

A Complexity Chasm for Solving Univariate Sparse Polynomial Equations Over p -adic Fields

J. Maurice Rojas*

Yuyu Zhu*

jmauricerojas@gmail.com

yuyu.zhu1213@gmail.com

Texas A&M University

College Station, Texas, USA

ABSTRACT

We reveal a complexity chasm, separating the trinomial and tetranomial cases, for solving univariate sparse polynomial equations over certain local fields. First, for any fixed field $K \in \{\mathbb{Q}_2, \mathbb{Q}_3, \mathbb{Q}_5, \dots\}$, we prove that any polynomial $f \in \mathbb{Z}[x]$ with exactly 3 monomial terms, degree d , and all coefficients having absolute value at most H , can be solved over K within deterministic time $\log^{O(1)}(dH)$ in the classical Turing model. (The best previous algorithms were of complexity exponential in $\log d$, even for just counting roots in \mathbb{Q}_p .) In particular, our algorithm generates approximations in \mathbb{Q} with bit-length $\log^{O(1)}(dH)$ to all the roots of f in K , and these approximations converge quadratically under Newton iteration. On the other hand, we give a unified family of *tetranomials* requiring $\Omega(d \log H)$ digits to distinguish the base- p expansions of their roots in K .

CCS CONCEPTS

• **Theory of computation** → **Algebraic complexity theory**; • **Mathematics of computing** → **Number-theoretic computations**; • **Computing methodologies** → **Number theory algorithms**.

KEYWORDS

p -adic, Newton's method, trinomial, approximate, root counting

ACM Reference Format:

J. Maurice Rojas and Yuyu Zhu. 2021. A Complexity Chasm for Solving Univariate Sparse Polynomial Equations Over p -adic Fields. In *Proceedings of the 2021 International Symposium on Symbolic and Algebraic Computation (ISSAC '21)*, July 18–23, 2021, Virtual Event, Russian Federation. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3452143.3465554>

*Partially supported by NSF grants CCF-1900881 and CCF-1409020.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISSAC '21, July 18–23, 2021, Virtual Event, Russian Federation

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8382-0/21/07...\$15.00

<https://doi.org/10.1145/3452143.3465554>

1 INTRODUCTION

Solving polynomial equations over the p -adic rational numbers \mathbb{Q}_p underlies many classical questions in number theory, and is close to numerous applications in cryptography, coding theory, and computational number theory. Furthermore, the complexity of solving *structured* systems — such as those with a fixed number of monomial terms or invariance with respect to a group action — arises naturally in many computational geometric applications and is closely related to a deeper understanding of circuit complexity (see, e.g., [17]). So we will classify when it is possible to separate and approximate roots in \mathbb{Q}_p in deterministic polynomial-time.

Recall that thanks to 17th century work of Descartes, and 20th century work of Lenstra [19] and Poonen [23], it is known that univariate polynomials with exactly t monomial terms have at most $t^{O(1)}$ roots in a fixed field K *only* when K is \mathbb{R} or a finite algebraic extension of \mathbb{Q}_p for some prime $p \in \mathbb{N}$. We'll use $|\cdot|_p$ (resp. $|\cdot|$) for the absolute value on the p -adic complex numbers \mathbb{C}_p [25] normalized so that $|p|_p = \frac{1}{p}$ (resp. the standard absolute value on \mathbb{C}). Recall also that for any function f analytic on K , the corresponding *Newton endomorphism* is $N_f(z) := z - \frac{f(z)}{f'(z)}$, and the corresponding sequence of *Newton iterates* of a *start-point* $z_0 \in K$ is the sequence $(z_i)_{i=0}^\infty$ where $z_{i+1} := N_f(z_i)$ for all $i \geq 0$. Finally, we call any polynomial in $\mathbb{Z}[x_1, \dots, x_n]$ having exactly t terms in its monomial term expansion an *n-variate t-nomial*. We will often use x in place of x_1 .

Our first main result is that we can efficiently count the roots of trinomials in \mathbb{Q}_p , and find succinct start-points in \mathbb{Q} under which Newton iteration converges quickly to all the roots in \mathbb{Q}_p . We use $\#S$ for the cardinality of a set S .

