

Faster p -adic Feasibility for Certain Multivariate Sparse Polynomials

Martín Avendaño* Ashraf Ibrahim† J. Maurice Rojas* Korben Rusek*

June 12, 2021

Abstract

We present algorithms revealing new families of polynomials admitting sub-exponential detection of p -adic rational roots, relative to the sparse encoding. For instance, we prove **NP**-completeness for the case of honest n -variate $(n+1)$ -nomials and, for certain special cases with p exceeding the Newton polytope volume, constant-time complexity. Furthermore, using the theory of linear forms in p -adic logarithms, we prove that the case of trinomials in one variable can be done in **NP**. The best previous complexity upper bounds for all these problems were **EXPTIME** or worse. Finally, we prove that detecting p -adic rational roots for sparse polynomials in one variable is **NP**-hard with respect to randomized reductions. The last proof makes use of an efficient construction of primes in certain arithmetic progressions. The smallest n where detecting p -adic rational roots for n -variate sparse polynomials is **NP**-hard appears to have been unknown.

1 Introduction

Paralleling earlier results over the real numbers [BRS09], we study the complexity of detecting p -adic rational roots for sparse polynomials. We find complexity lower bounds over \mathbb{Q}_p hitherto unattainable over \mathbb{R} (see Thm. 1.2 and Prop. 1.3 below), as well as new algorithms over \mathbb{Q}_p with complexity close to that of recent algorithms over \mathbb{R} (see Thm. 1.4 below).

More precisely, for any commutative ring R with multiplicative identity, we let \mathbf{FEAS}_R — the R -feasibility problem (a.k.a. Hilbert’s Tenth Problem over R [DLPvG00]) — denote the problem of deciding whether an input polynomial system $F \in \bigcup_{k,n \in \mathbb{N}} (\mathbb{Z}[x_1, \dots, x_n])^k$ has a root in R^n . Observe that $\mathbf{FEAS}_{\mathbb{R}}$, $\mathbf{FEAS}_{\mathbb{Q}}$, and $\{\mathbf{FEAS}_{\mathbb{F}_q}\}_{q \text{ a prime power}}$ are central problems respectively in algorithmic real algebraic geometry, algorithmic number theory, and cryptography.

Algorithmic results over the p -adics are useful in many settings: polynomial-time factoring algorithms over $\mathbb{Q}[x_1]$ [LLL82], computational complexity [Roj02, Che04, Koi11], the

*TAMU 3368, Math Dept., College Station, TX 77843-3368, USA. mavendar@yahoo.com.ar , rojas@math.tamu.edu , korben@rusek.org . Partially supported by NSF MCS grant DMS-0915245. J.M.R. and K.R. also partially supported by Sandia National Labs and DOE ASCR grant DE-SC0002505. Sandia is a multiprogram laboratory operated by Sandia Corp., a Lockheed Martin Company, for the US DOE under Contract DE-AC04-94AL85000.

†TAMU 3141, Aerospace Engineering Dept., College Station, TX 77843-3141, USA, ibrahim@aero.tamu.edu

study of prime ideals in number fields [Coh94, Ch. 4 & 6], elliptic curve cryptography [Lau04], and the computation of zeta functions [CDV06, LW08, Cha08]. Also, much work has gone into using p -adic methods to algorithmically detect rational points on algebraic curves via variations of the *Hasse Principle*¹ (see, e.g., [C-T98, Poo06]). We discuss further theoretical motivation in Section 1.2 below. Nevertheless, our knowledge of the complexity of deciding the existence of solutions for *sparse* polynomial equations over \mathbb{Q}_p is surprisingly coarse: good bounds for the number of solutions over \mathbb{Q}_p in one variable weren't even known until the late 1990s [Len99].

Definition 1.1 Let $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}$ denote the problem of deciding, for an input Laurent polynomial system $F \in \bigcup_{k,n \in \mathbb{N}} (\mathbb{Z}[x_1^{\pm 1}, \dots, x_n^{\pm 1}])^k$ and an input prime p , whether F has a p -adic rational root. Also let $\mathcal{P} \subset \mathbb{N}$ denote the set of primes, $p \in \mathcal{P}$, and, when \mathcal{I} is a family of such pairs (F, p) , we let $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{I})$ denote the restriction of $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}$ to inputs in \mathcal{I} .

When $a_j \in \mathbb{Z}^n$, the notation $a_j = (a_{1,j}, \dots, a_{n,j})$, $x^{a_j} = x_1^{a_{1,j}} \cdots x_n^{a_{n,j}}$, and $x = (x_1, \dots, x_n)$ will be understood. Also, when $f(x) := \sum_{j=1}^m c_j x^{a_j}$ with $c_j \in R \setminus \{0\}$ for all j , and the $a_j \in \mathbb{Z}^n$ are pair-wise distinct, we call f an n -variate m -nomial (over R), and we define $\text{Supp}(f) := \{a_1, \dots, a_m\}$ to be the support of f . We also define $\text{Newt}(f)$ — the (standard) Newton polytope of f — to be the convex hull of² $\text{Supp}(f)$ and let V_f denote its n -dimensional volume, normalized so that $[0, 1]^n$ has volume 1.

Let $\text{size}(f) := \sum_{i=1}^m \log_2 [(2 + |c_i|)(2 + |a_{1,i}|) \cdots (2 + |a_{n,i}|)]$ and, when $F := (f_1, \dots, f_k)$, $\text{size}(F) := \sum_{i=1}^k \text{size}(f_i)$. The underlying input sizes for $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}$ and $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{I})$ shall then be $\text{size}(F) + \log p$, but for $\text{FEAS}_{\mathbb{Q}_p}$ (and any prime p) we will instead use $\text{size}(F)$ as the input size. Finally, we let $\mathcal{F}_{n,m}$ denote the set of all n -variate m -nomials and, for any $m \geq n + 1$, we let $\mathcal{F}_{n,m}^* \subseteq \mathcal{F}_{n,m}$ denote the subset consisting of those f with $V_f > 0$. We call any $f \in \mathcal{F}_{n,m}^*$ an honest n -variate m -nomial (or honestly n -variate). \diamond

As an example, it is clear that upon substituting $y_1 := x_1^2 x_2 x_3^7 x_4^3$, the dishonestly 4-variate trinomial $-1 + 7x_1^2 x_2 x_3^7 x_4^3 - 43x_1^{198} x_2^{99} x_3^{693} x_4^{297}$ (with support contained in a line segment) has a root in $(\mathbb{Q}_p^*)^4$ if and only if the *honest* univariate trinomial $-1 + 7y_1 - 43y_1^{99}$ has a root in \mathbb{Q}_p^* . Via the use of Hermite Normal Form (as in Section 3 below), it is then easy to see that there is no loss of generality in restricting to $\mathcal{F}_{n,n+k}^*$ (with $k \geq 1$) when studying the algorithmic complexity of sparse polynomials. Note also that the degree, $\deg f$, of a polynomial f can sometimes be exponential in $\text{size}(f)$ for certain families of f , e.g., $d \geq 2^{\text{size}(1+5x_1^{126}+x_1^d)-16}$.

While there are now randomized algorithms of expected complexity polynomial in $\deg(f) + \text{size}(f) + \log p$ for factoring $f \in \mathbb{Z}[x_1]$ over $\mathbb{Q}_p[x_1]$ [CG00] (see also [Chi91]), no such algorithms are known to have complexity polynomial in just $\text{size}(f) + \log p$. Our first result shows that the existence of such algorithms would imply a collapse close to $\mathbf{P} = \mathbf{NP}$.

Theorem 1.2

If $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathcal{P}) \in \mathbf{ZPP}$ then $\mathbf{NP} \subseteq \mathbf{ZPP}$. Moreover, if the Wagstaff Conjecture is true, then the stronger implication $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathcal{P}) \in \mathbf{P} \implies \mathbf{P} = \mathbf{NP}$ holds

¹If $F(x_1, \dots, x_n) = 0$ is any polynomial equation and Z_K is its zero set in K^n , then the Hasse Principle is the assumption that [$Z_{\mathbb{C}}$ smooth, $Z_{\mathbb{R}} \neq \emptyset$, and $Z_{\mathbb{Q}_p} \neq \emptyset$ for all primes p] implies $Z_{\mathbb{Q}} \neq \emptyset$ as well. The Hasse Principle is a theorem when $Z_{\mathbb{C}}$ is a quadric hypersurface or a curve of genus zero, but fails in subtle ways already for curves of genus one (see, e.g., [Poo01]).

²i.e., smallest convex set containing...

The complexity class **ZPP** satisfies $\mathbf{P} \subseteq \mathbf{ZPP} \subseteq \mathbf{NP}$ and consists of decision problems admitting polynomial-time Las Vegas randomized algorithms (see Section 2, or [Pap95] for an excellent textbook treatment). The *Wagstaff Conjecture* is the assertion that, for any $\delta > 0$, the least prime congruent to $k \pmod N$ is bounded above by $\varphi(N) \log^{2+\delta} N$, where $\varphi(N)$ is the number of integers in $\{1, \dots, N\}$ relatively prime to N . (See the paragraph containing Inequality (1) in [Wag79] and the discussion in [BS96, Ch. 8, pp. 223–224 & 252–255].) This conjectural bound is unfortunately much stronger than the known implications of the Generalized Riemann Hypothesis.

The least n making $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x_1, \dots, x_n] \times \mathcal{P})$ be **NP**-hard appears to have been unknown. Theorem 1.2 thus comes close to settling this problem. In particular, while it is not hard to show that the broader problem $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}$ is **NP**-hard, the proof of Theorem 1.2 in Section 6 make essential use of a deep result of Alford, Granville, and Pomerance [AGP94] on primes in random arithmetic progressions. (See also Theorem 1.9 in Section 1.2 below).

One can also ask for the smallest k making it **NP**-hard to detect p -adic rational roots of honest n -variate $(n+k)$ -nomials. Curiously, $k=1$ suffices.

Proposition 1.3 *For any prime p we have that $\mathbf{FEAS}_{\mathbb{Q}_p}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*\right)$ is **NP**-hard.*

The preceding result, mentioned to Rojas by Bjorn Poonen during a conversation at the Extensions of Hilbert’s Tenth Problem workshop at the American Institute of Mathematics (March 21–25, 2005), does not appear to have been published. So we paraphrase Poonen’s proof in Section 7 below. We find Poonen’s result particularly interesting because **NP**-hardness occurs much less abruptly over the real numbers: we have instead the containment $\mathbf{FEAS}_{\mathbb{R}}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*\right) \in \mathbf{NC}^1$ [BRS09, Thm. 1.3]. (Recall that $\mathbf{NC}^1 \subseteq \mathbf{P}$.) In particular, **NP**-hardness of *real* root detection for n -variate $(n+k)$ -nomials is only known for $k = \Omega(n^\varepsilon)$ with any $\varepsilon > 0$ [BRS09, Thm. 1.3]. We discuss the case of fixed n in Section 1.1 below.

Our next main result reveals faster algorithms than previously known for detecting p -adic rational roots of certain sparse polynomials.

Theorem 1.4

0. $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{F}_{1,m} \times \mathcal{P}) \in \mathbf{P}$ for $m \in \{0, 1, 2\}$.
1. For any fixed prime p we have $\mathbf{FEAS}_{\mathbb{Q}_p}(\mathcal{F}_{1,3}) \in \mathbf{NP}$.
2. There is a countable union of algebraic hypersurfaces $\mathcal{E} \subsetneq \mathbb{Z}[x_1] \times \mathcal{P}$, with natural density 0, such that $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}((\mathbb{Z}[x_1] \times \mathcal{P}) \setminus \mathcal{E}) \in \mathbf{NP}$.
3. (a) $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}\left(\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*\right) \times \mathcal{P}\right) \in \mathbf{NP}$.
 (b) Letting $\mathcal{Q} := \{c_0 + c_1 x_1^2 + \dots + c_n x_n^2 \mid n \in \mathbb{N}; c_0, \dots, c_n \in \mathbb{Z} \setminus \{0\}\} \times \mathcal{P}$, we have $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{Q}) \in \mathbf{P}$.
 (c) Let $\mathcal{W} \subset \left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*\right) \times \mathcal{P}$ denote the subset consisting of those (f, p) with $f(x) = c_0 + c_1 x_1^{d_1} + \dots + c_n x_n^{d_n}$, $d_1, \dots, d_n \in \mathbb{N}$, $n \geq 2$, $p \geq (n! V_f)^{2/(n-1)}$, and p not dividing $n! V_f$ or any coefficient of f . Then for any such $(f, p) \in \mathcal{W}$, f always has a root in \mathbb{Q}_p^n , i.e., $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{W})$ is doable in constant time.

1.1 What is New About the Algorithms in Theorem 1.4?

As evinced by Parts (b) and (c) of Assertion (3) of Theorem 1.4, algorithms for $\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}\left(\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*\right) \times \mathcal{P}\right)$ clearly complement classical results on quadratic forms (see,

e.g., [Ser73, Ch. IV]) and the Weil Conjectures (see, e.g., [Wei49, FK88]). More to the point, the best previous complexity upper bound for $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^* \times \mathcal{P}\right))$ appears to be quadruply exponential, via an extension of Hensel's Lemma by Birch and McCann [BMc67] (see also [Gre74]).

The aforementioned real counter-part to Assertion (3) of Theorem 1.4 again presents interesting subtleties: simple tricks involving monomial changes of variables suffice to prove the dramatically sharper upper bound of $\text{FEAS}_{\mathbb{R}}(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*) \in \mathbf{NC}^1$ over \mathbb{R} [BRS09, Thm. 1.3], but these tricks are obstructed over \mathbb{Q}_p (see Example 1.7 below). Assertion (3) is thus much harder to prove and Proposition 1.3 shows that not much better is possible. Further speed-ups for detecting p -adic rational roots of n -variate $(n + 1)$ -nomials appear to hinge on a better understanding of the analogous problem over certain finite rings. In particular, by our development in Section 3.1, the truth of the following conjecture would imply $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{F}_{n,n+1}^*) \in \mathbf{P}$ for any fixed n .

Conjecture 1 *Suppose $\ell, n \in \mathbb{N}$ and p is a prime. Then $\text{FEAS}_{\mathbb{Z}/p^\ell\mathbb{Z}}(\mathcal{F}_{n,n+1}^*)$ admits a (deterministic) algorithm with complexity $(\log(p) + \ell + \text{size}(f))^{O(n)}$.*

Note that brute-force search easily attains a complexity bound of $p^{\ell n} \text{size}(f)^{O(1)}$. Also, basic group theory (see, e.g., Lemma 3.1 in Section 3 below) yields the $n = 1$ case of the conjecture. So the key difficulty in Conjecture 1 is the dependence on p^ℓ when $n \geq 2$.

While the real counter-part to Assertion (0) of Theorem 1.4 is easy to prove, $\text{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,3}) \in \mathbf{P}$ — a stronger real counter-part to Assertion (1) — was proved only recently [BRS09, Thm. 1.3] using linear forms in logarithms [Nes03]. It is thus worth noting that the proof of Assertion (1) in Section 5 uses linear forms in p -adic logarithms [Yu94] at a critical juncture, and suggests an approach to a significant speed-up.

Corollary 1.5 *Fix a prime p , let $\ell \geq 1$, and suppose $\text{FEAS}_{\mathbb{Z}/p^\ell\mathbb{Z}}(\mathcal{F}_{1,3})$ admits a (deterministic) algorithm³ with complexity $(\ell + \text{size}(f))^{O(1)}$. Then $\text{FEAS}_{\mathbb{Q}_p}(\mathcal{F}_{1,3}) \in \mathbf{P}$.*

The truth of the hypothesis to our corollary above appears to be an open question. (Note that brute-force search easily leads to an algorithm of complexity $p^\ell \text{size}(f)^{O(1)}$, so the main issue here is the dependence on ℓ .) Paraphrased in our notation, Erich Kaltofen asked in 2003 whether $\text{FEAS}_{\mathbb{Z}/p\mathbb{Z}}(\mathcal{F}_{1,3})$ admits a (deterministic) algorithm with complexity $(\log(p) + \text{size}(f))^{O(1)}$ [Kal03].⁴

Since trinomials are $(1 + 2)$ -nomials, it is worth noting the related real complexity upper bound $\text{FEAS}_{\mathbb{R}}(\mathcal{F}_{n,n+2}^*) \in \mathbf{P}$ for any fixed $n \in \mathbb{N}$ [BRS09, Thm. 1.3]. In fact, the proof there inspired our proof of Assertion (1) of Theorem 1.4, so it would be most interesting to extend our p -adic techniques here to the multivariate case.

Conjecture 2 *For any fixed $n \in \mathbb{N}$ and any prime p , we have $\text{FEAS}_{\mathbb{Q}_p}(\mathcal{F}_{n,n+2}^*) \in \mathbf{NP}$.*

As for general univariate polynomials, the best previous complexity upper bound for $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x_1] \times \mathcal{P})$ relative to the sparse input size appears to have been **EXPTIME** [MW99]. In particular, both $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{F}_{1,4} \times \mathcal{P}) \stackrel{?}{\in} \mathbf{NP}$ and $\text{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,4}) \stackrel{?}{\in} \mathbf{NP}$ are still open

³All algorithms discussed here are based on Turing machines [GJ79, Pap95].

⁴David A. Cox also independently asked Rojas the same question in august of 2004.

questions [BRS09, Sec. 1.2]. However, high probability speed-ups over \mathbb{R} paralleling Assertion (2) are now known for $\text{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,4})$ [BHPR11, Thm. 1.4]. For clarity, here is an example illustrating Assertion (2).

