Randomized $\mathbf{NP}$-Completeness for p-adic Rational Roots of Sparse Polynomials in One Variable

Martín Avendaño[*]
TAMU 3368
Mathematics Dept.
College Station, TX 77843-3368, USA
mavendar@yahoo.com.ar

Ashraf Ibrahim[*]
TAMU 3141
Aerspace Engineering Dept.
College Station, TX 77843-3141, USA
ibrahim@aero.tamu.edu

J. Maurice Rojas[*]
TAMU 3368
Mathematics Dept.
College Station, TX 77843-3368, USA
rojas@math.tamu.edu

Korben Rusek[*]
TAMU 3368
Mathematics Dept.
College Station, TX 77843-3368, USA
korben.rusek@gmail.com

ABSTRACT

Relative to the sparse encoding, we show that deciding whether a univariate polynomial has a p-adic rational root can be done in **NP** for most inputs. We also prove a sharper complexity upper bound of **P** for polynomials with suitably generic p-adic Newton polygon. We thus improve the best previous complexity upper bound of **EXPTIME**. We also prove an unconditional complexity lower bound of **NP**-hardness with respect to randomized reductions, for general univariate polynomials. The best previous lower bound assumed an unproved hypothesis on the distribution of primes in arithmetic progression. We also discuss how our results complement analogous results over the real numbers.

Categories and Subject Descriptors

F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems—*Number-theoretic computations*

Keywords

sparse, p-adic, feasibility, NP, arithmetic progression

1. INTRODUCTION

The fields $\mathbb{R}$ and $\mathbb{Q}_p$ (the reals and the p-adic rationals) bear more in common than just completeness with respect to a metric: increasingly, complexity results for one field have inspired and motivated analogous results in the other (see, e.g., [Coh69, DvdD88] and the pair of works [Kho91] and [Roj04]). We continue this theme by transposing recent algorithmic results for sparse polynomials over the real numbers [BRS09] to the p-adic rationals, sharpening the underlying complexity bounds along the way (see Theorem 1.5 below).

For any commutative ring R with multiplicative identity, let $\mathtt{FEAS}_R$ — the R-**feasibility problem** (a.k.a. Hilbert's Tenth Problem over R [DLPvG00]) — denote the problem of deciding whether an input $F \in \bigcup_{k,n\in\mathbb{N}}(\mathbb{Z}[x_1,\ldots,x_n])^k$ has a root in R^n. (The underlying input size is clarified in Definition 1.1 below.) Observe that $\mathtt{FEAS}_{\mathbb{R}}$, $\mathtt{FEAS}_{\mathbb{Q}}$, and $\{\mathtt{FEAS}_{\mathbb{F}_q}\}_{q \text{ a prime power}}$ are central problems respectively in algorithmic real algebraic geometry, algorithmic number theory, and cryptography.

For any prime p and $x \in \mathbb{Z}$, recall that the p-**adic valuation**, $\mathrm{ord}_p x$, is the greatest k such that $p^k | x$. We can extend $\mathrm{ord}_p(\cdot)$ to $\mathbb{Q}$ by $\mathrm{ord}_p \frac{a}{b} := \mathrm{ord}_p(a) - \mathrm{ord}_p(b)$ for any $a, b \in \mathbb{Z}$; and we let $|x|_p := p^{-\mathrm{ord}_p x}$ denote the p-**adic norm**. The norm $|\cdot|_p$ defines a natural metric satisfying the ultrametric inequality and $\mathbb{Q}_p$ is, tersely, the completion of $\mathbb{Q}$ with respect to this metric. $|\cdot|_p$ and $\mathrm{ord}_p(\cdot)$ extend 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].

We will also need to recall the following containments of complexity classes: $\mathbf{P} \subseteq \mathbf{ZPP} \subseteq \mathbf{NP} \subseteq \cdots \subseteq \mathbf{EXPTIME}$, and the fact that the properness of **every** inclusion above (save $\mathbf{P} \subsetneqq \mathbf{EXPTIME}$) is a major open problem [Pap95].

1.1 The Ultrametric Side: Relevance and Results

Algorithmic results over the p-adics are useful in many settings: polynomial-time factoring algorithms over $\mathbb{Q}[x]$ [LLL82], computational complexity [Roj02], studying prime ideals in number fields [Coh94, Ch. 4 & 6], elliptic curve cryptography [Lau04], and the computation of zeta functions [CDV06]. Also, much work has gone into using p-adic methods to algorithmically detect rational points on algebraic plane curves via variations of the **Hasse Principle**[1] (see, e.g., [C-T98, Poo06]). However, 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 [Len99b]. So we focus on precise complexity bounds for polynomials in one variable.

[*]Partially supported by Rojas' NSF CAREER grant DMS-0349309. Avendaño, Rojas, and Rusek were also partially supported by NSF MCS grant DMS-0915245. Rojas and Rusek were also partially supported by Sandia National Laboratories and DOE ASCR grant DE-SC0002505. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

ISSAC'10, 25–28 July 2010, Munich, Germany.
Copyright 2010 ACM 978-1-4503-0150-3/10/0007 ...$5.00.

[1] If $f \in K[x_1,\ldots,x_n]$ is any polynomial and Z_K is its zero set in K^n, then the Hasse Principle is the implication [$Z_{\mathbb{C}}$ smooth, $Z_{\mathbb{R}} \neq \emptyset$, and $Z_{\mathbb{Q}_p} \neq \emptyset$ for all primes p] $\Longrightarrow Z_{\mathbb{Q}} \neq \emptyset$. 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., [Poo01a]).

Definition 1.1. *Let $f(x):=\sum_{i=1}^{m} c_i x^{a_i}\in\mathbb{Z}[x]$ satisfy $c_i\neq 0$ for all i, with the a_i pair-wise distinct. We call such an f a* **(univariate) m-nomial**. *Let us also define* $\mathrm{size}(f):=\sum_{i=1}^{m}\log_2\left[(2+|c_i|)(2+|a_i|)\right]$ *and, for any $F:=(f_1,\ldots,f_k)\in(\mathbb{Z}[x])^k$, we define $\mathrm{size}(F):=\sum_{i=1}^{k}\mathrm{size}(f_i)$. Finally, we let $\mathcal{F}_{1,m}$ denote the subset of $\mathbb{Z}[x]$ consisting of polynomials with exactly m monomial terms* ⋄

The degree, $\deg f$, of a polynomial f can sometimes be exponential in $\mathrm{size}(f)$ for certain families of f, e.g., $d\geq\frac{2^{\mathrm{size}\left(1+5x^{126}+x^d\right)}}{2^{16}}$ for all $d\geq 127$. Note also that $\mathbb{Z}[x]$ is the disjoint union $\bigsqcup_{m\geq 0}\mathcal{F}_{1,m}$.

