 30 & -53 \\ 43 & -17 & -38 & -60 & -48 & 57 & -57 \end{bmatrix}.$$

Then a suitable b*-vector is* $(b_1,\ldots,b_7) :=$ $(6581562115, -23794818871, 732603963, 950736962, 13899922091, 1381315615, 248678125)$. *All our theory so far tells us that the positive zero set of* $f(x) = c_1x^{a_1} + \cdots + c_7x^{a_7}$*, with* a_i *the* i*th column of the matrix above, can be decided efficiently for "most"* $(c_1,\ldots,c_7) \in \mathbb{Z}^7$*. In particular, the proof of Theorem 1.4 tells us that we can accomplish this by deciding the sign of the linear combination of logarithms* $\log(\Xi_A(f)+1) = \sum_{i=1}^7 b_i\log|c_i/b_i|$*. So let us try to see concretely how big a fraction "most" is for the support above: Let us consider the distribution of the values of* $\log(\Xi_A(f)+1)$ *as the coefficients of* f *range uniformly over* $\{\pm1,\ldots,\pm1000\}$*: After a sample of* 10^7 *uniformly random trials, we observed a minimal value of* $134.0879\ldots$ *for* $\log(\Xi_A(f)+1)$*, attained at* $(|c_1|,\ldots,|c_7|) = (702, 646, 78, 856, 661, 821, 905)$*. In particular,* $\log 134.08979 = 4.8985\ldots$*, and thus the implied lower bound of* $-3.1255\ldots\times10^{25}$ *from Theorem 2.11 appears far more pessimistic than necessary. A histogram for the values of* $\log(\Xi_A(f)+1)$ *from our trial is plotted below:*

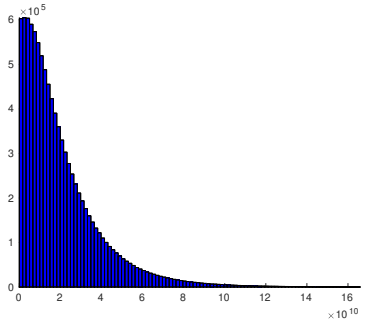

So the y-axis measures the frequency of a potential value for $\log(\Xi_A(f)+1)$*, and thus the smaller values appear to occur more often. While the minimum value of* $\log(\Xi_A(f)+1)$ *over our sample is not visible from the graph, one can easily see that about* 75% *of the values of* $\log(\Xi_A(f)+1)$ *are less than* 3×10^{10}. ⋄

3.1 Fast on Average vs. Fast for Most Inputs

In closing, we emphasize that the speed-up we have derived applies only to an *unknown but large fraction of our inputs*. While it would be even better to prove a speed-up that holds with high probability for all inputs, obvious tools for such an average-case speed-up are missing. For instance, one could hope for a simple way to check that a computed approximation for $\log(\Xi_A(f)+1)$ is incorrect. (So one could redo the approximation at higher accuracy and get a correct answer.) However, an *efficient* check of this kind is not yet available for linear combinations of logarithms of rational numbers. Furthermore, the best provable bounds for linear forms in logarithms are so large that, even with the ability to check correctness of the sign of $\log(\Xi_A(f)+1)$, averaging over all inputs does not yield an average-case complexity estimate polynomial in n.

So we hope that our new bounds avoiding "purple swans" spur further practical improvements to bounds for linear forms in logarithms, and other complexity results in real algebraic geometry.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the US National Science Foundation for their support through grants CCF-1900881 and CCF-2110075. The authors also humbly thank the anonymous referees who pointed out the valuable references [CS16, MSdW19, PKC22, Rez78].

REFERENCES

[ADRSD15] Divesh Aggarwal, Daniel Dadush, Oded Regev, Noah Stephens-Davidowitz, *"Solving the Shortest Vector Problem in* 2^n *Time Using Discrete Gaussian Sampling,"* Proceedings of STOC '15 (forty-seventh annual ACM symposium on Theory of Computing, June 2015), pp. 733–742, ACM Press.