Example 1.6 *Let T denote the family of pairs $(f, p) \in \mathbb{Z}[x_1] \times \mathcal{P}$ with $f(x_1) = a + bx_1^{11} + cx_1^{17} + x_1^{31}$ and let $T^* := T \setminus \mathcal{E}$ where \mathcal{E} is the set of pairs (f, p) where f does not admit a succinct certificate for p -adic rational feasibility. According to Assertion (2) of Theorem 1.4, \mathcal{E} has natural density 0. More precisely, there is a sparse 61×61 structured matrix \mathcal{S} (cf. Lemma 4.3 in Section 4 below), whose entries lie in $\{0, 1, 31, a, b, 11b, c, 17c\}$, such that $(f, p) \in T^* \iff p \nmid \det \mathcal{S}$. So $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(T^*) \in \mathbf{NP}$ and Corollary 4.6 in Section 4 below then tells us that for large coefficients, T^* occupies almost all of T . In particular, letting $T(H)$ (resp. $T^*(H)$) denote those pairs (f, p) in T (resp. T^*) with $|a|, |b|, |c|, p \leq H$, we obtain*

$$\frac{\#T^*(H)}{\#T(H)} \geq \left(1 - \frac{244}{2H+1}\right) \left(1 - \frac{1+61 \log(4H) \log H}{H}\right).$$

In particular, one can check via Maple that

$$(-973 + 21x_1^{11} - 2x_1^{17} + x_1^{31}, p) \in T^*$$

for all but 352 primes p . \diamond

1.2 Related Work, a Topological Observation, Weil's Conjecture, and Primes in Arithmetic Progression

Let us first recall that Emil Artin conjectured around 1935 that, for any prime p , homogeneous polynomials of degree d in $n > d^2$ variables always have non-trivial roots in \mathbb{Q}_p^n [Art65]. (The polynomials $x_1^2 + \dots + x_n^2$ show that Artin's conjecture is resoundingly false over the real numbers.) Artin's conjecture was already known to be true for $d = 2$ [Has24] and, in 1952, the $d = 3$ case was proved by Lewis [Lew52]. However, in 1966, Terjanian disproved the conjecture via an example with $(p, d, n) = (2, 4, 18)$.

The Ax-Kochen Theorem from 1965 provided a valid correction of Artin's conjecture: for any d , there is a constant p_d such that for all primes $p > p_d$, any homogeneous degree d polynomial in $n > d^2$ variables has a p -adic rational root [AK65, H-B10]. The hard cases of $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}$ thus appear to consist of high degree polynomials with few variables and p small.

It is interesting to observe that while it is easier for a polynomial in many variables to have roots over \mathbb{Q}_p than over \mathbb{R} , deciding the existence of roots appears to be much harder over \mathbb{Q}_p than over \mathbb{R} . In particular, while Tarski showed in 1939 that $\text{FEAS}_{\mathbb{R}}$ is decidable [Tar51], $\text{FEAS}_{\mathbb{Q}_p}$ wasn't shown to be decidable until work of Cohen in the 1960s [Coh69]. Now, the best general complexity upper bounds appear to be \mathbf{PSPACE} for $\text{FEAS}_{\mathbb{R}}$ [Can88] and *quadruply* exponential for $\text{FEAS}_{\mathbb{Q}_p}$ [BMc67, Gre74].

While the univariate problems $\text{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,2})$ and $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{F}_{1,2})$ are now both known to be in \mathbf{P} , their natural multivariate extensions $\text{FEAS}_{\mathbb{R}}(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*)$ and $\text{FEAS}_{\mathbb{Q}_p}(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*)$ already carry nuances distinguishing the real and p -adic settings: topological differences between the real and p -adic zero sets of polynomials in $\mathcal{F}_{n,n+1}^*$ force the underlying feasibility algorithms to differ. Concretely, positive zero sets for polynomials in $\mathcal{F}_{n,n+1}^*$ are always either empty or non-compact. This in turn allows one to solve $\text{FEAS}_{\mathbb{R}}(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*)$ by simply checking signs of coefficients, independent of the exponents [BRS09, Thm. 1.3]. On the other hand, solving $\text{FEAS}_{\mathbb{Q}_p}(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n,n+1}^*)$ depends critically on the exponents (see Corollary

3.2 of Section 3), and the underlying hypersurfaces in \mathbb{Q}_p^n can sometimes consist of just a single isolated point.

Example 1.7 Consider $f(x_1, x_2) := 1 + 2x_1^2 - 3x_2^2$. Then it is easy to see that $(1, 1)$ is the unique root of f in \mathbb{F}_7^2 . Via Hensel's Lemma (see Section 2 below), the root $(1, 1) \in \mathbb{F}_7^2$ can then be lifted to a unique root of f in \mathbb{Q}_7^2 . In particular, by checking valuations, any root of f in \mathbb{Q}_7^2 must be the lift of some root of f in \mathbb{F}_7^2 , and thus $(1, 1)$ is the only root of f in \mathbb{Q}_7^2 . \diamond

Our last example illustrates the importance of finite fields in studying p -adic rational roots. Deligne's Theorem on zeta functions over finite fields (née the Weil Conjectures) is the definitive statement on the connection between point counts over finite fields and complex geometry (see, e.g., [FK88]). The central result that originally motivated the Weil Conjectures will also prove useful in our study of $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}$.

Theorem 1.8 [Wei49, Pg. 502] Let p be any prime, $d_1, \dots, d_n \in \mathbb{N}$, and let c_0, \dots, c_n be integers not divisible by p . Then, defining $f(x) := c_0 + c_1x_1^{d_1} + \dots + c_nx_n^{d_n}$, the number, N , of roots of f in \mathbb{F}_p^n satisfies $|N - p^{n-1}| \leq (\prod_{i=1}^n (\gcd(d_i, p-1) - 1)) p^{(n-1)/2}$. ■

Finally, it is worth noting that our univariate **NP**-hardness proof requires the efficient construction of primes in certain arithmetic progressions. The following result, inspired by earlier work of von zur Gathen, Karpinski, and Shparlinski (see [vzGKS96, Fact 4.9]), may be of independent interest.

Theorem 1.9 For any $\delta > 0$, $\varepsilon \in (0, 1/2)$, and $n \in \mathbb{N}$, we can find — within

$$O\left((n/\varepsilon)^{\frac{3}{2}+\delta} + (n \log(n) + \log(1/\varepsilon))^{7+\delta}\right)$$

randomized bit operations — a sequence $P = (p_i)_{i=1}^n$ of consecutive primes and $c \in \mathbb{N}$ such that $p := 1 + c \prod_{i=1}^n p_i$ satisfies $\log p = O(n \log(n) + \log(1/\varepsilon))$ and, with probability $\geq 1 - \varepsilon$, p is prime.

2 Complexity Classes and p -adic Basics

Let us first recall briefly the following complexity classes (see also [Pap95] for an excellent textbook treatment):

NC¹ The family of functions computable by Boolean circuits with size polynomial⁵ in the input size and depth $O(\log^i \text{InputSize})$.

ZPP The family of decision problems admitting a randomized polynomial-time algorithm giving a correct answer, or a report of failure, the latter occurring with probability $\leq \frac{1}{2}$. Such algorithms are frequently called *Las Vegas* algorithms because one is never cheated (by a false answer, when an answer is given).⁶

⁵Note that the underlying polynomial depends only on the problem in question (e.g., matrix inversion, shortest path finding, primality detection) and not the particular instance of the problem.

⁶This now classical appellation likely involves some regional bias.

PSPACE The family of decision problems solvable within time polynomial in the input size, provided a number of processors exponential in the input size is allowed.

The following containments are standard:

$$\mathbf{NC}^1 \subseteq \mathbf{P} \subseteq \mathbf{ZPP} \subseteq \mathbf{NP} \subseteq \mathbf{PSPACE} \subseteq \mathbf{EXPTIME}.$$

The properness of each adjacent inclusion above (and even the properness of $\mathbf{P} \subseteq \mathbf{PSPACE}$) is a major open problem [Pap95].

Recall that for any ring R , we denote its unit group by R^* . For any prime p and $x \in \mathbb{Z}$, recall that the p -adic valuation, $\text{ord}_p x$, is the greatest k such that $p^k | x$. We can extend $\text{ord}_p(\cdot)$ to \mathbb{Q} by $\text{ord}_p \frac{a}{b} := \text{ord}_p(a) - \text{ord}_p(b)$ for any $a, b \in \mathbb{Z}$; and we let $|x|_p := p^{-\text{ord}_p x}$ denote the p -adic norm.

Recall also that, in any ring, x^n can be computed using just $O(\log n)$ bit operations and multiplication of powers of x , via recursive squaring [BS96, Thm. 5.4.1, pg. 103]. By considering the smallest k for which p^{2^k} divides an $x \in \mathbb{Z}$, and then repeating this calculation for $\frac{x}{p^{2^k}}$ (employing recursive squaring along the way), one can then compute $\text{ord}_p x$ efficiently. Recalling also that arithmetic in finite rings can be done efficiently [BS96, Ch. 5], we have the following statement:

Proposition 2.1 *Suppose p is any prime, $\ell \in \mathbb{N}$, and $x \in \mathbb{Z}$. Then we can compute $\text{ord}_p x$ in time polynomial in $\text{size}(x)$. Furthermore, all field operations in $\mathbb{Z}/p^\ell \mathbb{Z}$ can be done within a number of bit operations polynomial in $\log(p^\ell)$. ■*

The norm $|\cdot|_p$ defines a natural metric satisfying the ultrametric inequality and \mathbb{Q}_p is, to put it tersely, the completion of \mathbb{Q} with respect to this metric. This metric, along with $\text{ord}_p(\cdot)$, extends naturally to the field of p -adic complex numbers \mathbb{C}_p , which is the metric completion of the algebraic closure of \mathbb{Q}_p [Rob00, Ch. 3].

It will be useful to recall some classical invariants for treating quadratic polynomials over \mathbb{Q}_p .

Definition 2.2 [Ser73, Ch. I–IV, pp. 3–39] *For any prime p and $a \in \mathbb{Z}$ we define the Legendre symbol, $\left(\frac{a}{p}\right)$, to be $+1$ or -1 according as a has a square root mod p or not. Also, for any $b \in \mathbb{Z}$, we let the (p -adic) Hilbert symbol, $(a, b)_p$, be $+1$ or -1 according as $ax^2 + by^2 = z^2$ has a solution in $\mathbb{P}_{\mathbb{Q}_p}^2$ or not. Finally, for any $f(x) = c_0 + c_1 x_1^2 + \dots + c_n x_n^2 \in \mathbb{Z}[x_1, \dots, x_n]$, we define $d_f := \prod_{i=1}^n c_i$ and $\varepsilon_f := \prod_{1 \leq i < j \leq n} (c_i, c_j)_p$. ◇*

Theorem 2.3 [Ser73, Thm. 1, pg. 20 & Cor., pp. 37] *Following the notation of Definition 2.2, let $j := \text{ord}_p a$ and $k := \text{ord}_p b$. Then the Hilbert symbol $(a, b)_p$ is exactly*

$$(i) \quad (-1)^{jk \left(\frac{p-1}{2} \bmod 2\right)} \left(\frac{a/p^j}{p}\right)^k \left(\frac{b/p^k}{p}\right)^j, \text{ or}$$

$$(ii) \quad (-1)^{Z(a,b)} \text{ where } Z(a,b) := \left(\frac{a/2^j-1}{2}\right) \left(\frac{b/2^k-1}{2}\right) + j \left(\frac{(b/2^k)^2-1}{8}\right) + k \left(\frac{(a/2^j)^2-1}{8}\right) \bmod 2,$$

according as $p \neq 2$ or $p = 2$.

Finally, f has a root in \mathbb{Q}_p if and only if one of the following conditions holds:

1. $n = 1$, $\mu := \text{ord}_p(c_0/c_1)$ is even, and $\left(\frac{-c_0/(c_1 p^\mu)}{p}\right) = 1$.
2. $n = 2$ and $(-c_0, -d_f)_p = \varepsilon_f$ (viewing c_0 and d_f as elements of $\mathbb{Q}_p/(\mathbb{Q}_p^*)^2$).
3. $n = 3$ and either $c_0 \neq d_f$ or $c_0 = d_f$ and $(-1, -d_f) = \varepsilon_f$ (viewing c_0 and d_f as elements of $\mathbb{Q}_p/(\mathbb{Q}_p^*)^2$).
4. $n \geq 4$. ■

A key tool we will use throughout this paper is *Hensel's Lemma*, suitably extended to multivariate Laurent polynomials.

Hensel's Lemma *Suppose $f \in \mathbb{Z}_p[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$ and $\zeta_0 \in \mathbb{Z}_p^n$ satisfies $\text{ord}_p \frac{\partial f}{\partial x_i}(\zeta_0) = \ell < +\infty$ for some $i \in \{1, \dots, n\}$, and $f(\zeta_0) \equiv 0 \pmod{p^{2\ell+1}}$. Then there is a root $\zeta \in \mathbb{Z}_p^n$ of f with $\zeta \equiv \zeta_0 \pmod{p^\ell}$ and $\text{ord}_p \frac{\partial f}{\partial x_i}(\zeta) = \text{ord}_p \frac{\partial f}{\partial x_i}(\zeta_0)$. ■*

The special case of polynomials appears as Theorem 1 on the bottom of Page 14 of [Ser73]. (See also [BMc67].) The proof there extends almost verbatim to Laurent polynomials.

3 From Binomials to $(n + 1)$ -nomials: Proving Assertions (0) and (3) of Theorem 1.4

Let us first recall the following standard lemma on taking radicals in certain finite groups.

Lemma 3.1 *(See, e.g., [BS96, Thm. 5.7.2 & Thm. 5.6.2, pg. 109]) Given any cyclic group G , $a \in G$, and an integer d , the following 3 conditions are equivalent:*

1. *The equation $t^d = a$ has a solution $t \in G$.*
2. *The order of a divides $\frac{\#G}{\gcd(d, \#G)}$.*
3. *$a^{\#G/\gcd(d, \#G)} = 1$.*

Also, \mathbb{F}_q^ is cyclic for any prime power q , and $(\mathbb{Z}/p^\ell\mathbb{Z})^*$ is cyclic for any (p, ℓ) with p an odd prime or $\ell \leq 2$. Finally, for $\ell \geq 3$, $(\mathbb{Z}/2^\ell\mathbb{Z})^* = \{\pm 1, \pm 5, \pm 5^2, \pm 5^3, \dots, \pm 5^{2^{\ell-2}-1} \pmod{2^\ell}\}$. ■*

A direct consequence of Lemma 3.1 and Hensel's Lemma is the following characterization of univariate binomials with p -adic rational roots.

Corollary 3.2 *Suppose $c \in \mathbb{Q}_p^*$ and $d \in \mathbb{Z} \setminus \{0\}$. Let $k := \text{ord}_p c$, $\ell := \text{ord}_p d$, and (if $p = 2$ and d is even) $d' = \left(\frac{d}{2^\ell}\right)^{-1} \pmod{2^{2\ell-1}}$. Then the equation $x^d = c$ has a solution in \mathbb{Q}_p iff $d \mid \text{ord}_p c$ and one of the following two conditions hold:*

(a) *p is odd and $\left(\frac{c}{p^k}\right)^{p^{\ell(p-1)/\gcd(d, p-1)}} = 1 \pmod{p^{2\ell+1}}$.*

(b) *$p = 2$ and either (i) d is odd, or (ii) $\left(\frac{c}{p^k}\right)^{d'} = 1 \pmod{8}$ and $\left(\frac{c}{p^k}\right)^{d'2^{\max\{\ell-2, 0\}}} = 1 \pmod{2^{2\ell+1}}$.*

In particular, these conditions can be checked in time polynomial in $\log(d) + \log(p)$ when $\log c = (\log(d) + \log(p))^{O(1)}$. Furthermore, when $\text{ord}_p c = 0$, $x^d = c$ has a root in \mathbb{Q}_p if and only if $x^d = c$ has a root in $(\mathbb{Z}/p^{2\ell+1}\mathbb{Z})^$.*

Proof: Replacing x by $1/x$, we can clearly assume $d > 0$. Clearly, any p -adic root ζ of $x^d - c$ satisfies $d \text{ord}_p \zeta = \text{ord}_p c$. This accounts for the condition preceding Conditions (a) and (b).

Replacing x by $p^{\text{ord}_p c/d} x$ (which clearly preserves the existence of roots in \mathbb{Q}_p^*) we can assume further that $\text{ord}_p c = \text{ord}_p \zeta = 0$. Moreover, $\text{ord}_p f'(\zeta) = \text{ord}_p(d) + (d-1)\text{ord}_p \zeta = \text{ord}_p d$. So by Hensel's Lemma, $x^d - c$ has a root in \mathbb{Q}_p^* if and only if $x^d - c$ has a root in $(\mathbb{Z}/p^{2\ell+1}\mathbb{Z})^*$. Lemma 3.1 then tells us the order of the group $(\mathbb{Z}/p^{2\ell+1}\mathbb{Z})^*$, and how to decide solvability therein, thus accounting for Condition (a) when p is odd.

Condition (b) then follows routinely: First, one observes that exponentiating by an odd power is an automorphism of $(\mathbb{Z}/2^{2\ell+1}\mathbb{Z})^*$, and thus $x^d - c$ has a root in $(\mathbb{Z}/2^{2\ell+1}\mathbb{Z})^*$ if and

only if $x^{2^\ell} - c^{d'}$ does. Should $\ell=0$ then one has a root regardless of c . Otherwise, $c^{d'}$ must be a square for there to be a root. Since $\text{ord}_2 c=0$, c is odd and [BS96, Ex. 38, pg. 192] tells us that $c^{d'}$ is a square in $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ if and only if $c^{d'}=1 \pmod{8}$. Invoking Lemma 3.1 once more on the the cyclic subgroup $\{1, 5^2, 5^4, 5^6, \dots, 5^{2^{\ell-1}-2}\}$, it is clear that Condition (b) is exactly what we need when $p=2$.