Definition 1.2. *Let $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}$ denote the problem of deciding, for an input polynomial system $F\in\bigcup_{k,n\in\mathbb{N}}(\mathbb{Z}[x_1,\ldots,x_n])^k$* **and** *an input prime p, whether F has a root in $\mathbb{Q}_p^n$. Also let $\mathbb{P}\subset\mathbb{N}$ denote the set of primes and, when $\mathcal{I}$ is a family of such pairs (F,p), we let $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathcal{I})$ denote the restriction of $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}$ to inputs in $\mathcal{I}$. The underlying input sizes for $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}$ and $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathcal{I})$ shall then be $\mathrm{size}_p(F):=\mathrm{size}(F)+\log p$ (cf. Definition 1.1).* ⋄

To state our main results, we will also need a bit of arithmetic tropicalia.

Definition 1.3. *Given any polynomial $f(x):=\sum_{i=1}^{m}c_i x^{a_i}\in\mathbb{Z}[x]$, we define its* **$p$-adic Newton polygon**, $\mathrm{Newt}_p(f)$, *to be the convex hull of*[2] *the points $\{(a_i,\mathrm{ord}_p c_i)\mid i\in\{1,\ldots,m\}\}$. Also, a face of a polygon $P\subset\mathbb{R}^2$ is called* **lower** *iff it has an inner normal with positive last coordinate, and the* **lower hull** *of P is simply the union of all its lower edges. Finally, the polynomial given by summing the terms of f corresponding to points of the form $(a_i,\mathrm{ord}_p c_i)$ in some fixed lower face of $\mathrm{Newt}_p(f)$ is called a* **(p-adic) lower polynomial**. ⋄

Example 1.4. *For the polynomial $f(x)$ defined as $36-8868x+29305x^2-35310x^3+18240x^4-3646x^5+243x^6$, the polygon $\mathrm{Newt}_3(f)$ 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.* ⋄

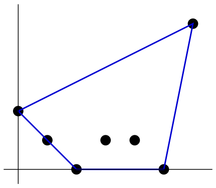

While there are now randomized algorithms for factoring $f\in\mathbb{Z}[x]$ over $\mathbb{Q}_p[x]$ with expected complexity polynomial in $\mathrm{size}_p(f)+\deg(f)$ [CG00], no such algorithms are known to have complexity polynomial in $\mathrm{size}_p(f)$ alone. Our main theorem below shows that the existence of such an algorithm would imply a complexity collapse nearly as strong as $\mathbf{P}=\mathbf{NP}$. Nevertheless, we obtain new sub-cases of $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathbb{Z}[x]\times\mathbb{P})$ lying in $\mathbf{P}$.

Theorem 1.5.

1. $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathcal{F}_{1,m}\times\mathbb{P})\in\mathbf{P}$ *for $m\in\{0,1,2\}$.*
2. *For any $(f,p)\in\mathbb{Z}[x]\times\mathbb{P}$ such that f has* **no** *p-adic lower m-nomials for $m\geq 3$, and p does* **not** *divide a_i-a_j for any lower binomial with exponents $\{a_i,a_j\}$, we can decide the existence of a root in $\mathbb{Q}_p$ for f in time polynomial in $\mathrm{size}_p(f)$.*
3. *There is a countable union of algebraic hypersurfaces $\mathcal{E}\subsetneq\mathbb{Z}[x]\times\mathbb{P}$, with natural density 0, such that $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}((\mathbb{Z}[x]\times\mathbb{P})\setminus\mathcal{E})\in\mathbf{NP}$. Furthermore, we can decide in $\mathbf{P}$ whether an $f\in\mathcal{F}_{1,3}$ lies in $\mathcal{E}$.*
4. *If $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathbb{Z}[x]\times\mathbb{P})\in\mathbf{ZPP}$ then $\mathbf{NP}\subseteq\mathbf{ZPP}$.*
5. *If the Wagstaff Conjecture is true, then $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathbb{Z}[x]\times\mathbb{P})\in\mathbf{P}\implies\mathbf{P}=\mathbf{NP}$, i.e., we can strengthen Assertion (4) above.*

Remark 1.6. *The* **Wagstaff Conjecture**, *dating back to 1979 (see, e.g., [BS96, Conj. 8.5.10, pg. 224]), is the assertion that the least prime congruent to k mod N is $O(\varphi(N)\log^2 N)$, where $\varphi(N)$ is the number of integers in $\{1,\ldots,N\}$ relatively prime to N. Such a bound is significantly stronger than the known implications of the* **Generalized Riemann Hypothesis (GRH)**. ⋄

While the real analogue of Assertion (1) is easy to prove, $\mathtt{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,3})\in\mathbf{P}$ was proved only recently [BRS09, Thm. 1.3]. That $\mathtt{FEAS}_{\mathbb{Q}_p}(\mathcal{F}_{1,3})\in\mathbf{NP}$ for any prime p is surprisingly subtle to prove, having been accomplished by the authors just as this paper went to press [AIRR10].

The intuition behind our algorithmic speed-ups (Assertions (1)–(3)) is that any potential hardness is caused by numerical ill-conditioning, quite similar to the sense long known in numerical linear algebra. Indeed, 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 $\mathbb{Q}_p$. This lies behind the hypotheses of Assertions (2) and (3) (see also Theorem 1.11 below). Note that the hypothesis of Assertion (2) is rather stringent: if one fixes $f\in\mathcal{F}_{1,m}$ with $m\geq 3$ and varies p, then it is easily checked that $\mathrm{Newt}_p(f)$ is a line segment (so the hypothesis fails) for all but finitely many p. On the other hand, the hypothesis for Assertion (3) holds for a significantly large fraction of inputs (see also Proposition 2.13 of Section 2.4).