[MP12] Daniele Micciancio and Chris Peikert, *"Trapdoors for lattices: Simpler, tighter, faster, smaller"* Annual International Conference on the Theory and Applications of Cryptographic Techniques, pg 700–718, 2012

[AL17] Dennis Amelunxen and Martin Lotz, *"Average-case complexity without the black swans,"* Journal of Complexity, Volume 41, August 2017, pp. 82–101.

[Bak77] Alan Baker, *"The Theory of Linear Forms in Logarithms,"* in Transcendence Theory: Advances and Applications: proceedings of a conference held at the University of Cambridge, Cambridge, Jan.–Feb., 1976, Academic Press, London, 1977.

[BW93] Alan Baker and Gisbert Wustholtz, *"Logarithmic forms and group varieties,"* J. Reine Angew. Math. **442** (1993), pp. 19–62.

[BHPR11] Osbert Bastani; Chris Hillar; Dimitar Popov; and J. Maurice Rojas, *"Randomization, Sums of Squares, and Faster Real Root Counting for Tetranomials and Beyond,"* Randomization, Relaxation, and Complexity in Polynomial Equation Solving, Contemporary Mathematics, vol. 556, pp. 145–166, AMS Press, 2011.

[BR14] Saugata Basu and Marie-Françoise Roy, *"Divide and conquer roadmap for algebraic sets,"* Discrete and Computational Geometry, 52:278–343, 2014.

[BPR06] Saugata Basu; Ricky Pollack; and Marie-Francoise Roy, *Algorithms in Real Algebraic Geometry,* Algorithms and Computation in Mathematics, vol. 10, Springer-Verlag, 2006.

[Ber03] Daniel J. Bernstein, *"Computing Logarithm Intervals with the Arithmetic-Geometric Mean Iterations,"* available from `http://cr.yp.to/papers.html` .

[Bih11] Frédéric Bihan, *"Topologie des variétés creuses,"* Habilitation à Diriger des Recherches, Universitde Savoie, France, 2011. Downloadable from http://www.lama.univ-savoie.f/˜bihan/MyPapers/HDRfinal.pdf

[BDPRRR24] Frédéric Bihan; Weixun Deng; Kaitlyn Phillipson; Erika Refsland; Robert J. Rennie; and J. Maurice Rojas, *"Quickly Computing Isotopy Type for Exponential Sums over Circuits,"* in preparation, 2024.

[BRS07] Frédéric Bihan; J. Maurice Rojas; and Frank Sottile, *"Sharpness of fewnomial bound and the number of components of a fewnomial hypersurface,"* IMA Volume 146: Algorithms in Algebraic Geometry edited by Alicia Dickenstein, Frank-Olaf Schreyer, and Andrew J. Sommese, pp. 15–20, Springer, New York, 2007.

[BRS09] Frédéric Bihan; J. Maurice Rojas; and Casey Stella, *"Faster real feasibility via circuit discriminants,"* proceedings of ISSAC 2009 (July 28-31, Seoul, Korea), pp. 39–46, ACM Press, 2009.

[BS07] Frédéric Bihan and Frank Sottile, *"New Fewnomial Upper Bounds from Gale Dual Polynomial Systems,"* Moscow Mathematical Journal, 7 (2007), no. 3, pp.

387–407.

[BS09] Frédéric Bihan and Frank Sottile, *"Betti number bounds for fewnomial hypersurfaces via stratified Morse theory,"* Proc. Amer. Math. Soc., 137, No. 9 (2009), pp. 2825–2833.

[BG06] Enrico Bombieri and Walter Gubler, *Heights in Diophantine Geometry,* new mathematical monographs: 4, Cambridge University Press, 2006.