The asserted time bound then follows immediately from Proposition 2.1. The final assertion follows immediately from setting $k=0$ in the conditions we've just derived. ■

At this point, the proof of Assertion (0) of Theorem 1.4 is trivial. By combining our last result with a classical integral matrix factorization, Assertion (3) then also becomes easy to prove. So let us first motivate the connection between n -variate $(n+1)$ -nomials and matrices.

Proposition 3.3 *Suppose K is any field of characteristic 0 and f is any honest n -variate $(n+1)$ -nomial over K of the form $f(x) := c_0 + c_1x^{a_1} + \dots + c_nx^{a_n}$. Then, letting A denote the matrix whose columns are a_1, \dots, a_n , letting $x = (x_1, \dots, x_n) \in (K^*)^n$, and defining $f_i := \frac{\partial f}{\partial x_i}$ for all i , we have:*

$$[f_1(x), \dots, f_n(x)] = [c_1x^{a_1}, \dots, c_nx^{a_n}]A^T \begin{bmatrix} x_1^{-1} & & & \\ & \ddots & & \\ & & \ddots & \\ & & & x_n^{-1} \end{bmatrix}.$$

In particular, all the roots of f in $(K^)^n$ are non-degenerate.*

Proof: The first assertion is routine. For the second assertion, observe that if $\zeta \in (K^*)^n$ is any root of f then, thanks to our first assertion, the vector $[f_1(\zeta), \dots, f_n(\zeta)]$ can not vanish because $\det A \neq 0$ (thanks to the condition of honesty). ■

Definition 3.4 *Let $\mathbb{Z}^{n \times n}$ denote the set of $n \times n$ matrices with all entries integral, and let $\mathbf{GL}_n(\mathbb{Z})$ denote the set of all matrices in $\mathbb{Z}^{n \times n}$ with determinant ± 1 (the set of unimodular matrices). Recall that any $n \times n$ matrix $[u_{ij}]$ with $u_{ij}=0$ for all $i > j$ is called upper triangular.*

Given any $M \in \mathbb{Z}^{n \times n}$, we then call an identity of the form $UM = H$, with $H = [h_{ij}] \in \mathbb{Z}^{n \times n}$ in row echelon form and $U \in \mathbf{GL}_n(\mathbb{Z})$, a Hermite factorization of M . Also, if we have the following conditions in addition:

1. *the left-most nonzero entry in any row of H is positive.*
2. *for any i , $h_{i,j}$ the left-most nonzero entry of row $i \implies 0 \leq h_{i',j} < h_{i,j}$ for all $i' < i$.*

then we call H the Hermite normal form of M .

Also, given any identity of the form $UMV = S$ with $U, V \in \mathbf{GL}_n(\mathbb{Z})$ and S diagonal a Smith factorization. In particular, if $S = [s_{i,j}]$ and we require additionally that $s_{i,i} \geq 0$ and $s_{i,i} | s_{i+1,i+1}$ for all $i \in \{1, \dots, n\}$ (setting $s_{n+1,n+1} := 0$), then S is uniquely determined and is called the Smith normal form of M .

Finally, defining $x^A = (x_1^{a_{1,1}} \dots x_n^{a_{n,1}}, \dots, x_1^{a_{1,n}} \dots x_n^{a_{n,n}})$, we call any map defined by $x \mapsto x^A$ a monomial change of variables. \diamond

Proposition 3.5 *We have that $x^{AB} = (x^A)^B$ for any $A, B \in \mathbb{Z}^{n \times n}$. Also, for any field K , the map defined by $m(x) = x^U$, for any unimodular matrix $U \in \mathbb{Z}^{n \times n}$, is an automorphism of $(K^*)^n$. Finally, for any column vector $v \in \mathbb{Z}^n$, the smallest valuation of an entry of Uv is k if and only if the smallest valuation of an entry of v is k . ■*

Theorem 3.6 [Sto00, Ch. 6 & 8, pg. 128] For any $A = [a_{i,j}] \in \mathbb{Z}^{n \times n}$, the Hermite and Smith factorizations of A can be computed within $O(n^{3.376} \log^2(n \max_{i,j} |a_{i,j}|))$ bit operations. Furthermore, the entries of all matrices in the Hermite and Smith factorizations have bit size $O(n \log(n \max_{i,j} |a_{i,j}|))$. ■

Lemma 3.7 Following the notation of Definition 3.4 and Proposition 3.5, suppose $\det A \neq 0$, $c_1, \dots, c_n \in \mathbb{Q}_p^*$, $c := (c_1, \dots, c_n)$, $c' := (c'_1, \dots, c'_n) := \left(\frac{c_1}{p^{\text{ord}_p c_1}}, \dots, \frac{c_n}{p^{\text{ord}_p c_n}}\right)$, $L := \max_i \text{ord}_p s_{i,i}$, and let v_1, \dots, v_n be the columns of V . Then $x^A = c$ has a solution in $(\mathbb{Q}_p^*)^n$ if and only if (a) $(\text{ord}_p c_1, \dots, \text{ord}_p c_n)v_i = 0 \pmod{s_{i,i}}$ for all i and (b) $x^A = c'$ has a solution in $((\mathbb{Z}/p^{2L+1})^*)^n$. In particular, the existence of a solution in $(\mathbb{Q}_p^*)^n$ for $x^A = c$ can be decided in time polynomial in n and $\log(n \max_{i,j} |a_{i,j}|)$.

Proof: The necessity of Condition (a) follows immediately from Proposition 3.5 upon observing that the valuations of the vector x^A are exactly the entries of $[\text{ord}_p x_1, \dots, \text{ord}_p x_n]A$. Conversely, should Condition (a) hold, we can reduce to the case where $\text{ord}_p c_i = 0$ for all i . So let us assume the last condition.

Observe now that $x^A = c$ if and only if $x^{AV} = c'$. Upon substituting $x := y^U$, we see that the latter equation holds if and only if $y^{UAV} = c^V$. In other words, $y^S = c^V$. By Proposition 3.5, the last system has a solution in $(\mathbb{Q}_p^*)^n$ if and only if the first system does. By Corollary 3.2 we thus see that Condition (b) is necessary and sufficient.

To prove the asserted complexity bound, note that we can find U , V , and S within the asserted time bound, thanks to Theorem 3.6. Note also that we can find the p -parts of the $s_{i,i}$ in polynomial time (by Proposition 2.1) so we can compute L in polynomial-time. Applying Corollary 3.2 n times, we can then decide in \mathbf{P} whether $y^S = c^V$ has a root in $(\mathbb{Q}_p^*)^n$. ■

A final ingredient we will need is a method to turn roots of Laurent polynomials on coordinate subspaces to roots in the algebraic torus.

Lemma 3.8 Suppose $\bar{f} \in \mathbb{Q}_p[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$, $c \in \mathbb{Q}_p^*$, $\alpha := (\alpha_1, \dots, \alpha_{n+1}) \in \mathbb{Z}^{n+1}$ with $\alpha_{n+1} > 0$, and \bar{f} has a non-degenerate root in $(\mathbb{Q}_p^*)^n$ (resp. $(\mathbb{Z}_p \setminus \{0\})^n$). Then the Laurent polynomial $f(x_1, \dots, x_n, x_{n+1}) := \bar{f}(x_1, \dots, x_n) + cx^\alpha$ has a non-degenerate root of f in $(\mathbb{Q}_p^*)^{n+1}$ (resp. $(\mathbb{Z}_p \setminus \{0\})^{n+1}$). ■

Proof: First note that the roots of \bar{f} (resp. f) in $(\mathbb{Q}_p^*)^n$ (resp. $(\mathbb{Q}_p^*)^{n+1}$) are unaffected if \bar{f} (resp. f) is multiplied by any monomial. Furthermore, if $\alpha_{n+1} = 1$ then the zero set of f is (up to rescaling) merely the graph of \bar{f} and the lemma follows immediately from the Schwartz-Zippel Lemma: pick any $(\zeta_1, \dots, \zeta_n) \in (\mathbb{Q}_p^*)^n$ with $\bar{f}(\zeta_1, \dots, \zeta_n) \neq 0$ and then the desired non-degenerate root of f is simply $\left(\zeta_1, \dots, \zeta_n, -\frac{\bar{f}(\zeta_1, \dots, \zeta_n)}{c \zeta_1^{\alpha_1} \dots \zeta_n^{\alpha_n}}\right)$. So we may also assume $\alpha_{n+1} \geq 2$.

Let $\zeta = (\zeta_1, \dots, \zeta_n) \in (\mathbb{Q}_p^*)^n$ be the stated non-degenerate root of \bar{f} . Observe that $\frac{\partial \bar{f}}{\partial x_i}(\zeta_1, \dots, \zeta_n) = \frac{\partial f}{\partial x_i}(\zeta_1, \dots, \zeta_n, 0)$ for all $i \in \{1, \dots, n\}$. So $(\zeta, 0)$ is a non-degenerate root of f . By the Implicit Function Theorem for analytic (i.e., C^∞) functions over \mathbb{Q}_p^{n+1} [Glo06, Thm. 7.4, pg. 237], there must then be a (non-degenerate) root $(\zeta'_1, \dots, \zeta'_k, p^\ell)$ of f for any sufficiently large $\ell \in \mathbb{N}$, with $\zeta'_i \rightarrow \zeta_i$ for all $i \in \{1, \dots, k\}$ as $\ell \rightarrow +\infty$. Thus, we can find a non-degenerate root of f in $(\mathbb{Q}_p^*)^{n+1}$. In conclusion, note that if ζ were in $(\mathbb{Z}_p \setminus \{0\})^n$ then the same argument yields a root in $(\mathbb{Z}_p \setminus \{0\})^{n+1}$. ■

Remark 3.9 Note that Example 1.7 from Section 1.2 shows that the converse of Lemma 3.8 need not hold. On the other hand, for honest n -variate $(n + 1)$ -nomials over the real numbers, both the corresponding analogue of Lemma 3.8 and its converse do hold [BRS09, Cor. 2.6]. \diamond

Henceforth, we will let \mathbf{O} denote the origin in whatever vector space we are working with.

Definition 3.10 Suppose K is a field, $f \in K[x_1^{\pm 1}, \dots, x_n^{\pm 1}]$, $\text{Supp}(f) = \{a_1, \dots, a_m\}$ has cardinality m , the coefficient of x^{a_i} in f is c_i for all i , and $w \in \mathbb{R}^n$. Let $\text{Supp}(f)^w$ denote the intersection of $\text{Supp}(f)$ with the face of $\text{Newt}(f)$ with inner normal w . We then define the initial term polynomial of f with respect to the weight w to be $\text{Init}_w(f) := \sum_{a_i \in \text{Supp}(f)^w} c_i x^{a_i}$. \diamond

Initial term polynomials are a natural (and classical) generalization of the lowest (or highest) degree part of a polynomial. As another example, following the notation above, suppose $J \subseteq \{1, \dots, n\}$ is such that, for all $j \in J$, the j^{th} coordinates of a_1, \dots, a_m are all nonnegative. Then substituting $x_j = 0$ into f for all $j \in J$ results in an initial term polynomial of f . In particular, f is an initial term polynomial of f (using the weight \mathbf{O}).

Corollary 3.11 Suppose f is an honest n -variate $(n + 1)$ -nomial over \mathbb{Q}_p . Then f has a root in $(\mathbb{Q}_p^*)^n$ if and only if some initial term polynomial of f with at least 2 terms has a root in $(\mathbb{Q}_p^*)^n$.

Proof: The (\implies) direction is trivial since f is an initial term polynomial by default. So let us focus on the (\impliedby) direction.

Since the roots of f (and any of its initial term polynomials) in $(\mathbb{Q}_p^*)^n$ are unaffected by multiplying by monomials, we can write $f(x) = c_0 + c_1 x^{a_1} + \dots + c_n x^{a_n}$ with $c_0, \dots, c_n \in \mathbb{Q}_p^*$, and also assume (reordering terms if need be) that the initial term polynomial from the hypothesis is of the form $\bar{f}(x) = c_0 + c_1 x^{a_1} + \dots + c_r x^{a_r}$ for some $r < n$. By Proposition 3.5, $\bar{f}(x)$ has a root in $(\mathbb{Q}_p^*)^n$ if and only if $\bar{f}(x^U)$ has a root in $(\mathbb{Q}_p^*)^n$ (and likewise for f). So via the Hermite Factorization, we may assume that the matrix A whose columns are a_1, \dots, a_n is upper-triangular. In other words, we may also assume that \bar{f} is independent of x_{r+1}, \dots, x_n , and that the $(r + 1)^{\text{st}}$ coordinate of a_{r+1} is positive. Letting $(\zeta_1, \dots, \zeta_n) \in (\mathbb{Q}_p^*)^n$ denote the non-degenerate root of \bar{f} from the hypothesis, we then obtain that $(\zeta_1, \dots, \zeta_r, 0, \dots, 0)$ is also a non-degenerate root of \bar{f} . By Lemma 3.8 (and induction) we then obtain that f must have a root in $(\mathbb{Q}_p^*)^n$. \blacksquare

3.1 The Proofs of Assertions (0) and (3) of Theorem 1.4

Assertion (0): First note that the case $m \leq 1$ is trivial: such a univariate m -nomial has no roots in \mathbb{Q}_p if and only if it is a nonzero constant.

The case $m = 2$ then follows immediately from Corollary 3.2. \blacksquare

Assertion (3):

Part (a): First note that if $\zeta = (\zeta_1, \dots, \zeta_n) \in \mathbb{Q}_p^n$ is a root of f then, modulo a permutation of coordinates, $\zeta \in (\mathbb{Q}_p^*)^r \times \{0\}^{n-r}$ for some r . In particular, ζ thus induces a root in $(\mathbb{Q}_p^*)^n$ of some initial term polynomial \bar{f} of f (upon replacing $\{0\}^{n-r}$ by, say, $\{1\}^{n-r}$). Furthermore,

by Corollary 3.11, such a root of \bar{f} in turn induces a root of f in $(\mathbb{Q}_p^*)^n$. So it suffices to certify the existence of a root in $(\mathbb{Q}_p^*)^n$ for some initial term polynomial of \bar{f} of f .

As before, since the roots of f (and its initial term polynomials) are unaffected by multiplying by monomials, we may assume that \bar{f} has a nonzero constant term. (Note also that enforcing this assumption induces at worst a factor of 2 growth in the absolute values of the entries of A .) Furthermore, by Theorem 3.6, we may assume that $\bar{f} \in \mathcal{F}_{r,r+1}^*$ for some $r \leq n$. So let us assume without loss of generality that $r = n$, $\bar{f}(x) = f(x) = c_0 + c_1 x^{\alpha_1} + \dots + c_n x^{\alpha_n}$, and the matrix A with columns a_1, \dots, a_n is upper-triangular.

Now set $L := \max_i \text{ord}_p(c_i) + \max_i \text{ord}_p s_{i,i}$ where the $s_{i,i}$ denote the diagonal entries of the Smith Normal Form of A . Our certificate for f having a root in $(\mathbb{Q}_p^*)^n$ will then be a root $\mu_0 \in (\mathbb{Z}/p^{2L+1}\mathbb{Z})^n \setminus \{\mathbf{0}\}$ of the mod p^{2L+1} reduction of $\bar{h}(x) := \bar{g}(x_1^{\pm 1}, \dots, x_n^{\pm 1})$, for some choice of reciprocals, where $\bar{g}(x) := x^{-\alpha_i} \bar{f}(x)$ for some i , and \bar{f} is an initial term polynomial of f with at least 2 terms. We will now show that f has a root $\zeta \in (\mathbb{Q}_p^*)^n$ if and only if a certificate of the preceding form exists.

To prove the (\implies) direction, let us first clarify the choice of reciprocals in $\bar{g}(x_1^{\pm 1}, \dots, x_n^{\pm 1})$: we place an exponent of -1 for all j where $\zeta_j \in \mathbb{Q}_p \setminus \mathbb{Z}_p$. Clearly then, with the preceding choice of reciprocals, $f(x_1^{\pm 1}, \dots, x_n^{\pm 1})$ has a root $\mu \in (\mathbb{Z}_p \setminus \{0\})^n$. The choice of i to define $\bar{h}(x)$ is also simple to pin down: pick any i with $\text{ord}_p(\mu^{\alpha_i})$ minimal. The roots of $h(x) := x^{-\alpha_i} f(x_1^{\pm 1}, \dots, x_n^{\pm 1})$ in $(\mathbb{Q}_p^*)^n$ are clearly independent of i .

To clarify the choice of \bar{f} let us first write $h(x) := \gamma_0 + \gamma_1 x^{\alpha_1} + \dots + \gamma_n x^{\alpha_n}$. The γ_i are then a re-ordering of the c_i , the α_i are differences of columns of A , and the matrix A' with columns $\alpha_1, \dots, \alpha_n$ is non-singular and has entries no larger in absolute value than twice those of A . We also have that $\text{ord}_p(\mu^{\alpha_i}) \geq 0$ for all i by construction. Moreover, by the ultrametric property (applied to the sum $\gamma_0 + (\gamma_1 \mu^{\alpha_1} + \dots + \gamma_n \mu^{\alpha_n})$), the root μ of h must satisfy $\text{ord}_p(\gamma_i \mu^{\alpha_i}) \leq \text{ord}_p \gamma_0 \leq \max_k \text{ord}_p c_k \leq L$ for some i . (Otherwise $\text{ord}_p h(\mu) = \text{ord}_p \gamma_0 < +\infty$). By Propositions 3.3 and 3.5, and the Smith factorization of the matrix A' , we must then have $\text{ord}_p h_j(\mu) \leq \text{ord}_p(\gamma_0) + \max_i \text{ord}_p(2s_{i,i}) \leq L = O(\text{size}(f))$ for some j , where $h_j = \frac{\partial h}{\partial x_j}$.