Theorem 1.1. *Suppose $K = \mathbb{Q}_p$ for some fixed¹ prime $p \in \mathbb{N}$. Then there is an algorithm which, for any input trinomial $f \in \mathbb{Z}[x]$ with degree d and all coefficients of (Archimedean) absolute value $\leq H$, outputs a set $\{\frac{a_1}{b_1}, \dots, \frac{a_m}{b_m}\} \subset \mathbb{Q}$ of cardinality $m = m(K, f)$ such that:*

1. For all j we have $a_j \neq 0 \implies \log |a_j|, \log |b_j| = O(\log^8(dH))$.
2. There is a $\mu = \mu(d, H) > 1$ such that for all j we have that $z_0 := a_j/b_j$ implies that f has a root $\zeta_j \in K$ such that the corresponding sequence of Newton iterates satisfies $|z_i - \zeta_j|_p \leq \mu^{-2^{i-1}} |z_0 - \zeta_j|_p$ for all $i \geq 1$.
3. $m = \#\{\zeta_1, \dots, \zeta_m\}$ is exactly the number of roots of f in K .

Moreover, the underlying algorithm (Algorithm 5.9 in Section 5 below) takes deterministic time $O(\log^{16}(dH) \log \log(dH))$.

We prove Theorem 1.1 in Section 5. (An analogue of Theorem 1.1 in fact holds for $K = \mathbb{R}$ as well, and will be presented in a sequel to this

¹We clarify the dependence of our complexity bounds on p in Section 5.

paper.) We will call the convergence condition on z_0 above *being an approximate root (in the sense of Smale²)*, with associated true root ζ_j . This type of convergence provides an efficient encoding of an approximation that can be quickly tuned to any desired accuracy.

Remark 1.2. *Defining the input size of a univariate polynomial $f(x) := \sum_{i=1}^t c_i x^{a_i} \in \mathbb{Z}[x]$ as $\sum_{i=1}^t \log(|c_i| + 2)(|a_i| + 2)$ we see that Theorem 1.1 implies that one can solve univariate trinomial equations, over any fixed p -adic field, in deterministic time polynomial in the input size. \diamond*

Remark 1.3. *Efficiently solving univariate t -nomial equations over K in the sense of Theorem 1.1 is easier for $t \leq 2$: The case $t = 1$ is clearly trivial (with 0 the only possible root) while the case $(K, t) = (\mathbb{R}, 2)$ is implicit in work on computer arithmetic from the 1970s (see, e.g., [9]). We review the case $(K, t) = (\mathbb{Q}_p, 2)$ with p prime in Corollary 2.6 and Theorem 2.16 of Section 2 below. \diamond*

Despite much work on factoring univariate polynomials over \mathbb{Q}_p (see, e.g., [7, 8, 10, 13]), all known general algorithms for solving (or even just counting the solutions of) arbitrary degree d polynomial equations over \mathbb{Q}_p have complexity exponential in $\log d$. So Theorem 1.1 presents a significant new speed-up, and improves an earlier complexity bound (membership in NP, for detecting roots in \mathbb{Q}_p) from [1]. We'll see in Section 3 how our speed-up depends on p -adic Diophantine approximation [35]. Another key new ingredient in proving Theorem 1.1 is an efficient encoding of roots in $\mathbb{Z}/(p^k)$ from [12, 18] (with an important precursor in [8]).

1.1 Why is the Field Fixed?

Much as real algebraic geometry fixes the underlying field to $K = \mathbb{R}$ once and for all, our results focus on $K = \mathbb{Q}_p$ with p fixed once and for all. In particular, while there are certainly number-theoretic algorithms with deterministic complexity having dependence $(\log p)^{O(1)}$ on an input prime p , solving sparse polynomial equations in one variable over \mathbb{Q}_p appears to have much larger complexity as a function of p . There are naive reasons, and subtle reasons, for this:

R1. *Whereas a binomial has at most 3 roots in \mathbb{R} (e.g., $x^3 - x$), a binomial can have as many as p roots in \mathbb{Q}_p (e.g., $x^p - x$). Furthermore, trinomials have at most 5, 7, 9, or $3p - 2$ roots in K , according as K is \mathbb{R} , \mathbb{Q}_2 [19], \mathbb{Q}_3 [36], or \mathbb{Q}_p with $p \geq 5$ [4], and each bound is sharp.*

R2. *Approximating square-roots of p -adic integers not divisible by p , within accuracy 1, is equivalent to finding square-roots in the finite field \mathbb{F}_p . The latter problem is still not known to be doable in deterministic time polynomial in $\log p$, even though the decision version is doable in deterministic polynomial-time (see, e.g., [5, 24]).*

In particular, even if one only wants to approximate just one or two roots in \mathbb{Q}_p , the minimal *currently* provable accuracy needed to decide if two approximations converge to the same root appears to have *quasi-linear* dependence on p . Interestingly, the truth of strong forms of the abc-Conjecture would imply a much smaller and practical dependence on p : See [6, 26] and Section 3 below.