Example 1.7. *Let T denote the family of pairs $(f,p)\in\mathbb{Z}[x]\times\mathbb{P}$ with $f(x)=a+bx^{11}+cx^{17}+x^{31}$ and let $T^*:=T\setminus\mathcal{E}$. Then there is a sparse 61×61 structured matrix $\mathcal{S}$ (cf. Lemma 2.8 in Section 2.3 below) such that $(f,p)\in T^*\iff p\not|\det\mathcal{S}$. So by Theorem 1.5, $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(T^*)\in$* **NP**, *and Proposition 2.13 in Section 3 below 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*

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

for all but 352 primes p. ⋄

One subtlety behind Assertion (3) is that $\mathbb{Q}_p$ is uncountable and thus, unlike $\mathtt{FEAS}_{\mathbb{F}_p}$, $\mathtt{FEAS}_{\mathbb{Q}_p}$ does **not** admit an obvious succinct certificate. Indeed, the best previous complexity bound relative to the sparse input size appears to have been $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathbb{Z}[x]\times\mathbb{P})\in\mathbf{EXPTIME}$ [MW99].[3] In particular, $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathcal{F}_{1,4}\times\mathbb{P})\overset{?}{\in}\mathbf{NP}$ and $\mathtt{FEAS}_{\mathbb{R}}(\mathcal{F}_{1,4})\overset{?}{\in}\mathbf{NP}$ are still open questions [BRS09, Sec. 1.2]. A real analogue for Assertion (3) is also unknown at this time.

As for lower bounds, while it is not hard to show that the full problem $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}$ is $\mathbf{NP}$-hard, the least n making $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathbb{Z}[x_1,\ldots,x_n]\times\mathbb{P})$ $\mathbf{NP}$-hard appears not to have been known unconditionally. In particular, a weaker version of Assertion (4) was found recently, but only under the truth of an unproved hypothesis on the distribution of primes in arithmetic progression [Roj07a, Main Thm.]. Assertion (4) thus also provides an interesting contrast to earlier work of

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

[3] An earlier result claiming $\mathtt{FEAS}_{\mathbb{Q}_{\mathrm{primes}}}(\mathbb{Z}[x]\times\mathbb{P})\in\mathbf{NP}$ for "most" inputs was announced without proof in [Roj07a, Main Thm.] (see Proposition 1 there).

H. W. Lenstra, Jr. [Len99a], who showed that one can actually find all **low** degree factors of a sparse polynomial (over $\mathbb{Q}[x]$ as opposed to $\mathbb{Q}_p[x]$) in polynomial time. Real analogues to Assertions (4) and (5) are unknown.

1.2 Primes in Random Arithmetic Progressions and a Tropical Trick

The key to proving our lower bound results (Assertions (4) and (5) of Theorem 1.5) is an efficient reduction from a problem discovered to be **NP**-hard by David Alan Plaisted: deciding whether a sparse univariate polynomial vanishes at a complex $D^{\underline{\text{th}}}$ root of unity [Pla84]. Reducing from this problem to its analogue over $\mathbb{Q}_p$ is straightforward, provided $\mathbb{Q}^*_p := \mathbb{Q}_p \setminus \{0\}$ contains a cyclic subgroup of order D where D has sufficiently many distinct prime divisors. We thus need to consider the factorization of $p-1$, which in turn leads us to primes congruent to 1 modulo certain integers.

While efficiently constructing random primes in **arbitrary** arithmetic progressions remains a famous open problem, we can now at least efficiently build random primes p such that p is moderately sized but $p-1$ has many prime factors. We use the notation $[j] := \{1, \ldots, j\}$ for any $j \in \mathbb{N}$.

Theorem 1.8. *For any $\delta > 0$, a failure probability $\varepsilon \in (0, 1/2)$, 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 a positive integer c 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.*

Theorem 1.8 and its proof are inspired in large part by an algorithm of von zur Gathen, Karpinski, and Shparlinski [vzGKS96, Algorithm following Fact 4.9]. (Theorem 4.10 of [vzGKS96] does not imply Theorem 1.8 above, nor vice-versa.) In particular, they use an intricate random sampling technique to prove that the enumerative analogue of $\mathtt{FEAS}_{\mathbb{F}_{\substack{\text{prime}\\\text{powers}}}}(\mathbb{Z}[x_1, x_2] \times \mathbb{P})$ is #**P**-hard [vzGKS96, Thm. 4.11].

Our harder upper bound results (Assertions (2) and (3) of Theorem 1.5) will follow in large part from an arithmetic analogue of a key idea from tropical geometry: **toric deformation**. Toric deformation, roughly speaking, means cleverly embedding an algebraic set into a family of algebraic sets 1 dimension higher, in order to invoke combinatorial methods (see, e.g., [EKL06]). Here, this simply means that we find ways to reduce problems involving general $f \in \mathbb{Z}[x]$ to similar problems involving binomials.

Lemma 1.9. *(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* $\mathrm{Newt}_p(f)$ *with inner normal $(v, 1)$.* ■

Example 1.10. *In Example 1.4 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 1.9 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. ◇

Theorem 1.11 *[AI10, Thm. 4.5] Suppose $(f, p) \in \mathbb{Z}[x] \times \mathbb{P}$, $(v, 1)$ is an inner normal to a lower edge E of* $\mathrm{Newt}_p(f)$, *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 $\mathrm{ord}_p\zeta = v$ is exactly the number of roots of g in $\mathbb{Q}_p$.* ■

Our main results are proved in Section 3, after the development of some additional theory below.

2. BACKGROUND

Our lower bounds will follow from a chain of reductions involving some basic problems we will review momentarily. We then show how to efficiently construct random primes p such that $p-1$ has many prime factors in Section 2.2, and then conclude with some quantitative results on resultants in Sections 2.3 and 2.4.