Clearly then, there are $u_{i_1}, \dots, u_{i_r} \in \mathbb{Z}_p \setminus \{0\}$ with $r \geq 1$, $L \geq \text{ord}_p u_{i_j} \geq \text{ord}_p \gamma_{i_j}$ for all j , $\gamma_0 + u_{i_1} + \dots + u_{i_r} = 0$, and $(\mu^{\alpha_{i_1}}, \dots, \mu^{\alpha_{i_r}}) = \left(\frac{u_{i_1}}{c_{i_1}}, \dots, \frac{u_{i_r}}{c_{i_r}}\right)$. So define \bar{f} to be the sum of terms of f corresponding to picking the i_1, \dots, i_r terms of h . By Lemma 3.7, μ then has a well-defined mod p^{2L+1} reduction $\mu_0 \in (\mathbb{Z}/p^{2L+1}\mathbb{Z})^n \setminus \{\mathbf{0}\}$ that is a root of the mod p^{2L+1} reduction of \bar{h} . So the (\implies) direction is proved.

To prove the (\impliedby) direction, let us suppose that the mod p^{2L+1} reduction of $\bar{h}(x) := \bar{g}(x_1^{\pm 1}, \dots, x_n^{\pm 1})$ has a root $\mu_0 \in (\mathbb{Z}/p^{2L+1}\mathbb{Z})^n \setminus \{\mathbf{0}\}$ for some choice of signs, some choice of i , and some choice of initial term polynomial \bar{f} of f so that $\bar{g}(x) = x^{-\alpha_i} \bar{f}(x)$. Writing $\bar{h}(x) = \gamma_0 + \gamma_{i_1} x^{\alpha_{i_1}} + \dots + \gamma_{i_r} x^{\alpha_{i_r}}$ as before, it is clear that $\text{ord}_p(\gamma_i \mu^{\alpha_i}) \leq \text{ord}_p \gamma_0$ for some i by the ultrametric inequality. So then, by Proposition 3.3, $\text{ord}_p \bar{h}'(\mu) \leq L$, and then by Hensel's Lemma, \bar{h} has a root $\mu' \in \mathbb{Z}_p^n \setminus \{\mathbf{0}\}$. By Corollary 3.11, $h(x) := \gamma_0 + \gamma_1 x^{\alpha_1} + \dots + \gamma_n x^{\alpha_n}$ must then have a root $\mu \in (\mathbb{Z}_p \setminus \{\mathbf{0}\})^n$. So by the definition of h , it is then clear that defining $\zeta_i = \mu_i^{\pm 1}$ for a suitable choice of signs, $\zeta := (\zeta_1, \dots, \zeta_n)$ is a root of f . ■

Part (b): Since the Legendre symbol $\left(\frac{a}{p}\right)$ can be evaluated within $O((\log a)(\log p))$ bit operations [BS96, Thm. 5.9.3, pg. 113], the criteria from Theorem 2.3 can clearly be checked in time polynomial in $\text{size}(f)$. ■

Part (c): By the succinct certificates we used to prove Part (a), the existence of a root of f in \mathbb{Q}_p^n is implied by the existence of a root of f in \mathbb{F}_p^n if $\text{ord}_p|c_0| = \cdots = \text{ord}_p|c_n| = \text{ord}_p(n!V_f) = 0$. By Theorem 1.8, a root for f in \mathbb{F}_p^n is guaranteed if $n \geq 2$, p does not divide any c_i , and $p \geq (n!V_f)^{2/(n-1)}$. ■

Remark 3.12 *The hypotheses of Theorem 1.8 clearly allow a slightly better lower bound for the p guaranteeing that f have a root in \mathbb{Q}_p^n . Also, it is likely that Assertion (c) remains true if f is any Laurent polynomial in $\mathcal{F}_{n,n+1}^*$. Proving this requires a refinement of Theorem 1.8 that, to the best of our knowledge, has not yet appeared in the literature. ◇*

4 Discriminants, p -adic Newton Polygons, and Assertion (2) of Theorem 1.4

The intuition behind the speed-up of Assertion (2) is that the hardness of instances of $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x_1] \times \mathcal{P})$ is governed by numerical conditioning, quite similar to the sense long known in numerical linear algebra (and extended more recently to real feasibility [CS99]). More concretely, the classical fact that Newton iteration converges more quickly for a root $\zeta \in \mathbb{C}$ of f with $f'(\zeta)$ having large norm (i.e., a *well-conditioned* root) persists over \mathbb{Q}_p .

To prepare for our next proof, let us first clarify the statement about natural density 0 in Assertion (2) of Theorem 1.4.

Definition 4.1 *Letting $\#$ denote set cardinality and $S \subseteq T \subseteq \mathbb{N}$, we say that S has (natural) density μ in T if and only if $\lim_{t \rightarrow \infty} \frac{\#S \cap \{1, \dots, t\}}{\#T \cap \{1, \dots, t\}} = \mu$. ◇*

Now let $(\mathbb{Z} \times (\mathbb{N} \cup \{0\}))^\infty$ denote the set of all infinite sequences of pairs $((c_i, a_i))_{i=1}^\infty$ with $c_i = a_i = 0$ for i sufficiently large. Note then that $\mathbb{Z}[x_1]$ admits a natural embedding into $(\mathbb{Z} \times (\mathbb{N} \cup \{0\}))^\infty$ by considering coefficient-exponent pairs in order of increasing exponents, e.g., $a + bx^{99} + x^{2001} \mapsto ((a, 0), (b, 99), (1, 2001), (0, 0), (0, 0), \dots)$. Then natural density for a set of pairs $\mathcal{I} \subseteq \mathbb{Z}[x_1] \times \mathcal{P}$ then simply means the corresponding natural density within $(\mathbb{Z} \times (\mathbb{N} \cup \{0\}))^\infty \times \mathcal{P}$.

The exceptional set to Assertion (2) can be made more precise once one introduces the \mathcal{A} -discriminant. But first we must introduce the resultant and some quantitative estimates.

Definition 4.2 *(See, e.g., [GKZ94, Ch. 12, Sec. 1, pp. 397–402].) Suppose $f(x_1) = a_0 + \cdots + a_d x_1^d$ and $g(x_1) = b_0 + \cdots + b_{d'} x_1^{d'}$ are polynomials with indeterminate coefficients. We define their Sylvester matrix to be the $(d + d') \times (d + d')$ matrix*

$$\mathcal{S}_{(d,d')}(f, g) := \left[\begin{array}{cccccc} a_0 & \cdots & a_d & 0 & \cdots & 0 \\ & & & & & \\ & & & & & \\ & & & & & \\ 0 & \cdots & 0 & a_0 & \cdots & a_d \\ b_0 & \cdots & b_{d'} & 0 & \cdots & 0 \\ & & & & & \\ & & & & & \\ & & & & & \\ 0 & \cdots & 0 & b_0 & \cdots & b_{d'} \end{array} \right] \left. \begin{array}{l} \right\} d' \text{ rows} \\ \left. \right\} d \text{ rows} \end{array}$$

and their Sylvester resultant to be $\text{Res}_{(d,d')}(f, g) := \det \mathcal{S}_{(d,d')}(f, g)$. ◇

Lemma 4.3 *Following the notation of Definition 4.2, assume $f, g \in K[x_1]$ for some field K , and that a_d and $b_{d'}$ are not both 0. Then $f = g = 0$ has a root in the algebraic closure of K if and only if $\text{Res}_{(d,d')}(f, g) = 0$. More generally, we have $\text{Res}_{(d,d')}(f, g) = a_d^{d'} \prod_{f(\zeta)=0} g(\zeta)$ where the product counts multiplicity. Finally, if we assume further that f and g have complex coefficients of absolute value $\leq H$, and f (resp. g) has exactly m (resp. m') monomial terms, then $|\text{Res}_{(d,d')}(f, g)| \leq m^{d'/2} m'^{d/2} H^{d+d'}$. ■*

The first 2 assertions are classical (see, e.g., [BPR06, Thm. 4.16, pg. 107], [RS02, pg. 9], and [GKZ94, Ch. 12, Sec. 1, pp. 397–402]). The last assertion follows easily from Hadamard's Inequality (see, e.g., [Mig82, Thm. 1, pg. 259]).

We are now ready to introduce discriminants.

Definition 4.4 *Let $\mathcal{A} := \{a_1, \dots, a_m\} \subset \mathbb{N} \cup \{0\}$ and $f(x_1) := \sum_{i=1}^m c_i x_1^{a_i}$, where $0 \leq a_1 < \dots < a_m$ and the c_i are algebraically independent indeterminates. We then define the (normalized) \mathcal{A} -discriminant of f , $\bar{\Delta}_{\mathcal{A}}(f)$, to be*

$$\text{Res}_{(\bar{a}_m - \bar{a}_2, \bar{a}_m)} \left(\frac{\partial \bar{f}}{\partial x_1} / x_1^{\bar{a}_2 - 1}, \bar{f} \right) / c_m^{\bar{a}_m - \bar{a}_m - 1},$$

where $\bar{a}_i := (a_i - a_1)/g$ for all i , $\bar{f}(x_1) := \sum_{i=1}^m c_i x_1^{\bar{a}_i}$, and $g := \gcd(a_2 - a_1, \dots, a_m - a_1)$ (see also [GKZ94, Ch. 12, pp. 403–408]). \diamond

Remark 4.5 *Note that when $\mathcal{A} = \{0, \dots, d\}$ we have*

$$\bar{\Delta}_{\mathcal{A}}(f) = \text{Res}_{(d-1, d)}(f', f) / c_d = (-1)^{a_3(a_3 - a_2)} \text{Res}_{(d, d-1)}(f, f') / c_d,$$

i.e., for dense polynomials, the \mathcal{A} -discriminant agrees with the classical discriminant, written $\Delta(f)$ in [GKZ94, Ch. 12], up to an explicit sign factor. \diamond

The claim of natural density 0 in Assertion (2) of Theorem 1.4 can then be made explicit as follows.

Corollary 4.6 *For any subset $\mathcal{A} = \{a_1, \dots, a_m\} \subset \mathbb{N} \cup \{0\}$ with $0 = a_1 < \dots < a_m$, let $T_{\mathcal{A}}$ denote the family of pairs $(f, p) \in \mathbb{Z}[x_1] \times \mathcal{P}$ with $f(x_1) = \sum_{i=1}^m c_i x_1^{a_i}$ and let $T_{\mathcal{A}}^*$ denote the subset of $T_{\mathcal{A}}$ consisting of those pairs (f, p) with $p \nmid \bar{\Delta}_{\mathcal{A}}(f)$. Also let $T_{\mathcal{A}}(H)$ (resp. $T_{\mathcal{A}}^*(H)$) denote those pairs (f, p) in $T_{\mathcal{A}}$ (resp. $T_{\mathcal{A}}^*$) where $|c_i| \leq H$ for all $i \in [m]$ and $p \leq H$. Finally, let $d := a_m / \gcd(a_2, \dots, a_m)$. Then for all $H \geq 17$ we have*


$$\frac{\#T_{\mathcal{A}}^*(H)}{\#T_{\mathcal{A}}(H)} \geq \left(1 - \frac{(2d-1)m}{2H+1}\right) \left(1 - \frac{1+(2d-1)\log(mH)\log H}{H}\right).$$

In particular, we will see in the proof of Assertion (2) of Theorem 1.4 that the exceptional set \mathcal{E} is merely the complement of the union $\bigcup_{\mathcal{A}} T_{\mathcal{A}}^*$ as \mathcal{A} ranges over all finite subsets of $\mathbb{N} \cup \{0\}$. Our corollary above is proved in Section 8.2.

Another bit of background we'll need to prove Assertion (2) of Theorem 1.4 is some arithmetic tropicalia.

Definition 4.7 *Given any polynomial $f(x_1) := \sum_{i=1}^m c_i x_1^{a_i} \in \mathbb{Z}[x_1]$, we define its p -adic Newton polygon, $\text{Newt}_p(f)$, to be the convex hull of the points $\{(a_i, \text{ord}_p c_i) \mid i \in \{1, \dots, m\}\}$. Also, a face of a polygon $Q \subset \mathbb{R}^2$ is called lower if and only if it has an inner normal with positive last coordinate, and the lower hull of Q is simply the union of all its lower edges. Finally, the polynomial associated to summing the terms of f corresponding to points of the form $(a_i, \text{ord}_p c_i)$ lying on a lower face of $\text{Newt}_p(f)$ is called a (p -adic) lower polynomial. \diamond*

Example 4.8 For $f(x_1) := 36 - 8868x_1 + 29305x_1^2 - 35310x_1^3 + 18240x_1^4 - 3646x_1^5 + 243x_1^6$, the polygon $\text{Newt}_3(f)$ has exactly 3 lower edges and can easily be verified to resemble the illustration to the right. The polynomial f thus has exactly 2 lower binomials, and 1 lower trinomial. \diamond



Note that the standard Newton polygon can be identified with the variant of the preceding construction that instead employs the *trivial* valuation (which sends all nonzero field elements to 0 and 0 to $+\infty$).

A remarkable fact true over \mathbb{C}_p but false over \mathbb{C} is that the norms of roots can be determined completely combinatorially.

Lemma 4.9 (See, e.g., [Rob00, Ch. 6, sec. 1.6].) *The number of roots of f in \mathbb{C}_p with valuation v , counting multiplicities, is exactly the horizontal length of the lower face of $\text{Newt}_p(f)$ with inner normal $(v, 1)$. ■*

Example 4.10 In Example 4.8 earlier, note that the 3 lower edges have respective horizontal lengths 2, 3, and 1, and inner normals $(1, 1)$, $(0, 1)$, and $(-5, 1)$. Lemma 4.9 then tells us that f has exactly 6 roots in \mathbb{C}_3 : 2 with 3-adic valuation 1, 3 with 3-adic valuation 0, and 1 with 3-adic valuation -5 . Indeed, one can check that the roots of f are exactly 6, 1, and $\frac{1}{243}$, with respective multiplicities 2, 3, and 1. \diamond

4.1 The Proof of Assertion (2) of Theorem 1.4

Let $f \in \mathbb{Z}[x_1]$, $\mathcal{A} := \text{Supp}(f)$, and assume $\mathcal{A} = \{a_1, \dots, a_m\}$. Since the roots of f in \mathbb{Q}_p^* are unaffected by multiplying f by a monomial, and since the existence of 0 as a root of f is clearly checkable in constant time, we may assume that $0 = a_1 < a_2 < \dots < a_m$. Via the reciprocal polynomial $f^*(x_1) := x_1^{\deg f} f(1/x_1)$, it is then enough to show that, for most f , having a root in $\mathbb{Z}_p \setminus \{0\}$ admits a succinct certificate. (Indeed, f has a root in \mathbb{Q}_p^* if and only if $[f$ has a root in $\mathbb{Z}_p \setminus \{0\}$ or f^* has a root in $\mathbb{Z}_p \setminus \{0\}]$.) Multiplying by another monomial if necessary, we can of course still continue to assume that $0 = a_1 < a_2 < \dots < a_m$. Letting $g := \gcd(a_2, \dots, a_m)$, we will also assume temporarily that $g = 1$ and handle the case $g > 1$ at the end of our proof.

Since convex hulls in the plane can be computed in quasi-linear time [OSvK00], it is clear by Proposition 2.1 that $\text{Newt}_p(f)$ can be computed in polynomial-time. Let c_i denote the coefficient of $x_1^{a_i}$ in f . Since $\text{ord}_p c_i \leq \log_p c_i \leq \text{size}(c_i)$, note also that every root $\zeta \in \mathbb{C}_p$ of f satisfies $|\text{ord}_p \zeta| \leq 2 \max_i \text{size}(c_i) \leq 2 \text{size}(f)$.

Since $\text{ord}_p(\mathbb{Z}_p) = \mathbb{N} \cup \{0\}$, we can clearly assume that $\text{Newt}_p(f)$ has an edge with non-positive integral slope, for otherwise f would have no roots in $\mathbb{Z}_p \setminus \{0\}$. Letting $\phi(x_1) := f'(x_1)/x_1^{a_2-1}$, and letting ζ be any nonzero p -adic integer root of f , note that $\text{ord}_p f'(\zeta) = (a_2 - 1)\text{ord}_p(\zeta) + \text{ord}_p \phi(\zeta)$. Note also that $\bar{\Delta}_{\mathcal{A}}(f) = \text{Res}_{(a_m, a_m - a_2)}(f, \phi)$ so if $p \nmid \bar{\Delta}_{\mathcal{A}}(f)$ then f and ϕ have no common roots in the algebraic closure of \mathbb{F}_p , by Lemma 4.3. In particular, $p \nmid \bar{\Delta}_{\mathcal{A}}(f) \implies \phi(\zeta) \not\equiv 0 \pmod{p}$; and thus $p \nmid \bar{\Delta}_{\mathcal{A}}(f, \phi) \implies \text{ord}_p f'(\zeta) = (a_2 - 1)\text{ord}_p(\zeta)$. Furthermore, by the convexity of the lower hull of $\text{Newt}_p(f)$, it is clear that $\text{ord}_p(\zeta) \leq \frac{\text{ord}_p c_0 - \text{ord}_p c_i}{a_i}$ where $(a_i, \text{ord}_p c_i)$ is the rightmost vertex of the lower edge of $\text{Newt}_p(f)$ with

least (non-positive and integral) slope. Clearly then, $\text{ord}_p(\zeta) \leq \frac{2 \max_i \log_p |c_i|}{a_1}$. So $p \nmid \bar{\Delta}_{\mathcal{A}}(f) \implies \text{ord}_p f'(\zeta) \leq 2 \text{size}(f)$.

Our fraction of inputs admitting a succinct certificate will then correspond precisely to those (f, p) such that $p \nmid \bar{\Delta}_{\mathcal{A}}(f)$. In particular, let us define \mathcal{E} to be the union of all pairs (f, p) such that $p \mid \bar{\Delta}_{\mathcal{A}}(f)$, as \mathcal{A} ranges over all finite subsets of $\mathbb{N} \cup \{0\}$. It is then easily checked that \mathcal{E} is a countable union of hypersurfaces, and the density 0 statement follows immediately from Corollary 4.6.