[BMS06] Yan Bugeaud; Maurice Mignotte; and Samir Siksek, *"Classical and modular approaches to exponential Diophantine equations, I, Fibonacci and Lucas perfect powers,"* Ann. of Math. (2) **163** (2006), pp. 969–1018.

[BET-C19] Peter Bürgisser, Alperen A. Ergür, and Josué Tonelli-Cueto, *"On the number of real zeros of random fewnomials,"* SIAM Journal on Applied Algebra and Geometry 3 (2019), no. 4, pp. 721–732.

[CS16] V. Chandrasekaran and P. Shah, *"Relative Entropy Relaxations for Signomial Optimzation,"* SIAM J. Optim. 26 (2), (2016), pp. 1147–1173.

[BSY14] Kousha Etessami, Alistair Stewart, Mihalis Yannakakis, *"A note on the complexity of comparing succinctly represented integers, with an application to maximum probability parsing,"* ACM Trans. Comput. Theory 6 (2), 9, 2014.

[GKZ94] Israel Moseyevitch Gel'fand; Misha M. Kapranov; and Andrei V. Zelevinsky, *Discriminants, Resultants and Multidimensional Determinants,* Birkhäuser, Boston, 1994.

[Kho91] Askold Khovanski, *Fewnomials,* AMS Press, Providence, Rhode Island, 1991.

[Lan78] Serge Lang, *Elliptic Curves: Diophantine Analysis,* Springer, 1978.

[LPP16] Galyna Livshyts, Grigoris Paouris, and Peter Pivovarov, *"On sharp bounds for marginal densities of product measures,"* Israel Journal of Mathematics, https://link.springer.com/article/10.1007/s11856-016-1431-5 .

[MSdW19] V. Magron, H. Seidler, and T. de Wolff, *"Exact optimization via sums of nonnegative circuits and arithmetic-geometric-mean-exponentials,"* Proceedings of ISSAC 2019, pp. 291–298, ACM Press.

[Mat00] E. M. Matveev, *"An explicit lower bound for a homogeneous rational linear form in logarithms of algebraic numbers, II"*, Izv. Ross. Akad. Nauk Ser. Mat. **64** (2000), pp. 125–180; English transl. in Izv.l Math. **64** (2000), pp. 1217–1269.

[Nes03] Yuri Nesterenko, *"Linear forms in logarithms of rational numbers,"* Diophantine approximation (Cetraro, 2000), pp. 53–106, Lecture Notes in Math., 1819, Springer, Berlin, 2003.

[PKC22] C. Pantea, H. Koeppl, and G. Craciun, *"Global injectivity and multiple equlibria,"* in uni- and bimolecular reaction networks," Discrete Contin. Dyn. Syst. Ser. B, 17 (6), (2012), pp. 2153–2170.

[PRT09] Philippe P. Pébay; J. Maurice Rojas; and David C. Thompson, *"Optimization and* $\mathbf{NP}_{\mathbb{R}}$*-Completeness of Certain Fewnomials,"* proceedings of SNC 2009 (August 3–5, 2009, Kyoto, Japan), pp. 133–142, ACM Press, 2009.

[Rez78] B. Reznick, *"Extremal PSD Forms with Few Terms,"* Duke Math. Journal, 45 (2), (1978), pp. 363–374.

[Roj24] J. Maurice Rojas, *"Counting Real Roots in Polynomial-Time via Diophantine Approximation,"* Found. Comput. Math. (2024), Vol. 24, pp. 639–681.

[RV15] Mark Rudelson and Roman Vershynin, *"Small Ball Probabilities for Linear Images of High-Dimensional Distributions,"* International Mathematics Research Notices, Vol. 2015, Issue 19, 2015, pp. 9594–9617, https://doi.org/10.1093/imrn/rnu243

[Sch92] Wolfgang M. Schmidt, *"Integer points on curves of genus* 1*,"* Compositio Mathematica, tome 81, no. 1 (1992), pp. 33–59.