2.1 Roots of Unity and NP-Completeness

Recall that any Boolean expression of one of the following forms:

$(\diamondsuit)$ $y_i \vee y_j \vee y_k$, $\neg y_i \vee y_j \vee y_k$, $\neg y_i \vee \neg y_j \vee y_k$, $\neg y_i \vee \neg y_j \vee \neg y_k$,
with $i, j, k \in [3n]$,

is a **`3CNFSAT` clause**. A **satisfying assigment** for an arbitrary Boolean formula $B(y_1, \ldots, y_n)$ is an assigment of values from $\{0, 1\}$ to the variables $y_1, \ldots, y_n$ which makes the equality $B(y_1, \ldots, y_n) = 1$ 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 2.1. *Letting $P := (p_1, \ldots, p_n)$ denote any strictly increasing sequence of primes, let us inductively define a semigroup homomorphism $\mathcal{P}_P$ — the* **Plaisted morphism with respect to** P *— from certain Boolean expressions in the variables $y_1, \ldots, y_n$ to $\mathbb{Z}[x]$, as follows:*[4] *(0) $D_P := \prod_{i=1}^n p_i$, (1) $\mathcal{P}_P(0) := 1$, (2) $\mathcal{P}_P(y_i) := x^{D_P/p_i} - 1$, (3) $\mathcal{P}_P(\neg B) := (x^{D_P} - 1)/\mathcal{P}_P(B)$, for any Boolean expression B for which $\mathcal{P}_P(B)$ has already been defined, (4) $\mathcal{P}_P(B_1 \vee B_2) := \mathrm{lcm}(\mathcal{P}_P(B_1), \mathcal{P}_P(B_2))$, for any Boolean expressions B_1 and B_2 for which $\mathcal{P}_P(B_1)$ and $\mathcal{P}_P(B_2)$ have already been defined.* ◇

Lemma 2.2. *[Pla84, Sec. 3, pp. 127–129] Suppose $P = (p_i)_{k=1}^n$ 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 $(\diamondsuit)$, we have* $\mathrm{size}(\mathcal{P}_P(C))$ *polynomial in n^γ. In particular, $\mathcal{P}_P$ can be evaluated at any such C in time polynomial in n. Furthermore, if K is any field possessing D_P distinct ${D_P}^{\underline{th}}$ roots of unity, then a* `3CNFSAT` *instance $B(y) := C_1(y) \wedge \cdots \wedge C_k(y)$ has a satisfying assignment iff the univariate polynomial system $F_B := (\mathcal{P}_P(C_1), \ldots, \mathcal{P}_P(C_k))$ has a root $\zeta \in K$ satisfying $\zeta^{D_P} - 1$.* ■

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.

2.2 Randomization to Avoid Riemann Hypotheses

The result below allows us to prove Theorem 1.8 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 $1 \bmod M$.

[4] Throughout this paper, for Boolean expressions, we will always identify 0 with "`False`" and 1 with "`True`".

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 $a \nmid M$ for all $a \in \mathcal{D}(x)$ then $\pi(x; M, 1) \geq \frac{\pi(x)}{2\varphi(M)}$.* ■

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 2.3.

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 $P = (p_j)_{j=1}^{n}$ 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 $p_1, \ldots$, 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.`” ◇

Remark 2.4. *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). ◇

Remark 2.5. *An anonymous referee suggested that one can employ a faster probabilistic primality test in Step 3 (e.g, [Mor07]), reserving the AKS algorithm solely for so-called* **pseudoprimes**. *This can likely reduce the complexity bound from Theorem 1.8 slightly.* ◇

Proof of Theorem 1.8: It clearly suffices to prove that Algorithm 2.3 is correct, has a success probability that is at least $1 - \varepsilon$, and works within

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

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

Proving Correctness and the Success Probability Bound for Algorithm 2.3: First observe that $M_1, \ldots, M_L$ 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 $i \in [\lceil 2/\varepsilon \rceil \ell]$ is uniformly random), the arithmetic progression $\{1 + M_i, \ldots, 1 + KM_i\}$ 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 $\{1 + M_i, \ldots, 1 + KM_i\}$ 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 $\left(1 - \frac{1}{t}\right)^{ct} \leq e^{-c}$ (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 2.3 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 2.3): Let $L' := nL$ and, for the remainder of our proof, let p_i denote the $i^{\underline{\text{th}}}$ 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

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

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

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

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

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

randomized bit operations. ■

2.3 Transferring from Complex Numbers to p-adics

The proposition below is a folkloric way to reduce systems of univariate polynomial equations to a single polynomial equation, and was already used by Plaisted at the beginning of his proof of Theorem 5.1 in [Pla84].

Proposition 2.6. *Given any $f_1,\ldots,f_k\in\mathbb{Z}[x]$ with maximum coefficient absolute value H, let $d:=\max_i \deg f_i$ and*

$$\tilde{f}(x):=x^d(f_1(x)f_1(1/x)+\cdots+f_k(x)f_k(1/x)).$$

Then $f_1=\cdots=f_k=0$ has a root on the complex unit circle iff $\tilde{f}$ has a root on the complex unit circle.

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

By introducing the classical univariate resultant we will be able to derive the explicit quantitative bounds we need.

Definition 2.7. *(See, e.g., [GKZ94, Ch. 12, Sec. 1, pp. 397–402].) Suppose $f(x)=a_0+\cdots+a_dx^d$ and $g(x)=b_0+\cdots+b_{d'}x^{d'}$ are polynomials with indeterminate coefficients. We define their* **Sylvester matrix** *to be the $(d+d')\times(d+d')$ matrix*

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

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

Lemma 2.8. *Following the notation of Definition 2.7, assume $f,g\in K[x]$ 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 iff $\mathcal{R}_{(d,d')}(f,g)=0$. More generally, we have $\mathcal{R}_{(d,d')}(f,g)=a_d^{d'}\prod\limits_{f(\zeta)=0}g(\zeta)$ 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 $|\mathcal{R}_{(d,d')}(f,g)|\leq m^{d'/2}m'^{d/2}H^{d+d'}$.* ■

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

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

Lemma 2.9. *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 $D^{\underline{th}}$ root of unity $\Longleftrightarrow$ f vanishes at a $D^{\underline{th}}$ root of unity in $\mathbb{Q}_p$.*

Remark 2.10 *Note that x^2+x+1 vanishes at a $3^{\underline{rd}}$ 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.* ⋄