Now fix $\ell = 4 \text{size}(f) + 1$. Clearly then, by Hensel's Lemma, for any $(f, p) \in (\mathbb{Z}[x_1] \times \mathcal{P}) \setminus \mathcal{E}$, f has a root $\zeta \in \mathbb{Z}_p$ if and only if f has a root $\zeta_0 \in \mathbb{Z}/p^\ell \mathbb{Z}$. Since $\log(p^\ell) = O(\text{size}(f) \log p) = O((\text{size}(f) + \log p)^2)$, and since arithmetic in $\mathbb{Z}/p^\ell \mathbb{Z}$ can be done in time polynomial in $\log(p^\ell)$ [BS96, Ch. 5], we have thus at last found our desired certificate: a root $\zeta_0 \in (\mathbb{Z}/p^\ell \mathbb{Z})^*$ of f with $\ell = 4 \text{size}(f) + 1$.

To conclude, let us address the case $g > 1$: by our preceding construction, a certificate that clearly works for this case will simply be a root $\zeta_0 \in (\mathbb{Z}/p^\ell \mathbb{Z})^*$ of f (with $\ell = 4 \text{size}(f) + 1$) *also* satisfying the condition that $x^g - \zeta_0$ has a root in $(\mathbb{Z}/p^\ell \mathbb{Z})^*$ (thanks to the binomial case of Assertion (0) of Theorem 1.4). ■

5 Degenerate Trinomials, Linear Forms in p-adic Logarithms, and Assertion (1) of Theorem 1.4

We will first need to recall the concept of a *gcd-free basis*. In essence, a gcd-free basis is nearly as powerful as factorization into primes, but far easier to compute.

Definition 5.1 [BS96, Sec. 8.4] *For any subset $\{\alpha_1, \dots, \alpha_N\} \subset \mathbb{N}$, a gcd-free basis for $\{\alpha_1, \dots, \alpha_N\}$ is a pair of sets $(\{\gamma_i\}_{i=1}^\eta, \{e_{ij}\}_{(i,j) \in [N] \times [\eta]})$ such that (1) $\text{gcd}(\gamma_i, \gamma_j) = 1$ for all $i \neq j$, and (2) $\alpha_i = \prod_{j=1}^\eta \gamma_j^{e_{ij}}$ for all i . ◊*

Theorem 5.2 *Following the notation of Definition 5.1, we can compute a gcd-free basis for $\{\alpha_1, \dots, \alpha_N\}$ (with η linear in $N + \max_i \log \alpha_i$) in time linear in $N + \max_i \log^2 \alpha_i$. In particular, if $u_1, \dots, u_N \in \mathbb{Z}$ then we can decide $\alpha_1^{u_1} \cdots \alpha_N^{u_N} \stackrel{?}{=} 1$ in time linear in $N + (\max_i \log(\alpha_i) + \max_i \log(u_i))^2$. ■*

The first assertion of Theorem 5.2 follows immediately from [BS96, Thm. 4.8.7, Sec. 4.8] and the naive bounds for the complexity of integer multiplication. The second assertion then follows immediately by checking whether the linear combinations $\sum_{i=1}^N e_{ij} u_i$ are all 0 or not.

We now make some final observations about the roots of trinomials before proving Assertion (1) of Theorem 1.4. Recall that a degenerate root of f is a $\zeta \in \mathbb{C}_p$ with $f(\zeta) = f'(\zeta) = 0$.

Lemma 5.3 *Suppose $f(x_1) = c_1 + c_2 x_1^{a_2} + c_3 x_1^{a_3} \in \mathcal{F}_{1,3}$, $\mathcal{A} := \{0, a_2, a_3\}$, $0 < a_2 < a_3$, and $\text{gcd}(a_2, a_3) = 1$. Then:*

$$(0) \bar{\Delta}_{\mathcal{A}}(f) = a_3^{a_3} c_1^{a_3 - a_2} c_3^{a_2} - a_2^{a_2} (a_3 - a_2)^{a_3 - a_2} (-c_2)^{a_3}.$$

(1) $\bar{\Delta}_{\mathcal{A}}(f) \neq 0 \iff f$ has no degenerate roots in \mathbb{C}_p . In which case, we also have

$$\bar{\Delta}_{\mathcal{A}}(f) = \left(\frac{c_3}{c_1}\right)^{a_2 - 1} \prod_{f(\zeta)=0} f'(\zeta) = (-1)^{a_3(a_3 - a_2)} \prod_{f(\zeta)=0} (a_2 c_2 + a_3 c_3 \zeta^{a_3 - a_2}).$$

- (2) Deciding whether f has a degenerate root in \mathbb{C}_p can be done in time polynomial in $\text{size}(f) + \log p$.
- (3) If f has a degenerate root $\zeta \in \mathbb{C}_p^*$ then $(\zeta^{a_2}, \zeta^{a_3}) = \frac{c_1}{a_3 - a_2} \left(-\frac{a_3}{c_2}, \frac{a_2}{c_3} \right)$. In particular, such a ζ is unique and lies in \mathbb{Q} .
- (4) The polynomial $q(x_1) := (a_3 - a_2) - a_3x_1^{a_2} + a_2x_1^{a_3}$ has 1 as its unique degenerate root and satisfies $\lim_{x_1 \rightarrow 1} \frac{q(x_1)}{(x_1 - 1)^2} = a_2a_3(a_3 - a_2)/2$ and $\bar{\Delta}_{\{0, \dots, a_3 - 2\}} \left(\frac{q(x_1)}{(x_1 - 1)^2} \right) = a_3(a_2a_3(a_3 - a_2))^{a_3 - 4} J$, where $J = O(a_2^2 a_3^3 (a_3 - a_2)^2)$ is a nonzero integer.

Proof of Lemma 5.3:

Part (0): One simply mimicks the argument from [GKZ94, pp. 406–407]. ■

Part (1): The first assertion follows directly from Definition 4.4 and the vanishing criterion for $\text{Res}_{(a_3, a_3 - a_2)}$ from Lemma 4.3. To prove the second assertion, observe that the product formula from Lemma 4.3 implies that $\bar{\Delta}_{\mathcal{A}}(f) = (-1)^{a_3(a_3 - a_2)} c_3^{a_3 - a_2} \prod_{f(\zeta)=0} (a_2 c_2 + a_3 c_3 \zeta^{a_3 - a_2})$.

Combining with the basic identity $\prod_{f(\zeta)=0} \zeta = (-1)^{a_3} \frac{c_1}{c_3}$, and the fact that $a_3^2 = a_3 \pmod{2}$ and $-1 = 1 \pmod{2}$, we are done. ■

Part (2): From Part (1) it suffices to detect the vanishing of $\bar{\Delta}_{\mathcal{A}}(f)$. However, while Part (0) implies that one can evaluate $\bar{\Delta}_{\mathcal{A}}(f)$ with a small number of arithmetic operations, the bit-size of $\bar{\Delta}_{\mathcal{A}}(f)$ can be quite large. Nevertheless, we can decide within time polynomial in $\text{size}(f)$ whether these particular $\bar{\Delta}_{\mathcal{A}}(f)$ vanish for integer c_i via gcd-free bases (invoking Theorem 5.2). ■

Part (3): It is easily checked that if $\zeta \in \mathbb{C}_p$ is a degenerate root of f then the vector $[c_1, c_2 \zeta^{a_2}, c_3 \zeta^{a_3}]$ must be a right null vector for the matrix $M := \begin{bmatrix} 1 & 1 & 1 \\ 0 & a_2 & a_3 \end{bmatrix}$. Since $[a_3 - a_2, -a_3, a_2]$ is clearly a right null vector for M , $[c_1, c_2 \zeta^{a_2}, c_3 \zeta^{a_3}]$ must then be a multiple of $[a_3 - a_2, -a_3, a_2]$. Via the extended Euclidean algorithm [BS96, Sec. 4.3], we can then find A and B (also of size polynomial in $\text{size}(f)$) with $Aa_2 + Ba_3 = 1$. So then we obtain that $\left(\frac{c_2 \zeta^{a_2}}{c_1} \right)^A \left(\frac{c_3 \zeta^{a_3}}{c_1} \right)^B = \frac{c_2^A c_3^B}{c_1^{A+B}} \zeta = \left(\frac{-a_3}{a_3 - a_2} \right)^A \left(\frac{a_2}{a_3 - a_2} \right)^B$. ■

Part (4): That 1 is a root of q is obvious. Uniqueness follows directly from Part (3) and our assumption that $\text{gcd}(a_2, a_3) = 1$. The limit formula follows easily from two application of L'Hôpital's Rule.

To prove the final assertion, first note that a routine long division reveals that $\frac{q(x_1)}{(x_1 - 1)^2}$ has coefficients rising by one arithmetic progression and then falling by another. Explicitly,

$$\frac{q(x_1)}{(x_1 - 1)^2} = \left(\sum_{i=1}^{a_2 - 1} (a_3 - a_2) i x_1^{i-1} \right) + \left(\sum_{i=1}^{a_3 - a_2} (a_3 - a_2 + 1 - i) a_2 x_1^{a_2 - 2 + i} \right).$$

Definition 4.2 then implies that $\bar{\Delta}_{\{0, \dots, a_3 - 2\}} \left(\frac{q(x_1)}{(x_1 - 1)^2} \right)$ is exactly $\frac{1}{a_2}$ times the determinant of the following quasi-Toeplitz matrix which we will call \mathcal{M} :

$$\begin{bmatrix} a_3 - a_2 & 2(a_3 - a_2) & \cdots & (a_2 - 1)(a_3 - a_2) & (a_3 - a_2)a_2 & \cdots & 2a_2 & a_2 & 0 & \cdots & 0 \\ & \ddots & \ddots & & \ddots & \ddots & & \ddots & \ddots & & \\ 1 \cdot 2 \cdot (a_3 - a_2) & 2 \cdot 3 \cdot (a_3 - a_2) & \cdots & (a_2 - 2)(a_2 - 1)(a_3 - a_2) & (a_2 - 1)(a_3 - a_2)a_2 & \cdots & (a_3 - 2) \cdot 1 \cdot a_2 & 0 & \cdots & & 0 \\ & \ddots & \ddots & & \ddots & \ddots & & \ddots & \ddots & & \end{bmatrix},$$

tions of a column from another column and subtractions of a row from another row, we can then reduce our matrix to a $(2a_3 - 5) \times (2a_3 - 5)$ permutation matrix with the a_3^{rd} row and $(2a_3 - 5)^{\text{th}}$ row resembling the corresponding rows above. In particular, these 2 new rows have entries at worst $O(a_3)$ times larger than before. Clearly then, our final determinant is a nonzero integer with absolute value $O(a_2^2(a_3 - a_2)^2 a_3^3)$. ■

We now quote the following important result on lower binomials.

Theorem 5.4 (See [AI11, Thm. 3.10 & Prop. 4.4].) *Suppose $(f, p) \in \mathbb{Z}[x_1] \times \mathcal{P}$, $(v, 1)$ is an inner normal to a lower edge E of $\text{Newt}_p(f)$, the lower polynomial g corresponding to E is a binomial with exponents $\{a_i, a_j\}$, and p does not divide $a_i - a_j$. Then the number of roots $\zeta \in \mathbb{Q}_p$ of f with $\text{ord}_p \zeta = v$ is exactly the number of roots of g in \mathbb{Q}_p^* . ■*

Finally, we recall a deep theorem from Diophantine approximation that allows us to carefully bound from above the p -adic valuation of certain high degree binomials.

Yu's Theorem. [Yu94, pg. 242] *Suppose $p \in \mathbb{N}$ is any prime; $\frac{\alpha_1^+}{\alpha_1}, \dots, \frac{\alpha_m^+}{\alpha_m} \in \mathbb{Q} \setminus \{0\}$ are fractions in lowest terms; and β_1, \dots, β_m are integers not all zero. Then $\left(\frac{\alpha_1^+}{\alpha_1}\right)^{\beta_1} \cdots \left(\frac{\alpha_m^+}{\alpha_m}\right)^{\beta_m} \neq 1$ implies that $\text{ord}_p \left(\left(\frac{\alpha_1^+}{\alpha_1}\right)^{\beta_1} \cdots \left(\frac{\alpha_m^+}{\alpha_m}\right)^{\beta_m} - 1 \right)$ is strictly less than*

$$22000 \left(\frac{9.5(m+1)}{\sqrt{\log p}} \right)^{2(m+1)} (p-1) \log(10mh) \log \max\{3, \max_i |\beta_i|\} \prod_{i=1}^m \max\{\log |\alpha_i^\pm|, \log p\},$$

where $h = \max\{\max_i \log |\alpha_i^\pm|, \log p, 1\}$ and the imaginary part of \log lies in $(-\pi, \pi)$. ■

Let us call any $\text{Newt}_p(f)$ such that f has no lower m -nomials with $m \geq 3$ *generic*. Oppositely, we call $\text{Newt}_p(f)$ *flat* if it is a line segment. Finally, if $p|(a_i - a_j)$ with $\{a_i, a_j\}$ the exponents of some lower binomial of f then we call $\text{Newt}_p(f)$ *ramified*. We will see later that certain ramified cases are where one begins to see the surprising complexity behind proving $\text{FEAS}_{\mathbb{Q}_p}(\mathcal{F}_{1,3}) \in \mathbf{NP}$, including the need for Yu's Theorem above.

5.1 The Proof of Assertion (1) of Theorem 1.4

Our underlying certificate will ultimately be a root $\zeta_0 \in \mathbb{Z}/p^\ell \mathbb{Z}$ for f (or a slight variant thereof) with $\ell = O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))$. Certain cases will force us to use tools that can currently only yield complexity upper bounds of the preceding magnitude.

Let us write $f(x_1) = c_1 + c_2 x_1^{a_2} + c_3 x_1^{a_3}$. Just as in Section 4.1, we may assume $c_1 \neq 0$ and reduce to certifying roots in \mathbb{Z}_p . We may also assume that the rightmost (or only) lower edge of f is a horizontal line segment at height 0. (And thus $\text{ord}_p c_1 \geq 0$ in particular.) This is because, as already observed earlier in the proof of Assertion (2) of Theorem 1.4, we can compute $\text{Newt}_p(f)$ in time polynomial in $\text{size}(f) + \log p$. So we can rescale f (in polynomial-time) without increasing $\text{size}(f)$. More precisely, if $\text{Newt}_p(f)$ has no lower edges of integral slope then we can immediately conclude that f has no roots in \mathbb{Q}_p by Lemma 4.9. So, replacing f by the reciprocal polynomial f^* if necessary, we may assume that the rightmost lower edge of f has integral slope and then set $g(x_1) := p^{-\text{ord}_p c_2} f\left(p^{\frac{\text{ord}_p(c_2) - \text{ord}_p(c_3)}{a_3 - a_2}} x_1\right)$. The lower hull of $\text{Newt}_p(g)$ then clearly has the desired shape, and it is clear that f has a root in \mathbb{Q}_p if and only if g has a root in \mathbb{Q}_p . In particular, it is easily checked that $\text{size}(g) \leq \text{size}(f)$.

To simplify our proof we will assume that $\gcd(a_2, a_3) = 1$ (unless otherwise noted), and recover the case $\gcd(a_2, a_3) > 1$ at the very end of our proof. The vanishing of $\bar{\Delta}_{\mathcal{A}}(f)$, which can be detected in \mathbf{P} thanks to Lemma 5.3, then determines 2 cases:

Case (a): $\bar{\Delta}_{\mathcal{A}}(f) \neq 0$

Since $\gcd(a_2, a_3) = 1$ we may clearly assume that p divides at most one of $\{a_2, a_3, a_3 - a_2\}$. The shape of the lower hull of $\text{Newt}_p(f)$ (which we've already observed can be computed in time polynomial in $\text{size}(f) + \log p$) then determines 2 subcases:

↙ If $\text{Newt}_p(f)$ has lower hull a line segment then we may also assume (by rescaling f as detailed above) that $p \nmid c_1, c_3$ and $e := \text{ord}_p c_2 \geq 0$.

When p divides either a_2 or $a_3 - a_2$ then we can easily find certificates for solvability of f over \mathbb{Q}_p : If $e = 0$ then $p \nmid \bar{\Delta}_{\mathcal{A}}(f)$ by Lemma 5.3 (since $p \nmid a_3$) and thus f has no degenerate roots mod p . So Hensel's Lemma implies that we can use a root of f in $\mathbb{Z}/p\mathbb{Z}$ as a certificate for f having a root in \mathbb{Q}_p . If $e > 0$ then we can in fact detect roots in \mathbb{Q}_p for f in \mathbf{P} by the binomial case, thanks to Theorem 5.4.

So let us now assume p does not divide a_2 or $a_3 - a_2$, and set $e' := \text{ord}_p a_3$. If $e > e'$ then observe that $f'(x) = a_3 c_3 x^{a_3 - 1} \pmod{p^e}$. By Lemma 4.9, any putative root $\zeta \in \mathbb{Q}_p$ of f must satisfy $\text{ord}_p \zeta = 0$. So $f'(\zeta) \neq 0 \pmod{p^e}$ and Hensel's Lemma implies that a root of f in $\mathbb{Z}/p^{2e+1}\mathbb{Z}$ is clearly a certificate for f having a root in \mathbb{Q}_p . Our certificate can also clearly be verified in time polynomial in $\text{size}(f) + \log p$ since $\text{size}(p^{2e+1}) \leq (2e + 1) \log p \leq \text{size}(f) \log p$.