Proof of Lemma 2.9: First note that by our assumption on p, $\mathbb{Q}_p$ has D distinct $D^{\underline{\text{th}}}$ roots of unity: This follows easily from Hensel's Lemma (see, e.g., [Rob00, Pg. 48]) and $\mathbb{F}_p$ having D distinct $D^{\underline{\text{th}}}$ roots of unity. Since $\mathbb{Z}\hookrightarrow\mathbb{Q}_p$ and $\mathbb{Q}_p$ contains all $D^{\underline{\text{th}}}$ roots of unity by construction, the equivalence then follows directly from Lemma 2.8. ■

2.4 Good Inputs and Bad Trinomials

Definition 2.11. *For any field K, write any $f\in K[x]$ as $f(x)=\sum_{i=1}^m c_ix^{a_i}$ with $0\leq a_1<\cdots<a_m$. Letting $\mathcal{A}=\{a_1,\ldots,a_m\}$, and following the notation of Lemma 2.9, we then define the* **$\mathcal{A}$-discriminant** *of f, $\Delta_{\mathcal{A}}(f)$, to be*

$$\mathcal{R}_{(\bar{a}_m,\bar{a}_m-\bar{a}_2)}\left(\bar{f},\,\frac{\partial\bar{f}}{\partial x}\Big/x^{\bar{a}_2-1}\right)\Big/c_m^{\bar{a}_m-\bar{a}_{m-1}},$$

where $\bar{a}_i:=(a_i-a_1)/g$ for all i, $\bar{f}(x):=\sum_{i=1}^m c_ix^{\bar{a}_i}$, and $g:=\gcd(a_2-a_1,\ldots,a_m-a_1)$ (see also [GKZ94, Ch. 12, pp. 403–408]). Finally, if $c_i\neq 0$ for all i, then we call $\mathrm{Supp}(f):=\{a_1,\ldots,a_m\}$ the **support** *of f.* ⋄

Remark 2.12 *Note that when $\mathcal{A}=\{0,\ldots,d\}$ we have $\Delta_{\mathcal{A}}(f)=\mathcal{R}_{(d,d-1)}(f,f')/c_d$, i.e., for dense polynomials, the $\mathcal{A}$-discriminant agrees with the classical discriminant* ⋄

Let us now clarify our statement about natural density 0 from Assertion (4) of Theorem 1.5: First, 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 $\mathbb{Z}[x]$ 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.,
$a+bx^{99}+x^{2001}\mapsto((a,0),(b,99),(1,2001),(0,0),(0,0),\ldots)$.
Then natural density for a set of pairs $\mathcal{I}\subseteq\mathbb{Z}[x]\times\mathbb{P}$ simply means the corresponding natural density within $(\mathbb{Z}\times(\mathbb{N}\cup\{0\}))^\infty\times\mathbb{P}$. In particular, our claim of natural density 0 can be made explicit as follows.

Proposition 2.13. *For any subset $\mathcal{A}=\{a_1,\ldots,a_m\}\subset\mathbb{N}\cup\{0\}$ with $0=a_1<\cdots<a_m$, let $T_{\mathcal{A}}$ denote the family of pairs $(f,p)\in\mathbb{Z}[x]\times\mathbb{P}$ with $f(x)=\sum_{i=1}^m c_ix^{a_i}$ and let $T^*_{\mathcal{A}}$ denote the subset of $T_{\mathcal{A}}$ consisting of those pairs (f,p) with $p\not|\Delta_{\mathcal{A}}(f)$. 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 $d:=a_m/\gcd(a_2,\ldots,a_m)$. Then for all $H\geq 17$ we have*

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

Note that each $T^*_{\mathcal{A}}(H)$ is the complement of a union of hypersurfaces (one for each mod p reduction of $\Delta_{\mathcal{A}}(f)$) in a "brick" in $\mathbb{Z}^m\times\mathbb{P}$. We will see in the proof of Assertion (3) of Theorem 1.5 that the exceptional set $\mathcal{E}$ is then merely the complement of the union $\bigcup_{\mathcal{A}}\mathcal{T}^*_{\mathcal{A}}$ as $\mathcal{A}$ ranges over all finite subsets of $\mathbb{N}\cup\{0\}$. Our proposition above is proved in Section 3.2.

Before proving our main results, let us make some final observations about the roots of trinomials.

Corollary 2.14. *Suppose $f(x)=c_1+c_2x^{a_2}+c_3x^{a_3}\in\mathcal{F}_{1,3}$, $\mathcal{A}:=\{0,a_2,a_3\}$, $0<a_2<a_3$, $a_3\geq 3$, and $\gcd(a_2,a_3)=1$. Then: (0) $\Delta_{\mathcal{A}}(f)=(a_3-a_2)^{a_3-a_2}a_2^{a_2}c_2^{a_3}-(-a_3)^{a_3}c_1^{a_3-a_2}c_3^{a_2}$. (1) $\Delta_{\mathcal{A}}(f)\neq 0\Longleftrightarrow f$ has no degenerate roots. In which*

case, we also have $\Delta_{\mathcal{A}}(f) = \frac{(-1)^{a_3}c_3^{a_2-1}}{c_1^{a_2-1}}\prod_{f(\zeta)=0} f'(\zeta)$.

(2) *Deciding whether f has a degenerate root in $\mathbb{C}_p$ can be done in time polynomial in* $\text{size}_p(f)$.

Proof:
Assertion (0): [GKZ94, Prop. 1.8, pg. 274]. ■
Assertion (1): The first assertion follows directly from Definition 2.11 and the vanishing criterion for $\text{Res}_{(a_3,a_3-a_2)}$ from Lemma 2.8. To prove the second assertion, observe that the product formula from Lemma 2.8 implies that

$$\Delta_{\mathcal{A}}(f) = c_3^{a_3-a_2}\left(\prod_{f(\zeta)=0}\frac{f'(\zeta)}{\zeta^{a_2-1}}\right)\Big/\, c_3^{a_3-a_2}$$
$$= (-1)^{a_3}\left(\prod_{f(\zeta)=0} f'(\zeta)\right)\Big/\,(c_1/c_3)^{a_2-1}. \blacksquare$$

Assertion (2): From Assertion (1) it suffices to detect the vanishing of $\Delta_{\mathcal{A}}(f)$. However, while Assertion (0) implies that one can evaluate $\Delta_{\mathcal{A}}(f)$ with a small number of arithmetic operations, the bit-size of $\Delta_{\mathcal{A}}(f)$ can be quite large. Nevertheless, we can decide within time polynomial in size(f) whether these particular $\Delta_{\mathcal{A}}(f)$ vanish for integer c_i via **gcd-free bases** (see, e.g., [BRS09, Sec. 2.4]). ■

We will also need a concept that is essentially the opposite of a degenerate root: Given any $f \in \mathbb{Z}[x]$, we call a $\zeta_0 \in \mathbb{Z}/p^\ell\mathbb{Z}$ an **approximate root** iff $f(\zeta_0) = 0$ mod p^ℓ and $\text{ord}_p f'(\zeta_0) < \ell/2$, i.e., ζ_0 satisfies the hypotheses of Hensel's Lemma (see, e.g., [Rob00, Pg. 48]), and thus ζ_0 can be lifted to a ***p*-adic integral** root ζ of f. The terminology "approximate root" is meant to be reminiscent of an Archimedean analogue guaranteeing that $\zeta_0 \in \mathbb{C}$ converge quadratically to a true (non-degenerate) complex root of f (see, e.g., [Sma86]).

We call any $\text{Newt}_p(f)$ such that f has no lower m-nomials with $m \geq 3$ **generic**. 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**.

3. PROVING OUR MAIN RESULTS

3.1 The Proof of Theorem 1.5

Assertion (1) ($\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathcal{F}_{1,m} \times \mathbb{P}) \in \mathbf{P}$ for $m \leq 2$):
First note that the case $m \leq 1$ is trivial: such a univariate m-nomial has no roots in $\mathbb{Q}_p$ iff it is a nonzero constant. So let us now assume $m = 2$.

We can easily reduce to the special case $f(x) := x^d - \alpha$ with $\alpha \in \mathbb{Q}^*$, since we can divide any input by a suitable monomial term, and arithmetic over $\mathbb{Q}$ is doable in polynomial time. Clearly then, any p-adic root ζ of $x^d - \alpha$ satisfies $d\text{ord}_p\zeta = \text{ord}_p\alpha$. Since we can compute $\text{ord}_p\alpha$ and reductions of integers mod d in polynomial-time [BS96, Ch. 5], we can then assume that $d|\text{ord}_p\alpha$ (for otherwise, f would have no roots over $\mathbb{Q}_p$). Replacing $f(x)$ by $p^{-\text{ord}_p\alpha}f(p^{\text{ord}_p\alpha/d}x)$, we can assume further that $\text{ord}_p\alpha = \text{ord}_p\zeta = 0$. In particular, if $\text{ord}_p\alpha$ was initially a nonzero multiple of d, then $\log\alpha \geq d\log_2 p$. So size$(f) \geq d$ and our rescaling at worst doubles size(f).

Letting $k := \text{ord}_p d$, note that $f'(x) = dx^{d-1}$ and thus $\text{ord}_p f'(\zeta) = \text{ord}_p(d) + (d-1)\text{ord}_p\zeta = k$. So by Hensel's Lemma it suffices to decide whether the mod p^ℓ reduction of f has a root in $(\mathbb{Z}/p^\ell\mathbb{Z})^*$, for $\ell = 1 + 2k$. Note in particular that $\text{size}(p^\ell) = O(\log(p)\text{ord}_p d) = O(\log(p)\log(d)/\log p) = O(\log d)$ which is linear in our notion of input size. Since the equation $x^d = \alpha$ can be solved in any cyclic group via a fast exponentiation, we can then clearly decide whether $x^d - \alpha$ has a root in $(\mathbb{Z}/p^\ell\mathbb{Z})^*$ within $\mathbf{P}$, provided $p^\ell \notin \{8, 16, 32, \ldots\}$. This is because of the classical structure theorem for the multiplicative group of $\mathbb{Z}/p^\ell\mathbb{Z}$ (see, e.g., [BS96, Thm. 5.7.2 & Thm. 5.6.2, pg. 109]).

To dispose of the remaining cases $p^\ell \in \{8, 16, 32, \ldots\}$, first recall that the multiplicative group of $\mathbb{Z}/2^\ell$ is exactly

$$\left\{\pm 1, \pm 5, \pm 5^2, \pm 5^3, \ldots, \pm 5^{2^{\ell-2}-1} \text{ mod } 2^\ell\right\}$$

(see, e.g., [BS96, Thm. 5.7.2 & Thm. 5.6.2, pg. 109]). So we can replace d by its reduction mod $2^{\ell-2}$, since every element of $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ has order dividing $2^{\ell-2}$, and this reduction can certainly be computed in polynomial-time. Let us then write $d = 2^h d'$ where $2 \not| d'$ and $h \in \{0, \ldots, \ell - 3\}$, and compute $d'' := 1/d'$ mod $2^{\ell-2}$. Clearly then, $x^d - \alpha$ has a root in $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ iff $x^{2^h} - \alpha'$ has a root in $(\mathbb{Z}/2^\ell\mathbb{Z})^*$, where $\alpha' := \alpha^{d''}$ (since exponentiation by any odd power is an automorphism of $(\mathbb{Z}/2^\ell\mathbb{Z})^*$). Note also that α', d', and d'' can be computed in polynomial time via recursive squaring and standard modular arithmetic, and $h \leq \log_2 d$.

Since $x^{2^h} - \alpha'$ always has a root in $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ when $h = 0$, we can then restrict our root search to the cyclic subgroup $\left\{1, 5^2, 5^4, 5^6, \ldots, 5^{2^{\ell-2}-2}\right\}$ when $h \geq 1$ and α' is a square (since there can be no roots when $h \geq 1$ and α' is not a square). Furthermore, we see that $x^{2^h} - \alpha'$ can have no roots in $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ if $\text{ord}_2\alpha'$ is odd. So, by rescaling x, we can assume further that $\text{ord}_2\alpha' = 0$, and thus that α' is odd. Now an odd α' is a square in $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ iff $\alpha' \equiv 1$ mod 8 [BS96, Ex. 38, pg. 192], and this can clearly be checked in $\mathbf{P}$. So we can at last decide the existence of a root in $\mathbb{Q}_2$ for $x^d - \alpha$ in $\mathbf{P}$: Simply use fast exponentiation to solve the equation $x^{2^h} = \alpha'$ over the cyclic subgroup $\left\{1, 5^2, 5^4, 5^6, \ldots, 5^{2^{\ell-2}-2}\right\}$ of $(\mathbb{Z}/2^\ell\mathbb{Z})^*$ [BS96, Thm. 5.7.2 & Thm. 5.6.2, pg. 109]. ■