If $e < e'$ then $f'(x) = a_2 c_2 x^{a_2 - 1} \pmod{p^{e'}}$. Similar to the last paragraph, $f'(\zeta) \neq 0 \pmod{p^{e'}}$ and we then instead employ a root of f in $\mathbb{Z}/p^\ell\mathbb{Z}$ with $\ell = 2e' + 1$ as a certificate for f having a root in \mathbb{Q}_p .

Now, if $e = e'$, observe that $\text{ord}_p f'(\zeta) = \text{ord}_p \frac{f'(\zeta)}{\zeta^{a_2 - 1}}$ since Lemma 4.9 tells us that $\text{ord}_p \zeta = 0$ for any root $\zeta \in \mathbb{C}_p$. Since $\bar{\Delta}_{\mathcal{A}}(f) \neq 0$, Lemma 5.3 then tells us that $\text{ord}_p(a_2 c_2 + a_3 c_3 \zeta^{a_3 - a_2}) < +\infty$. So $\text{ord}_p f'(\zeta) < +\infty$ for any root $\zeta \in \mathbb{C}_p$ of f and then Lemma 5.3 tells us that

$$\begin{aligned} \text{ord}_p \prod_{f(\zeta)=0} \frac{f'(\zeta)}{\zeta^{a_2 - 1}} &= \sum_{f(\zeta)=0} \text{ord}_p f'(\zeta) = \text{ord}_p \left((a_3 - a_2)^{a_3 - a_2} a_2^{a_2} c_2^{a_3} - (-a_3)^{a_3} c_1^{a_3 - a_2} c_3^{a_2} \right) \\ &= a_3 e + \text{ord}_p \left(\left(\frac{a_3 - a_2}{c_1} \right)^{a_3 - a_2} \left(\frac{a_2}{c_3} \right)^{a_2} \left(\frac{c_2}{-a_3} \right)^{a_3} - 1 \right). \end{aligned}$$

(since $p \nmid (c_1 c_3)$).

So by the $m = 3$ case of Yu's Theorem we obtain

$$\sum_{f(\zeta)=0} \text{ord}_p f'(\zeta) = a_3 e + O(p(\text{size}(f)/\log p)^4 \log \text{size}(f)).$$

Now, since $p^e \mid c_2, a_3$, we have $\text{ord}_p f'(\zeta) \geq e$ for any root $\zeta \in \mathbb{C}_p$ of f . So all roots $\zeta \in \mathbb{C}_p$ of f must satisfy $\text{ord}_p f'(\zeta) = e + O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))$
 $= O(\text{size}(f)) + O(p(\text{size}(f)/\log p)^4 \log \text{size}(f)) = O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))$.

In other words, a root of f in $\mathbb{Z}/p^{O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))}\mathbb{Z}$ suffices as a certificate, thanks to Hensel's Lemma.

↘ If the lower hull of $\text{Newt}_p(f)$ is not a line segment then (by rescaling f as detailed above), we may also assume that $p \mid c_1$ but $p \nmid c_2, c_3$. Since $\gcd(a_2, a_3) = 1$, we may also assume (via rescaling and/or reciprocals) that $p \nmid a_2 a_3$, i.e., if p divides the length of any lower edge of $\text{Newt}_p(f)$ then it is the rightmost (now horizontal) edge.

Via Theorem 5.4 and the binomial case of Assertion (0) we can easily decide (within time polynomial in $\text{size}(f) + \log p$) the existence of a root of f in \mathbb{Z}_p with valuation v , where $(v, 1)$ is an inner normal of the left lower edge of $\text{Newt}_p(f)$. So now we need only efficiently detect

roots in \mathbb{Z}_p of valuation 0. Toward this end, let us now set $e := \text{ord}_p c_1$ and $e' := \text{ord}_p(a_3 - a_2)$. Clearly, $e > 0$ or else we would be in the earlier case where $\text{Newt}_p(f)$ has lower hull a single edge.

If $e > e'$ then $f(x) = c_2 x^{a_2} + c_3 x^{a_3} \pmod{p^e}$ and thus $f'(\zeta) = a_2 c_2 \zeta^{a_2-1} + a_3 c_3 \zeta^{a_3-1} = -a_2 c_3 \zeta^{a_3-1} + a_3 c_3 \zeta^{a_3-1} = c_3(a_3 - a_2) \zeta^{a_3-1} \pmod{p^e}$ for any root $\zeta \in \mathbb{C}_p$ of f . So $f'(\zeta) \neq 0 \pmod{p^e}$ for any root $\zeta \in \mathbb{Z}_p$ of valuation 0 and thus, by Hensel's Lemma, we can certify the existence of such a ζ in \mathbf{NP} by a root of f in $\mathbb{Z}/p^{2e+1}\mathbb{Z}$.

If $e < e'$ then $f'(x) = a_2 c_2 x^{a_2-1} + a_3 c_3 x^{a_3-1} = a_3 c_2 x^{a_2-1} + a_3 c_3 x^{a_3-1} \pmod{p^{e'}}$ since $a_3 = a_2 \pmod{p^{e'}}$. So $f'(\zeta) = a_3 c_2 \zeta^{a_2-1} - a_3(c_1 \zeta^{-1} + c_2 \zeta^{a_2-1}) = -\frac{a_3 c_1}{\zeta} \neq 0 \pmod{p^{e'}}$ for any root $\zeta \in \mathbb{C}_p$ of f . So a root of f in $\mathbb{Z}/p^{2e'+1}\mathbb{Z}$ serves as a certificate for a root of f in \mathbb{Z}_p .

Finally, if $e = e'$, observe that $f'(x) = a_2 c_2 x^{a_2-1} + a_3 c_3 x^{a_3-1}$ and there are exactly a_2 (resp. $a_3 - a_2$) roots of f in \mathbb{C}_p of valuation $\frac{e}{a_2}$ (resp. 0) by Lemma 4.9. Using the fact that $p \nmid a_2 a_3 c_2 c_3$, it is then easy to see that $\text{ord}_p f'(\zeta) = \left(\frac{a_2-1}{a_2}\right) e$ for any root $\zeta \in \mathbb{C}_p$ of f with valuation $\frac{e}{a_2}$.

The value of $\text{ord}_p f'(\zeta)$ is harder to control when $\zeta \in \mathbb{C}_p$ is root of valuation 0. So let us observe that, at such a ζ , $f'(\zeta) = \frac{a_3 c_1}{\zeta} + f'(\zeta) \pmod{p^e}$ and thus:

(\star)
$$f'(\zeta) = \frac{a_3 c_1}{\zeta} + a_2 c_2 \zeta^{a_2-1} + a_3 c_3 \zeta^{a_3-1} = \frac{a_3 c_1}{\zeta} + a_3 c_2 \zeta^{a_2-1} + a_3 c_3 \zeta^{a_3-1} = \frac{a_3}{\zeta} f(\zeta) = 0 \pmod{p^e},$$
 since $e = e'$ and $a_2 = a_3 \pmod{p^{e'}}$. So $e \leq \text{ord}_p f'(\zeta)$ at any such ζ . Similar to our earlier flat case, Part (1) of Lemma 5.3 then implies the following:

$$\text{ord}_p \bar{\Delta}_{\mathcal{A}}(f) = -(a_2 - 1)e + \sum_{f(\zeta)=0} f'(\zeta) = \sum_{\substack{f(\zeta)=0 \\ \text{ord}_\zeta=0}} f'(\zeta).$$

On the other hand, since $e = \text{ord}_p(a_3 - a_2) = \text{ord}_p c_1$, Part (0) of Lemma 5.3 combined with the $m=3$ case of Yu's Theorem implies that $\text{ord}_p \bar{\Delta}_{\mathcal{A}}(f) = (a_3 - a_2)e + O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))$. So any root $\zeta \in \mathbb{C}_p$ of f having valuation 0 must satisfy

$$\text{ord}_p f'(\zeta) \leq e + O(p(\text{size}(f)/\log p)^4 \log \text{size}(f)) \leq \text{size}(f) + O(p(\text{size}(f)/\log p)^4 \log \text{size}(f)).$$

So again, a root of f in $\mathbb{Z}/p^{O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))}\mathbb{Z}$ suffices as a certificate, thanks to Hensel's Lemma.

Remark 5.5 *Note that if $\text{Newt}_p(f)$ is unramified as well as generic, then Theorem 5.4 implies that we can in fact decide the existence of roots in \mathbb{Q}_p for f in \mathbf{P} . \diamond*

Case (b): $\bar{\Delta}_{\mathcal{A}}(f) = 0$

First note that, independent of $\gcd(a_2, a_3)$, a degenerate root of f in \mathbb{Q}_p admits a very simple certificate: a $\zeta \in \mathbb{Z}/p^{4\text{size}(f)+1}\mathbb{Z}$ satisfying $c_2(a_3 - a_2)\zeta^{a_2} + c_1 a_3 = c_3(a_3 - a_2)\zeta^{a_3} - c_1 a_2 = 0 \pmod{p^{4\text{size}(f)+1}}$. Thanks to Lemma 5.3 and our proof of Assertion (0) in Section 3, it is clear that the preceding 2×1 binomial system has a solution if and only if f has a degenerate root in \mathbb{Q}_p .

So now we resume our assumption that $\gcd(a_2, a_3) = 1$ and build certificates for the *non-degenerate* roots of f in \mathbb{Z}_p . Toward this end, observe that the proof of Lemma 5.3 tells us that the unique degenerate root ζ of f lies in \mathbb{Q}^* and satisfies $[c_1, c_2 \zeta^{a_2}, c_3 \zeta^{a_3}] = \gamma[a_3 - a_2, -a_3, a_2]$ for some $\gamma \in \mathbb{Q}$. Clearly then, $q(x_1) = \frac{1}{\gamma} f(\zeta x_1)$, and f has exactly the same number of roots in \mathbb{Q}_p as q does.

So we can henceforth restrict to the special case $f(x_1) = (a_3 - a_2) - a_3 x_1^{a_2} + a_2 x_1^{a_3}$, and let $r(x_1) := \frac{f(x_1)}{(x_1-1)^2}$ and $\bar{\Delta} := \bar{\Delta}_{\{0, \dots, a_3-2\}}(r)$. Should $p \nmid a_2 a_3 (a_3 - a_2)$ then f is clearly flat

and thus all the roots of f have valuation 0. Part (4) of Lemma 5.3 then tells us that $\text{ord}_p \bar{\Delta} = O(\log(a_2) + \log(a_3)) = O(\text{size}(f))$ and thus the product formula from Lemma 4.3 implies that $\text{ord}_p r'(\zeta) = O(\text{size}(f))$ at any root $\zeta \in \mathbb{C}_p$ of r . So a root $\zeta_0 \in \mathbb{Z}/p^{O(\text{size}(f))}\mathbb{Z}$ of r suffices as a certificate for f to have a root in \mathbb{Q}_p other than 1. (Note also that by construction, r can clearly be evaluated mod $p^{O(\text{size}(f))}$ within a number of arithmetic operations quadratic in $\text{size}(f) + \log p$.)

So let us now assume that p divides exactly one number from $\{a_2, a_3, a_3 - a_2\}$. (Otherwise, p would divide all 3 numbers, thus contradicting the assumption $\text{gcd}(a_2, a_3) = 1$.)

↙ If $p|a_3$ then f is clearly flat and, by Lemma 4.9, every root of r has valuation 0. This implies $\text{ord}_p r'(\zeta) \geq 0$ at any root $\zeta \in \mathbb{C}_p$ of r . So by the product formula from Lemma 4.3 and Part (4) of Lemma 5.3, combined with the fact that $\text{ord}_p t \leq \log_p t$ for any integer t , we obtain that

$$\text{ord}_p \bar{\Delta} = (a_3 - 3)\text{ord}_p(a_2) + \sum_{r(\zeta)=0} \text{ord}_p r'(\zeta) = (a_3 - 4)\text{ord}_p(a_2) + O(\log(a_2) + \log(a_3)).$$

So $\text{ord}_p r'(\zeta) = O(\log(a_2) + \log(a_3)) = O(\text{size}(f))$ at any root $\zeta \in \mathbb{C}_p$ and we can again use a root $\zeta_0 \in \mathbb{Z}/p^{O(\text{size}(f))}\mathbb{Z}$ of r as a certificate for f to have a root in \mathbb{Q}_p other than 1.

↘ Replacing f by the reciprocal polynomial f^* if need be, we are left with the case $p|(a_3 - a_2)$. By Lemma 4.9, f clearly has exactly a_2 (resp. $a_3 - a_2$) roots of valuation $\frac{\text{ord}_p(a_3 - a_2)}{a_2}$ (resp. 0) in \mathbb{C}_p .

Since $p \nmid a_2$, Theorem 5.4 tells us that we can apply the binomial case of Assertion (0) of Theorem 1.4 to detect roots of f in \mathbb{Q}_p with valuation $\frac{\text{ord}_p(a_3 - a_2)}{a_2}$ in polynomial-time. So let us now focus on roots $\zeta \in \mathbb{C}_p \setminus \{1\}$ of f having valuation 0.

For any such root, we then obtain $\text{ord}_p f'(\zeta) \geq \text{ord}_p(a_3 - a_2)$, thanks to identity (\star) from the non-degenerate case. Note also that $r'(\zeta) = \frac{f'(\zeta)}{(\zeta-1)^2} - 2\frac{f(\zeta)}{(\zeta-1)^3} = \frac{f'(\zeta)}{(\zeta-1)^2}$. Employing the product formula from Lemma 4.3 we then obtain

$$\text{ord}_p \bar{\Delta} = \left(\sum_{r(\zeta)=0} \text{ord}_p f'(\zeta) \right) - 2\text{ord}_p \prod_{r(\zeta)=0} (\zeta - 1) = \left(\sum_{r(\zeta)=0} \text{ord}_p f'(\zeta) \right) - 2\text{ord}_p r(1)$$

since $p \nmid a_2$. Part (4) of Lemma 5.3 tells us that $\text{ord}_p r(1)$ is $\text{ord}_p(a_3 - a_2)$ or $\text{ord}_p(a_3 - a_2) - 1$, according as $p \geq 3$ or $p = 2$. So applying Part (4) of Lemma 5.3 one last time we obtain

$$\sum_{r(\zeta)=0} \text{ord}_p f'(\zeta) = (a_3 - 4)\text{ord}_p(a_3 - a_2) + O(\log(a_2) + \log(a_3)) + 2\text{ord}_p(a_3 - a_2).$$

Now note that $f'(\zeta) = a_2 a_3 \zeta^{a_2-1} (\zeta^{a_3-a_2} - 1)$. Since f has exactly a_2 roots of valuation $\frac{\text{ord}_p(a_3 - a_2)}{a_2}$, we thus obtain

$$\begin{aligned} \sum_{\substack{r(\zeta)=0 \\ \text{ord}_p \zeta=0}} \text{ord}_p f'(\zeta) &= (a_3 - 2)\text{ord}_p(a_3 - a_2) + O(\log_p(a_2) + \log_p(a_3)) - (a_2 - 1)\text{ord}_p(a_3 - a_2) \\ &= (a_3 - a_2 - 1)\text{ord}_p(a_3 - a_2) + O(\log_p(a_2) + \log_p(a_3)). \end{aligned}$$

Since $\text{ord}_p f'(\zeta) \geq \text{ord}_p(a_3 - a_2)$ at a valuation 0 root $\zeta \in \mathbb{C}_p$ of f , and there are exactly $a_3 - a_2$ such roots, the value of $\text{ord}_p f'(\zeta)$ at such a root must therefore admit an upper bound of

$$\text{ord}_p f'(\zeta) = -\text{ord}_p(a_3 - a_2) + O(\log_p(a_2) + \log_p(a_3)) = O(\text{size}(f)).$$

So we can certify a non-degenerate root $\zeta \in \mathbb{Q}_p \setminus \{1\}$ of f with valuation 0 by a root $\zeta_0 \in \mathbb{Z}/p^{O(\text{size}(f))}\mathbb{Z}$ of r mod $p^{O(\text{size}(f))}$ not divisible by p .

Wrapping up the case $\gcd(a_2, a_3) > 1$: From our preceding arguments, we see that we are left with certifying the existence of *non*-degenerate roots in the case $g := \gcd(a_2, a_3) > 1$. Fortunately, this is simple: we merely find a non-degenerate root $\zeta_0 \in \mathbb{Z}/p^\ell\mathbb{Z}$ of $\bar{f} := c_1 + c_2x^{a_2/g} + c_3x^{a_3/g}$ as before (with ℓ depending on the case \bar{f} falls into), *also* satisfying the condition that $x^g - \zeta_0$ has a root in $\mathbb{Z}/p^\ell\mathbb{Z}$. Thanks to Corollary 3.2, we are done. ■

6 NP-hardness in One Variable: Proving Theorem 1.2

We will first need to develop two key ingredients: (A) Plaisted’s beautiful connection between Boolean satisfiability and roots of unity, and (B) an algorithm for constructing moderately small primes p with $p - 1$ having many prime factors.

6.1 Roots of Unity and NP-Completeness

Let us define $[n] := \{1, \dots, n\}$. Recall that any Boolean expression of one of the following forms:

$$(\diamond) y_i \vee y_j \vee y_k, \quad \neg y_i \vee y_j \vee y_k, \quad \neg y_i \vee \neg y_j \vee y_k, \quad \neg y_i \vee \neg y_j \vee \neg y_k, \quad \text{with } i, j, k \in [3n],$$

is a **3CNFSAT clause**. A *satisfying assignment* for an arbitrary Boolean formula $B(y_1, \dots, y_n)$ is an assignment of values from $\{0, 1\}$ to the variables y_1, \dots, y_n which makes the equality $B(y_1, \dots, y_n) = 1$ true. Let us now refine slightly Plaisted’s elegant reduction from **3CNFSAT** to feasibility testing for univariate polynomial systems over the complex numbers [Pla84, Sec. 3, pp. 127–129].