Assertion (2) ($\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathbb{P}) \in \mathbf{P}$ for generic, unramified $\text{Newt}_p(f)$):
Assertion (2) follows directly from Theorem 1.11, since we can apply the $m = 2$ case of Assertion (1) to the resulting lower binomials. In particular, note that the number of lower binomials of f is no more than the number of monomial terms of f, which is in turn bounded above by size(f), so the complexity is indeed $\mathbf{P}$. ■

Assertion (3) ($\mathbf{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathbb{P}) \in \mathbf{NP}$ usually):
Let us first observe that it suffices to prove that, for most inputs, we can detect roots in $\mathbb{Z}_p$ in $\mathbf{NP}$. This is because $x \in \mathbb{Q}_p \setminus \mathbb{Z}_p \iff \frac{1}{x} \in p\mathbb{Z}_p$, so letting $f^*(x) := x^{\deg f}f(1/x)$ denote the **reciprocal polynomial** of f, the set of p-adic rational roots of f is simply the union of the p-adic integer roots of f and the reciprocals of the p-adic integer roots of f^*. We may also assume that f is not divisible by x.

Note also that we can find the p-parts of the c_i in polynomial-time via gcd-free bases [BRS09, Sec. 2.4] and thus compute $\text{Newt}_p(f)$ in time polynomial in $\text{size}_p(f)$ (via standard convex hull algorithms, e.g., [Ede87]). Since $\text{ord}_p c_i \leq \log_p c_i \leq \text{size}(c_i)$, note also that that every root $\zeta \in \mathbb{C}_p$ of f satisfies $|\text{ord}_p\zeta| \leq 2\max_i \text{size}(c_i) \leq 2\text{size}(f) < 2\text{size}_p(f)$.

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 $\mathbb{Z}_p$. Letting $g(x) := f'(x)/x^{a_1-1}$, and $\zeta \in \mathbb{Z}_p$ be any p-adic integer root of f, note then that

$$(\star) \qquad \text{ord}_p f'(\zeta) = (a_1 - 1)\text{ord}_p(\zeta) + \text{ord}_p g(\zeta).$$

Note also that $\Delta_{\mathcal{A}}(f) = \text{Res}_{a_m, a_m - a_1}(f, g)$ so if $p \not| \Delta_{\mathcal{A}}(f)$ then f and g have no common roots in the algebraic closure of $\mathbb{F}_p$, by Lemma 2.8. In particular, $p \nmid \Delta_{\mathcal{A}}(f) \Longrightarrow g(\zeta) \not\equiv 0 \mod p$; and thus $p \nmid \Delta_{\mathcal{A}}(f,g) \Longrightarrow \text{ord}_p f'(\zeta) = (a_1-1)\text{ord}_p(\zeta)$. 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 $p \nmid \Delta_{\mathcal{A}}(f) \Longrightarrow \text{ord}_p f'(\zeta) \leq 2\text{size}(f)$, thanks to $(\star)$.

Our fraction of inputs admitting a succinct certificate will then correspond precisely to those (f,p) such that $p \nmid \Delta_{\mathcal{A}}(f)$. In particular, let us define $\mathcal{E}$ to be the union of all pairs (f,p) such that $p|\Delta_{\mathcal{A}}(f)$, 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.

Now fix $\ell = 4\text{size}(f)+1$. Clearly then, by Hensel's Lemma, for any $(f,p) \in (\mathbb{Z}[x] \times \mathbb{P}) \setminus \mathcal{E}$, f has a root $\zeta \in \mathbb{Z}_p \Longleftrightarrow f$ has a root $\zeta_0 \in \mathbb{Z}/p^\ell\mathbb{Z}$. Since $\log(p^\ell) = O(\text{size}(f)\log p) = O\big(\text{size}_p(f)^2\big)$, 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: an approximate root $\zeta_0 \in (\mathbb{Z}/p^\ell\mathbb{Z})^*$ of f with $\ell = 4\text{size}(f)+1$. ■

Assertion (4) ($\texttt{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathbb{P})$ is NP-hard under ZPP-reductions):
We will prove a (**ZPP**) randomized polynomial-time reduction from $\texttt{3CNFSAT}$ to $\texttt{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathbb{P})$, making use of the intermediate input families $\{(\mathbb{Z}[x])^k \mid k \in \mathbb{N}\} \times \mathbb{P}$ and $\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\} \times \mathbb{P}$ along the way.

Toward this end, suppose $B(y) := C_1(y) \wedge \cdots \wedge C_k(y)$ is any $\texttt{3CNFSAT}$ instance. The polynomial system $(\mathcal{P}_P(C_1), \ldots, \mathcal{P}_P(C_k))$, for P the first n primes (employing Lemma 2.2), then clearly yields $\texttt{FEAS}_{\mathbb{C}}(\{(\mathbb{Z}[x])^k \mid k \in \mathbb{N}\}) \in \mathbf{P} \Longrightarrow \mathbf{P} = \mathbf{NP}$. Composing this reduction with Proposition 2.6, we then immediately obtain $\texttt{FEAS}_{\mathbb{C}}(\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\}) \in \mathbf{P} \Longrightarrow \mathbf{P} = \mathbf{NP}$.

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 P: by applying Theorem 1.8 with $\varepsilon = 1/3$ (cf. Lemma 2.9) we immediately obtain the implication $\texttt{FEAS}_{\mathbb{Q}_{\text{primes}}}((\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\}) \times \mathbb{P}) \in \mathbf{ZPP} \Longrightarrow \mathbf{NP} \subseteq \mathbf{ZPP}$.

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 $\texttt{FEAS}_{\mathbb{Q}_{\text{primes}}}((\mathbb{Z}[x] \times \{x^D - 1 \mid D \in \mathbb{N}\}) \times \mathbb{P})$ to $\texttt{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathbb{P})$: simply map any instance $(f(x), x^D - 1, p)$ of the former problem to $(f(x)^2 - (x^D - 1)^2 p, p)$. So we are done. ■

Assertion (5) ($\texttt{FEAS}_{\mathbb{Q}_{\text{primes}}}(\mathbb{Z}[x] \times \mathbb{P})$ is NP-hard, assuming Wagstaff's Conjecture):
If we also have the truth of the Wagstaff Conjecture then 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 $k = O\Big(\frac{\varphi(D)}{D} \log^2 D\Big)$. (Note that $1 \leq \frac{\varphi(D)}{D} < D$ 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. ■

3.2 The Proof of Proposition 2.13

By the Schwartz-Zippel Lemma [Sch80], $\Delta_{\mathcal{A}}(f)$ vanishes for at most $(2d-1)m(2H+1)^{m-1}$ selections of coefficients from $\{-H, \ldots, H\}$. In other words, $\Delta_{\mathcal{A}}(f) = 0$ for a fraction of at most $\frac{(2d-1)m}{2H+1}$ of the pairs $(f,p) \in T_{\mathcal{A}}(H)$.

Clearly, a pair $(f,p) \in T_{\mathcal{A}}(H)$ for which $p \nmid \Delta_{\mathcal{A}}(f)$ must satisfy $\Delta_{\mathcal{A}}(f) \neq 0$. We have just shown that the fraction of $T_{\mathcal{A}}(H)$ satisfying the last condition is at least $1 - \frac{(2d-1)m}{2H+1}$. Once we show that, amongst these pairs, at least

$$1 - \frac{1+(2d-1)\log(mH)}{H/\log H}$$

of them actually satisfy $p \nmid \Delta_{\mathcal{A}}(f)$, then we will be done.

To prove the last lower bound, note that $\Delta_{\mathcal{A}}(f)$ has degree at most $2d-1$ in the coefficients of f by Lemma 2.8. Also, for any fixed $f \in T_{\mathcal{A}}(H)$, $\Delta_{\mathcal{A}}(f)$ is an integer as well, and is thus divisible by no more than $1 + (2d-1)\log(mH))$ primes if $\Delta_{\mathcal{A}}(f) \neq 0$. (This follows from Lemma 2.8 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 $\Delta_{\mathcal{A}}(f)$ is bounded above by $\frac{1+(2d-1)\log(mH)}{H/\log H}$. ■

Acknowledgements

The authors would like to thank David Alan Plaisted for his kind encouragement, and Eric Bach, Sidney W. Graham, and Igor Shparlinski for many helpful comments on primes in arithmetic progression. We also thank Matt Papanikolas for valuable p-adic discussions. Finally, we thank the anonymous referees for insightful comments and corrections.

4. 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.

[AI10] Avendaño, Martín and Ibrahim, Ashraf, *"Ultrametric Root Counting,"* submitted for publication, also available as Math ArXiV preprint `0901.3393v3` .

[AIRR10] Avendaño, Martín; Ibrahim, Ashraf; Rojas, J. Maurice; Rusek, Korben, *"Succinct Certificates and Maximal Root Counts for p-adic Trinomials and Beyond,"* in progress.

[BS96] Bach, Eric and Shallit, Jeff, *Algorithmic Number Theory, Vol. I: Efficient Algorithms,* MIT Press, Cambridge, MA, 1996.

[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.

[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.

[CDV06] Castrick, Wouter; Denef, Jan; and Vercauteren, Frederik, *"Computing Zeta Functions of Nondegenerate Curves,"* International Mathematics Research Papers, vol. 2006, article ID 72017, 2006.

[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.

[DvdD88] Denef, Jan and van den Dries, Lou, *"p-adic and Real Subanalytic Sets,"* Annals of Mathematics (2) **128** (1988), no. 1, pp. 79–138.

[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 Denef, Leonard Lipshitz, Thanases Pheidas and Jan Van Geel. Contemporary Mathematics, 270, American Mathematical Society, Providence, RI, 2000.

[Ede87] Edelsbrunner, Herbert, *Algorithms in combinatorial geometry,* EATCS Monographs on Theoretical Computer Science, 10, Springer-Verlag, Berlin, 1987.

[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.

[EKL06] Einsiedler, Manfred; Kapranov, Misha M.; Lind, Doug, *"Non-archimedean amoebas and tropical varieties,"* J. reine und angew. Math. 601 (2006), pp. 139–158.

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

[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.

[Len99a] Lenstra (Jr.), Hendrik W., *"Finding Small Degree Factors of Lacunary Polynomials,"* Number Theory in Progress, Vol. 1 (Zakopane-Kóscielisko, 1997), pp. 267–276, de Gruyter, Berlin, 1999.

[Len99b] ______________________, *"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, Dartmouth University, 2005.

[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, 2$^{\underline{\text{nd}}}$ ed., (edited by B. Buchberger, G. E. Collins, and R. Loos, in cooperation with R. Albrecht), Springer-Verlag 1982.

[Mor07] Morain, Francois, *"Implementing the asymptotically fast version of the elliptic curve primality proving algorithm,"* Math. Comp. 76 (2007), pp. 493–505.

[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.

[Poo01a] 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.

[Roj04] _______________, *"Arithmetic Multivariate Descartes' Rule,"* American Journal of Mathematics, vol. 126, no. 1, February 2004, pp. 1–30.

[Roj07a] _______________, *"On Interpolating Between Quantum and Classical Complexity Classes,"* Proceedings of Mathematics of Quantum Computation and Quantum Technology (November 13-16, 2005, Texas A&M University), pp. 67–88, Taylor & Francis, 2007.

[Roj07b] _______________, *"Efficiently Detecting Torsion Points and Subtori,"* proceedings of MAGIC 2005 (Midwest Algebra, Geometry, and their Interactions Conference, Oct. 7–11, 2005, Notre Dame University, Indiana), edited by A. Corso, J. Migliore, and C. Polini), pp. 213–233, Contemporary Mathematics, vol. 448, AMS Press, 2007.

[Sch80] Schwartz, Jacob T., *"Fast Probabilistic Algorithms for Verification of Polynomial Identities,"* J. of the ACM 27, 701–717, 1980.

[Sma86] Smale, Steve, *"Newton's Method Estimates from Data at One Point,"* The Merging of Disciplines: New Directions in Pure, Applied, and Computational Mathematics (Laramie, Wyo., 1985), pp. 185–196, Springer, New York, 1986.