Definition 6.1 *Letting $\bar{p} := (p_1, \dots, p_n)$ denote any strictly increasing sequence of primes, let us inductively define a semigroup homomorphism $\rho_{\bar{p}}$ — the Plaisted morphism with respect to \bar{p} — from certain Boolean expressions in the variables y_1, \dots, y_n to $\mathbb{Z}[x]$, as follows:⁷ (0) $D_{\bar{p}} := \prod_{i=1}^n p_i$, (1) $\rho_{\bar{p}}(0) := 1$, (2) $\rho_{\bar{p}}(y_i) := x^{D_{\bar{p}}/p_i} - 1$, (3) $\rho_{\bar{p}}(\neg B) := (x^{D_{\bar{p}}} - 1) / \rho_{\bar{p}}(B)$, for any Boolean expression B for which $\rho_{\bar{p}}(B)$ has already been defined, (4) $\rho_{\bar{p}}(B_1 \vee B_2) := \text{lcm}(\rho_{\bar{p}}(B_1), \rho_{\bar{p}}(B_2))$, for any Boolean expressions B_1 and B_2 for which $\rho_{\bar{p}}(B_1)$ and $\rho_{\bar{p}}(B_2)$ have already been defined. \diamond*

Lemma 6.2 [Pla84, Sec. 3, pp. 127–129] *Suppose $\bar{p} = (p_i)_{i=1}^n$ is an increasing sequence of primes with $\log(p_k) = O(k^\gamma)$ for some constant γ . Then, for all $n \in \mathbb{N}$ and any clause C of the form (\diamond) , we have $\text{size}(\rho_{\bar{p}}(C))$ polynomial in n^γ . In particular, $\rho_{\bar{p}}$ can be evaluated at any such C in time polynomial in n . Furthermore, if K is any field possessing $D_{\bar{p}}$ distinct $D_{\bar{p}}^{\text{th}}$ roots of unity, then a **3CNFSAT** instance $B(y) := C_1(y) \wedge \dots \wedge C_k(y)$ has a satisfying assignment if and only if the univariate polynomial system $F_B := (\rho_{\bar{p}}(C_1), \dots, \rho_{\bar{p}}(C_k))$ has a root $\zeta \in K$ satisfying $\zeta^{D_{\bar{p}}} - 1$. ■*

Plaisted actually proved the special case $K = \mathbb{C}$ of the above lemma, in slightly different language, in [Pla84]. However, his proof extends verbatim to the more general family of fields detailed above.

A simple consequence of the resultant is that vanishing at a D^{th} root of unity is algebraically the same thing over \mathbb{C} or \mathbb{Q}_p , provided p lies in the right arithmetic progression.

⁷Throughout this paper, for Boolean expressions, we will always identify 0 with “False” and 1 with “True”.

Lemma 6.3 *Suppose $D \in \mathbb{N}$, $f \in \mathbb{Z}[x]$, and p is any prime congruent to 1 mod D . Then f vanishes at a complex D^{th} root of unity $\iff f$ vanishes at a D^{th} root of unity in \mathbb{Q}_p .*

Remark 6.4 *Note that $x^2 + x + 1$ vanishes at a 3^{rd} root of unity in \mathbb{C} , but has **no** roots at all in \mathbb{F}_5 or \mathbb{Q}_5 . So our congruence assumption on p is necessary. \diamond*

Proof of Lemma 6.3: First note that by our assumption on p , \mathbb{Q}_p has D distinct D^{th} roots of unity: This follows easily from Hensel’s Lemma and \mathbb{F}_p having D distinct D^{th} roots of unity. Since $\mathbb{Z} \hookrightarrow \mathbb{Q}_p$ and \mathbb{Q}_p contains all D^{th} roots of unity by construction, the equivalence then follows directly from Lemma 4.3. \blacksquare

Finally, let us recall a folkloric way (see, e.g., Plaisted’s proof of Theorem 5.1 in [Pla84]) to reduce systems of univariate polynomial equations to a single polynomial equation.

Proposition 6.5 *Given any $f_1, \dots, f_k \in \mathbb{Z}[x_1]$, let $d := \max_i \deg f_i$ and define $\tilde{f}(x_1)$ to be $x_1^d(f_1(x_1)f_1(1/x_1) + \dots + f_k(x_1)f_k(1/x_1))$. Then $f_1 = \dots = f_k = 0$ has a root on the complex unit circle if and only if \tilde{f} has a root on the complex unit circle.*

Proof: Trivial, upon observing that for any $x_1 \in \mathbb{C}$ with $|x_1| = 1$ and $i \in [k]$ we have $f_i(x_1)f_i(1/x_1) = |f_i(x_1)|^2$. \blacksquare

6.2 Random Primes in Arithmetic Progressions: Proving Theorem 1.9

The result below allows us to prove Theorem 1.9 and further tailor Plaisted’s clever reduction to our purposes. We let $\pi(x)$ denote the number of primes $\leq x$, and let $\pi(x; M, 1)$ denote the number of primes $\leq x$ that are congruent to 1 mod M .

The AGP Theorem (*very special case of [AGP94, Thm. 2.1, pg. 712]*) *There exist $x_0 > 0$ and an $\ell \in \mathbb{N}$ such that for each $x \geq x_0$, there is a subset $\mathcal{D}(x) \subset \mathbb{N}$ of finite cardinality ℓ with the following property: If $M \in \mathbb{N}$ satisfies $M \leq x^{2/5}$ and $a \not\equiv 1 \pmod{M}$ for all $a \in \mathcal{D}(x)$ then $\pi(x; M, 1) \geq \frac{\pi(x)}{2\varphi(M)}$. \blacksquare*

For those familiar with [AGP94, Thm. 2.1, pg. 712], the result above follows immediately upon specializing the parameters there as follows:

$$(A, \varepsilon, \delta, y, a) = (49/20, 1/2, 2/245, x, 1)$$

(see also [vzGKS96, Fact 4.9]).

The AGP Theorem enables us to construct random primes from certain arithmetic progressions with high probability. An additional ingredient that will prove useful is the famous *AKS algorithm* for deterministic polynomial-time primality checking [AKS02]. Consider now the following algorithm.

Algorithm 6.6

Input: *A constant $\delta > 0$, a failure probability $\varepsilon \in (0, 1/2)$, a positive integer n , and the constants x_0 and ℓ from the AGP Theorem.*

Output: *An increasing sequence $\bar{p} = (p_j)_{j=1}^n$ of primes, and $c \in \mathbb{N}$, such that $p := 1 + c \prod_{i=1}^n p_i$*

satisfies $\log p = O(n \log(n) + \log(1/\varepsilon))$ and, with probability $1 - \varepsilon$, p is prime. In particular, the output always gives a true declaration as to the primality of p .

Description:

0. Let $L := \lceil 2/\varepsilon \rceil \ell$ and compute the first nL primes p_1, \dots, p_{nL} in increasing order.
1. Define (but do not compute) $M_j := \prod_{k=(j-1)n+1}^{jn} p_k$ for any $j \in \mathbb{N}$. Then compute M_L, M_i for a uniformly random $i \in [L]$, and $x := \max \left\{ x_0, 17, 1 + M_L^{5/2} \right\}$.
2. Compute $K := \lfloor (x - 1)/M_i \rfloor$ and $J := \lceil 2 \log(2/\varepsilon) \log x \rceil$.
3. Pick uniformly random $c \in [K]$ until one either has $p := 1 + cM_i$ prime, or one has J such numbers that are each composite (using primality checks via the AKS algorithm along the way).
4. If a prime p was found then output

“ $1 + c \prod_{j=(i-1)n+1}^{in} p_j$ is a prime that works!”

 and stop. Otherwise, stop and output

“I have failed to find a suitable prime. Please forgive me.” \diamond

Remark 6.7 In our algorithm above, it suffices to find integer approximations to the underlying logarithms and square-roots. In particular, we restrict to algorithms that can compute the $\log_2 \mathcal{L}$ most significant bits of $\log \mathcal{L}$, and the $\frac{1}{2} \log_2 \mathcal{L}$ most significant bits of $\sqrt{\mathcal{L}}$, using $O((\log \mathcal{L})(\log \log \mathcal{L}) \log \log \log \mathcal{L})$ bit operations. Arithmetic-Geometric Mean Iteration and (suitably tailored) Newton Iteration are algorithms that respectively satisfy our requirements (see, e.g., [Ber03] for a detailed description). \diamond

Proof of Theorem 1.9: It clearly suffices to prove that Algorithm 6.6 is correct, has a success probability that is at least $1 - \varepsilon$, and works within

$$O\left(\left(\frac{n}{\varepsilon}\right)^{\frac{3}{2}+\delta} + (n \log(n) + \log(1/\varepsilon))^{7+\delta}\right)$$

randomized bit operations, for any $\delta > 0$. These assertions are proved directly below. \blacksquare

Proving Correctness and the Success Probability Bound for Algorithm 6.6: First observe that M_1, \dots, M_L are relatively prime. So at most ℓ of the M_i will be divisible by elements of $\mathcal{D}(x)$. Note also that $K \geq 1$ and $1 + cM_i \leq 1 + KM_i \leq 1 + ((x - 1)/M_i)M_i = x$ for all $i \in [L]$ and $c \in [K]$.

Since $x \geq x_0$ and $x^{2/5} \geq (x - 1)^{2/5} \geq \left(M_i^{5/2}\right)^{2/5} = M_i$ for all $i \in [L]$, the AGP Theorem implies that with probability at least $1 - \frac{\varepsilon}{2}$ (since $i \in \lceil [2/\varepsilon] \ell \rceil$ is uniformly random), the arithmetic progression $\{1 + M_i, \dots, 1 + KM_i\}$ contains at least $\frac{\pi(x)}{2\varphi(M_i)} \geq \frac{\pi(x)}{2M_i}$ primes. In which case, the proportion of numbers in $\{1 + M_i, \dots, 1 + KM_i\}$ that are prime is $\frac{\pi(x)}{2KM_i} > \frac{\pi(x)}{2+2KM_i} > \frac{x/\log x}{2x} = \frac{1}{2\log x}$, since $\pi(x) > x/\log x$ for all $x \geq 17$ [BS96, Thm. 8.8.1, pg. 233]. So let us now assume that i is fixed and M_i is not divisible by any element of $\mathcal{D}(x)$.

Recalling the inequality $(1 - \frac{1}{t})^{ct} \leq e^{-c}$ (valid for all $c \geq 0$ and $t \geq 1$), we then see that the AGP Theorem implies that the probability of *not* finding a prime of the form $p = 1 + cM_i$ after picking J uniformly random $c \in [K]$ is bounded above by $\left(1 - \frac{1}{2\log x}\right)^J \leq \left(1 - \frac{1}{2\log x}\right)^{2\log(2/\varepsilon)\log x} \leq e^{-\log(2/\varepsilon)} = \frac{\varepsilon}{2}$.

In summary, with probability $\geq 1 - \frac{\varepsilon}{2} - \frac{\varepsilon}{2} = 1 - \varepsilon$, Algorithm 6.6 picks an i with M_i not divisible by any element of $\mathcal{D}(x)$ and a c such that $p := 1 + cM_i$ is prime. In particular, we clearly have that

$$\log p = O(\log(1 + KM_i)) = O(n \log(n) + \log(1/\varepsilon)). \blacksquare$$

Complexity Analysis of Algorithm 6.6: Let $L' := nL$ and, for the remainder of our proof, let p_i denote the i^{th} prime. Since $L' \geq 6$, we have that

$$p_{L'} \leq L'(\log(L') + \log \log L')$$

by [BS96, Thm. 8.8.4, pg. 233]. Recall that the primes in $[\mathcal{L}]$ can be listed simply by deleting all multiples of 2 in $[\mathcal{L}]$, then deleting all multiples of 3 in $[\mathcal{L}]$, and so on until one reaches multiples of $\lfloor \sqrt{\mathcal{L}} \rfloor$. (This is the classic sieve of Eratosthenes.) Recall also that one can multiply an integer in $[\mu]$ and an integer $[\nu]$ within

$$O((\log \mu)(\log \log \nu)(\log \log \log \nu) + (\log \nu)(\log \log \mu) \log \log \log \mu)$$

bit operations (see, e.g., [BS96, Table 3.1, pg. 43]). So let us define the function $\lambda(a) := (\log \log a) \log \log \log a$.

Step 0: By our preceding observations, it is easily checked that Step 0 takes $O(L'^{3/2} \log^3 L')$ bit operations.

Step 1: This step consists of $n - 1$ multiplications of primes with $O(\log L')$ bits (resulting in M_L , which has $O(n \log L')$ bits), multiplication of a small power of M_L by a square root of M_L , division by an integer with $O(n \log L')$ bits, a constant number of additions of integers of comparable size, and the generation of $O(\log L)$ random bits. Employing Remark 2.4 along the way, we thus arrive routinely at an estimate of

$$O(n^2(\log L')\lambda(L') + \log(1/\varepsilon)\lambda(1/\varepsilon))$$

for the total number of bit operations needed for Step 1.

Step 2: Similar to our analysis of Step 1, we see that Step 2 has bit complexity

$$O((n \log(L') + \log(1/\varepsilon))\lambda(n \log L')).$$

Step 3: This is our most costly step: Here, we require

$$O(\log K) = O(n \log(L') + \log(1/\varepsilon))$$

random bits and $J = O(\log x) = O(n \log(L') + \log(1/\varepsilon))$ primality tests on integers with

$$O(\log(1 + cM_i)) = O(n \log(L') + \log(1/\varepsilon))$$

bits. By an improved version of the AKS primality testing algorithm [AKS02, LP05] (which takes $O(N^{6+\delta})$ bit operations to test an N bit integer for primality), Step 3 can then clearly be done within

$$O((n \log(L') + \log(1/\varepsilon))^{7+\delta})$$

bit operations, and the generation of $O(n \log(L') + \log(1/\varepsilon))$ random bits.

Step 4: This step clearly takes time on the order of the number of output bits, which is just $O(n \log(n) + \log(1/\varepsilon))$ as already observed earlier.

Conclusion: We thus see that Step 0 and Step 3 dominate the complexity of our algorithm, and we are left with an overall randomized complexity bound of

$$\begin{aligned} & O\left(L'^{3/2} \log^3(L') + (n \log(L') + \log(1/\varepsilon))^{7+\delta}\right) \\ &= O\left(\left(\frac{n}{\varepsilon}\right)^{3/2} \log^3(n/\varepsilon) + (n \log(n) + \log(1/\varepsilon))^{7+\delta}\right) \\ &= O\left(\left(\frac{n}{\varepsilon}\right)^{\frac{3}{2}+\delta} + (n \log(n) + \log(1/\varepsilon))^{7+\delta}\right) \end{aligned}$$

randomized bit operations. \blacksquare

6.3 The Proof of Theorem 1.2

We will prove a (**ZPP**) randomized polynomial-time reduction from **3CNFSAT** to $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathcal{P})$, making use of the intermediate input families $\{(\mathbb{Z}[x])^k \mid k \in \mathbb{N}\} \times \mathcal{P}$ and $\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\} \times \mathcal{P}$ along the way.

Toward this end, suppose $B(y) := C_1(y) \wedge \cdots \wedge C_k(y)$ is any **3CNFSAT** instance. The polynomial system $(\rho_{\bar{p}}(C_1), \dots, \rho_{\bar{p}}(C_k))$, for \bar{p} the first n primes (employing Lemma 6.2), then clearly yields $\text{FEAS}_{\mathbb{C}}(\{(\mathbb{Z}[x])^k \mid k \in \mathbb{N}\}) \in \mathbf{P} \implies \mathbf{P} = \mathbf{NP}$. Composing this reduction with Proposition 6.5 we then immediately obtain $\text{FEAS}_{\mathbb{C}}(\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\}) \in \mathbf{P} \implies \mathbf{P} = \mathbf{NP}$.

We now need only find a means of transferring from \mathbb{C} to \mathbb{Q}_p . This we do by preceding our reductions above by a judicious (possibly new) choice of \bar{p} : by applying Theorem 1.9 with $\varepsilon = 1/3$ (cf. Lemma 6.3) we immediately obtain the implication

$$\text{FEAS}_{\mathbb{Q}_{\text{primes}}}((\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\}) \times \mathcal{P}) \in \mathbf{ZPP} \implies \mathbf{NP} \subseteq \mathbf{ZPP}.$$

To conclude, observe that any root $(x, y) \in \mathbb{Q}_p^2 \setminus \{(0, 0)\}$ of the quadratic form $x^2 - py^2$ must satisfy $2\text{ord}_p x = 1 + 2\text{ord}_p y$ (an impossibility). So the only p -adic rational root of $x^2 - py^2$ is $(0, 0)$ and we easily obtain a polynomial-time reduction from $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}((\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\}) \times \mathcal{P})$ to $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathcal{P})$: simply map any instance $(f(x), x^D - 1, p)$ of the former problem to $(f(x)^2 - (x^D - 1)^2, p)$. So we have proved the first implication.

To prove the second (conditional) implication, we simply repeat our last proof, replacing our AGP Theorem-based algorithm with a simple brute-force search. More precisely, letting $D := 2 \cdot 3 \cdots p_n$, we simply test the integers $1 + kD$ for primality, starting with $k = 1$ until one finds a prime. If Wagstaff's Conjecture is true then we need not proceed any farther than $k = O\left(\frac{\varphi(D)}{D} \log^2 D\right)$. (Note that $\frac{1}{2} \leq \frac{\varphi(D)}{D} < 1$ for all $D \geq 2$.) Using the AKS algorithm, this brute-force search clearly has (deterministic) complexity polynomial in $\log D$ which in turn is polynomial in n . ■

7 The Proof of Proposition 1.3

Let us first recall that in the *Subset Sum Problem* one is given nonzero integers c_1, \dots, c_n and one must decide whether there is a non-empty subset $I \subseteq \{1, \dots, n\}$ with $\sum_{i \in I} c_i = 0$. Using $\sum_{i=1}^n \log(2 + |c_i|)$ (or the number of bits needed to specify c_1, \dots, c_n) as the underlying input size, the Subset Sum Problem is then **NP**-complete [GJ79].

We will present a polynomial-time reduction from the Subset Sum Problem to $\text{FEAS}_{\mathbb{Q}_p}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n, n+1}^*\right)$, thereby proving Proposition 1.3.

Proof of Proposition 1.3 (Poonen): Suppose we fix a prime p . Then, given any instance of the Subset Sum Problem as described above, we can create an instance of $\text{FEAS}_{\mathbb{Q}_p}\left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n, n+1}^*\right)$ as follows: Let $\ell \geq 3$ be the smallest integer such that p^ℓ exceeds $\sum_{i=1}^n |c_i|$. Also let $P := p^{\ell-1}(p-1)$, which is the order of $(\mathbb{Z}/p^\ell \mathbb{Z})^*$ (cf. Lemma 3.1). By Proposition 2.1, we can construct ℓ and P in polynomial-time. We then consider the polynomial $f(x) := c_1 x_1^P + \cdots + c_n x_n^P$, which clearly has size linear in the size of our Subset Sum instance. Note also that by homogeneity, f has a non-trivial root in \mathbb{Q}_p^n if and only if f has

a root $(x_1, \dots, x_n) \in \mathbb{Z}_p^n$ with $\text{ord}_p x_i = 0$ for some i . In particular, for each j , the value of $x_j^P \pmod{p^\ell}$ is 0 or 1 according as $\text{ord}_p x_j > 0$ or $\text{ord}_p x_j = 0$.

Now, should f have such a root in \mathbb{Z}_p^n , then f must have a root in $(\mathbb{Z}/p^\ell\mathbb{Z})^n$. A root of f in \mathbb{Q}_p^n thus induces a non-empty I with $\sum_{i \in I} c_i = 0 \pmod{p^\ell}$. By our choice of ℓ , this in turn implies that $\sum_{i \in I} c_i = 0$ as an integer.

Conversely, a non-empty I with $\sum_{i \in I} c_i = 0$ in \mathbb{Z} implies a zero-one vector $x = (x_1, \dots, x_n)$ that is a root of f in \mathbb{Z}^n and thus a root of f in \mathbb{Q}_p^n .

We have thus shown that a Subset Sum instance can always be converted in polynomial-time to a particular kind of n -variate n -nomial f (with size linear in the Subset Sum Instance) such that the Subset Sum Instance has a “Yes” answer if and only if f has a root in \mathbb{Q}_p^n . The

latter decision problem can then be reduced to n instance of $\text{FEAS}_{\mathbb{Q}_p} \left(\bigcup_{n \in \mathbb{N}} \mathcal{F}_{n, n+1}^* \right)$ as follows: merely check if any of the dehomogenizations $f(1, x_2, \dots, x_n), \dots, f(x_1, \dots, x_{n-1}, 1)$ have a root in \mathbb{Q}_p^{n-1} . Each of these dehomogenizations is an honest $(n-1)$ -variate n -nomial, so we are done. ■

8 The Final Corollaries

8.1 Proof of Corollary 1.5

Our proof of Assertion (1) of Theorem 1.4 is, in retrospect, a polynomial-time reduction from $\text{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{F}_{1,3})$ to $\text{FEAS}_{\mathbb{Z}/p^\ell\mathbb{Z}}(\mathcal{F}_{1,3})$ with $\ell = O(p(\text{size}(f)/\log p)^4 \log \text{size}(f))$. Combining this reduction with the hypothesis of Corollary 1.5 then clearly implies that $\text{FEAS}_{\mathbb{Q}_p}(\mathcal{F}_{1,3})$ can be solved in time polynomial in $p + (\text{size}(f)/\log p)^4 \log \text{size}(f)$, so we are done. ■

8.2 Proof of Corollary 4.6

By Lemma 4.3 we know that $\bar{\Delta}_{\mathcal{A}}(f)$ has degree at most $2d-1$ in the coefficients of f . We also know that for any fixed $f \in T_{\mathcal{A}}(H)$, $\bar{\Delta}_{\mathcal{A}}(f)$ is an integer as well, and is thus divisible by no more than $1 + (2d-1) \log(mH)$ primes. (The last assertion follows from Lemma 4.3 again, and the elementary fact that an integer N has no more than $1 + \log N$ distinct prime factors.) Recalling that $\pi(x) > x/\log x$ for all $x \geq 17$ [BS96, Thm. 8.8.1, pg. 233], we thus obtain that the fraction of primes $\leq H$ dividing a nonzero $\bar{\Delta}_{\mathcal{A}}(f)$ is bounded above by $\frac{1 + (2d-1) \log(mH)}{H/\log H}$.

Now by the Schwartz-Zippel Lemma [Sch80], $\bar{\Delta}_{\mathcal{A}}(f)$ vanishes for at most $(2d-1)m(2H)^{m-1}$ selections of coefficients from $\{-H, \dots, H\}$. In other words, $\bar{\Delta}_{\mathcal{A}}(f) = 0$ for a fraction of at most $\frac{(2d-1)m}{2H+1}$ of the polynomials in $T_{\mathcal{A}}(H)$.

Combining our last two fractional bounds, we are done. ■

Acknowledgements

We thank Bjorn Poonen for his permission to include Proposition 1.3 and its proof. We also thank Jan Denef for pointing out the reference [BMc67], and Matt Papanikolas and

Paula Tretkoff for valuable discussions on the Weil Conjectures. Finally, we thank the two anonymous referees for their valuable comments that helped improve this paper.

References

- [AKS02] Agrawal, Manindra; Kayal, Neeraj; and Saxena, Nitin, “*PRIMES is in P*,” Ann. of Math. (2) 160 (2004), no. 2, pp. 781–793.
- [AGP94] Alford, W. R.; Granville, Andrew; and Pomerance, Carl, “*There are Infinitely Many Carmichael Numbers*,” Ann. of Math. (2) 139 (1994), no. 3, pp. 703–722.
- [Art65] Artin, Emil, *The collected papers of Emil Artin*, edited by Serge Lang and John T. Tate, Addison–Wesley Publishing Co., Inc., Reading, Mass.-London, 1965.
- [AI11] Avendaño, Martin and Ibrahim, Ashraf, “*Multivariate ultrametric root counting*,” to appear in an upcoming AMS Contemporary Mathematics volume, also downloadable from www.math.tamu.edu/~avendano/murc.pdf.
- [AK65] Ax, James and Kochen, Simon, “*Diophantine problems over local fields I*,” Amer. J. Math. 87, 1965, pp. 605–630.
- [BS96] Bach, Eric and Shallit, Jeff, *Algorithmic Number Theory, Vol. I: Efficient Algorithms*, MIT Press, Cambridge, MA, 1996.
- [BHPR11] Bastani, Osbert; C. Hillar, D. Popov, and J. M. Rojas, “*Randomization, Sums of Squares, and Faster Real Root Counting for Tetranomials and Beyond*,” submitted for publication, 2011. Also available as Math ArXiv preprint 1101.2642 .
- [BPR06] Basu, Saugata; Pollack, Ricky; and Roy, Marie-Francoise, *Algorithms in Real Algebraic Geometry*, Algorithms and Computation in Mathematics, vol. 10, Springer-Verlag, 2006.
- [Ber03] Bernstein, Daniel J., “*Computing Logarithm Intervals with the Arithmetic-Geometric Mean Iterations*,” available from <http://cr.yp.to/papers.html> .
- [BRS09] Bihan, Frederic; Rojas, J. Maurice; Stella, Case E., “*Faster Real Feasibility via Circuit Discriminants*,” proceedings of International Symposium on Symbolic and Algebraic Computation (ISSAC 2009, July 28–31, Seoul, Korea), pp. 39–46, ACM Press, 2009.
- [BMc67] Birch, B. J. and McCann, K., “*A Criterion for the p -adic Solubility of Diophantine Equations*,” Quart. J. Math. Oxford (2), 18 (1967), pp. 59–63.
- [Can88] Canny, John F., “*Some Algebraic and Geometric Computations in PSPACE*,” Proc. 20th ACM Symp. Theory of Computing, Chicago (1988), ACM Press.
- [CG00] Cantor, David G. and Gordon, Daniel M., “*Factoring polynomials over p -adic fields*,” Algorithmic number theory (Leiden, 2000), pp. 185–208, Lecture Notes in Comput. Sci., 1838, Springer, Berlin, 2000.

- [Cha08] Chambert-Loir, Antoine, “*Compter (rapidement) le nombre de solutions d’équations dans les corps finis*,” Séminaire Bourbaki, Vol. 2006/2007, Astérisque No. 317 (2008), Exp. No. 968, vii, pp. 39–90.
- [Che04] Cheng, Qi, “*Straight Line Programs and Torsion Points on Elliptic Curves*,” Computational Complexity, 12:3-4, Sept. 2004, pp. 150–161.
- [CDV06] Castryck, Wouter; Deneff, Jan; and Vercauteren, Frederik, “*Computing Zeta Functions of Nondegenerate Curves*,” International Mathematics Research Papers, vol. 2006, article ID 72017, 2006.
- [Chi91] Chistov, Alexander L., “*Efficient Factoring [of] Polynomials over Local Fields and its Applications*,” in I. Satake, editor, Proc. 1990 International Congress of Mathematicians, pp. 1509–1519, Springer-Verlag, 1991.
- [Coh94] Cohen, Henri, *A course in computational algebraic number theory*, Graduate Texts in Mathematics, 138, Springer-Verlag, Berlin, 1993.
- [Coh69] Cohen, Paul J., “*Decision procedures for real and p -adic fields*,” Comm. Pure Appl. Math. 22 (1969), pp. 131–151.
- [C-T98] Colliot-Thelene, Jean-Louis, “*The Hasse principle in a pencil of algebraic varieties*,” Number theory (Tiruchirapalli, 1996), pp. 19–39, Contemp. Math., 210, Amer. Math. Soc., Providence, RI, 1998.
- [CS99] Cucker, Felipe and Smale, Steve, “*Complexity estimates depending on condition and round-off error*,” J. ACM 46 (1999), no. 1, pp. 113–184.
- [DLPvG00] *Hilbert’s Tenth Problem: Relations with Arithmetic and Algebraic Geometry*, Papers from a workshop held at Ghent University, Ghent, November 2–5, 1999. Edited by Jan Deneff, Leonard Lipshitz, Thanases Pheidas and Jan Van Geel. Contemporary Mathematics, 270, American Mathematical Society, Providence, RI, 2000.
- [FK88] Freitag, Eberhard and Kiehl, Reinhardt, “*Etale cohomology and the Weil conjecture*,” (translated from the German by Betty S. Waterhouse and William C. Waterhouse, with an historical introduction by J. A. Dieudonné), Ergebnisse der Mathematik und ihrer Grenzgebiete (3), 13, Springer-Verlag, Berlin, 1988.
- [GJ79] Garey, Michael R. and Johnson, David S. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, A Series of Books in the Mathematical Sciences, W. H. Freeman and Co., San Francisco, Calif., 1979, x+338 pp.
- [vzGKS96] von zur Gathen, Joachim; Karpinski, Marek; and Shparlinski, Igor, “*Counting curves and their projections*,” Computational Complexity 6, no. 1 (1996/1997), pp. 64–99.
- [GKZ94] Gel’fand, Israel Moseyevitch; Kapranov, Misha M.; and Zelevinsky, Andrei V.; *Discriminants, Resultants and Multidimensional Determinants*, Birkhäuser, Boston, 1994.

- [Glo06] Glöckner, Helge, “*Implicit functions from topological vector spaces to Banach spaces*,” Israel J. Math. 155 (2006), pp. 205–252.
- [Gre74] Greenberg, Marvin J., “*Strictly local solutions of Diophantine equations*,” Pacific J. Math. 51 (1974), pp. 143–153.
- [Has24] Hasse, H., “*Darstellbarkeit von Zahlen durch quadratische Formen in einem beliebigen algebraischen Zahlkörper*,” J. Reine Angew. Math. 153 (1924), pp. 113–130.
- [H-B10] Heath-Brown, D. R., “*Zeros of p -adic Forms*,” Proc. Lond. Math. Soc. (3) 100 (2010), no. 2, pp. 560–584.
- [Kal03] Kaltofen, Erich, “*Polynomial factorization: a success story*,” In ISSAC 2003 Proc. 2003 Internat. Symp. Symbolic Algebraic Comput. (New York, N.Y., 2003), J. R. Sendra, Ed., ACM Press, pp. 3–4.
- [Koi11] Koiran, Pascal, *personal e-mail communication*, March, 2011.
- [Lau04] Lauder, Alan G. B., “*Counting solutions to equations in many variables over finite fields*,” Found. Comput. Math. 4 (2004), no. 3, pp. 221–267.
- [LW08] Lauder, Alan G. B. and Wan, Daqing, “*Counting points on varieties over finite fields of small characteristic*,” Algorithmic number theory: lattices, number fields, curves and cryptography, pp. 579–612, Math. Sci. Res. Inst. Publ., 44, Cambridge Univ. Press, Cambridge, 2008.
- [Len99] _____, “*On the Factorization of Lacunary Polynomials*,” Number Theory in Progress, Vol. 1 (Zakopane-Kóscielisko, 1997), pp. 277–291, de Gruyter, Berlin, 1999.
- [LLL82] Lenstra, Arjen K.; Lenstra (Jr.), Hendrik W.; Lovász, L., “*Factoring polynomials with rational coefficients*,” Math. Ann. 261 (1982), no. 4, pp. 515–534.
- [LP05] Lenstra (Jr.), Hendrik W., and Pomerance, Carl, “*Primality Testing with Gaussian Periods*,” manuscript, downloadable from <http://www.math.dartmouth.edu/~carlp/PDF/complexity12.pdf>
- [Lew52] Lewis, D. J., “*Cubic homogeneous polynomials over p -adic number fields*,” Ann. of Math. (2) 56, no. 3, November 1952, pp. 473–478.
- [MW99] Maller, Michael and Whitehead, Jennifer, “*Efficient p -adic cell decomposition for univariate polynomials*,” J. Complexity 15 (1999), pp. 513–525.
- [Mig82] Mignotte, Maurice, “*Some Useful Bounds*,” in Computer Algebra: Symbolic and Algebraic Computation, 2nd ed., (edited by B. Buchberger, G. E. Collins, and R. Loos, in cooperation with R. Albrecht), Springer-Verlag 1982.
- [Nes03] Nesterenko, Yuri, “*Linear forms in logarithms of rational numbers*,” Diophantine approximation (Cetraro, 2000), pp. 53–106, Lecture Notes in Math., 1819, Springer, Berlin, 2003.

- [OSvK00] Overmars, Mark; Schwarzkopf, Otfried; and van Kreveld, Marc, *Computational Geometry: Algorithms and Applications*, Springer Verlag, 2000.
- [Pap95] Papadimitriou, Christos H., *Computational Complexity*, Addison-Wesley, 1995.
- [Pla84] Plaisted, David A., “*New NP-Hard and NP-Complete Polynomial and Integer Divisibility Problems*,” *Theoret. Comput. Sci.* 31 (1984), no. 1–2, 125–138.
- [Poo01] Poonen, Bjorn, “*An explicit algebraic family of genus-one curves violating the Hasse principle*,” 21st Journées Arithmétiques (Rome, 2001), *J. Théor. Nombres Bordeaux* 13 (2001), no. 1, pp. 263–274.
- [Poo06] _____, “*Heuristics for the Brauer-Manin Obstruction for Curves*,” *Experimental Mathematics*, Volume 15, Issue 4 (2006), pp. 415–420.
- [RS02] Rahman, Qazi Ibadur; and Schmeisser, Gerhard, *Analytic Theory of Polynomials*, Clarendon Press, London Mathematical Society Monographs 26, 2002.
- [Rob00] Robert, Alain M., *A course in p-adic analysis*, Graduate Texts in Mathematics, 198, Springer-Verlag, New York, 2000.
- [Roj02] Rojas, J. Maurice, “*Additive Complexity and the Roots of Polynomials Over Number Fields and p-adic Fields*,” *Proceedings of ANTS-V (5th Annual Algorithmic Number Theory Symposium, University of Sydney, July 7–12, 2002)*, *Lecture Notes in Computer Science #2369*, Springer-Verlag (2002), pp. 506–515.
- [Sch80] Schwartz, Jacob T., “*Fast Probabilistic Algorithms for Verification of Polynomial Identities*,” *J. of the ACM* 27, 701–717, 1980.
- [Ser73] Serre, Jean-Pierre, “*A course in arithmetic*,” *Graduate Texts in Mathematics*, No. 7, Springer-Verlag, New York-Heidelberg, 1973.
- [Sto00] Storjohann, Arne, “*Algorithms for Matrix Canonical Forms*,” doctoral dissertation, Swiss Federal Institute of Technology, Zurich, 2000.
- [Tar51] Tarski, Alfred, *A Decision Method for Elementary Algebra and Geometry*, prepared for publication by J. C. C. McKinsey, University of California Press, Berkeley and Los Angeles, California, 1951.
- [Ter66] Terjanian, Guy, “*Un contre-exemple à une conjecture d’Artin*,” *C. R. Acad. Sci. Paris Sér. A-B* 262, 1966, A612.
- [Wag79] Wagstaff, Samuel S., Jr., “*Greatest of the Least Primes in Arithmetic Progressions Having a Given Modulus*,” *Mathematics of Computation*, Vol. 33, No. 147, July 1979, pp. 1073–1080.
- [Wei49] Weil, André, “*Numbers of solutions of equations in finite fields*,” *Bull. Amer. Math. Soc.* 55, (1949), pp. 497–508.
- [Yu94] Yu, Kunrui, “*Linear forms in p-adic logarithms III*,” *Compositio Mathematica*, tome 91, no. 3 (1994), pp. 241–276.