1.2 The Separation Chasm at Four Terms

The p -adic rational roots of sparse polynomials can range from well-separated to (possibly) tightly spaced, already with just 4 terms.

²This terminology has only been applied over \mathbb{C} with $\mu = 2$ so far [31], so we take the opportunity here to extend it to the p -adic rationals.

Theorem 1.4. *Consider the family of tetranomials*

$$f_{d,\varepsilon}(x) := x^d - \varepsilon^{-2h}x^2 + 2\varepsilon^{-(h+1)}x - \varepsilon^{-2}$$

with $h \in \mathbb{N}$, $h \geq 3$, and $d \in \{4, \dots, \lfloor e^h \rfloor\}$ even. Let $H := \max\{\varepsilon^{\pm 2h}\}$. Then $f_{d,\varepsilon}$ has distinct roots $\zeta_1, \zeta_2 \in K$ with $|\log|\zeta_1 - \zeta_2||$ or $|\log|\zeta_1 - \zeta_2||$ of order $\Omega(d \log H)$, according as $(K, \varepsilon) = (\mathbb{Q}_p, p)$ or $(K, \varepsilon) = (\mathbb{R}, 1/2)$. In particular, the coefficients of $p^{2h}f_{d,p}$ (resp. $f_{d,1/2}$) all lie in \mathbb{Z} and have $O(\log H)$ base- p digits (resp. bits).

We prove Theorem 1.4 in Section 4. The special case $K = \mathbb{R}$ was derived earlier (in different notation) by Mignotte [21]. (See also [27].) The case $K = \mathbb{Q}_p$ with p prime appears to be new, and our proof unifies the Archimedean and non-Archimedean cases via tropical geometry. Note that Theorem 1.4 implies that the roots in K of a tetranomial can be so close that one needs $\Omega(d \log H)$ many digits to distinguish their base- p expansions in the worst case.

Mignotte used the tetranomial $f_{d,1/2}$ in [21] to show that an earlier root separation bound of Mahler [20], for *arbitrary* degree d polynomials in $\mathbb{Z}[x]$, is asymptotically near-optimal. We recall the following paraphrased version:

Mahler's Theorem. *Suppose $f \in \mathbb{Z}[x]$ has degree d , all coefficients of (Archimedean) absolute value at most H , and is irreducible in $\mathbb{Z}[x]$. Let $\zeta_1, \zeta_2 \in \mathbb{C}$ be distinct roots of f . Then $|\log|\zeta_1 - \zeta_2|| = O(d \log(dH))$. ■*

Our new algorithmic results are enabled by our third and final main result: Mahler's bound can be dramatically improved for trinomials.

Theorem 1.5. *Suppose p is prime and $f \in \mathbb{Z}[x]$ is square-free, has exactly 3 monomial terms, degree d , and all coefficients of (Archimedean) absolute value at most H . Let $\zeta_1, \zeta_2 \in \mathbb{C}_p$ be distinct roots of f . Then $|\log|\zeta_1 - \zeta_2||_p = O\left(\frac{p}{\log^2 p} \log(d) \log^2(dH + p) \log \log(dH + p)\right)$.*

We prove Theorem 1.5 in Section 3. Theorem 1.5 is in fact a p -adic analogue of a separation bound of Koiran for roots in \mathbb{R} [16]. Even sharper bounds can be derived for binomials: We review these bounds in Section 2.2.

1.3 Previous Complexity and Sparsity Results

Deciding the existence of roots over \mathbb{Q}_p for univariate polynomials with an *arbitrary* number of monomial terms is already NP-hard with respect to randomized (ZPP, a.k.a. Las Vegas) reductions [1]. On the other hand, detecting roots over \mathbb{Q}_p for n -variate $(n + 1)$ -nomials is known to be doable in NP [1]. Speeding this up to polynomial-time, even for $n = 2$ and fixed p , hinges upon detecting roots in $(\mathbb{Z}/(p^k))^2$ for bivariate trinomials of degree d in time $(k + \log d)^{O(1)}$. The latter problem remains open, but some progress has been made in author Zhu's Ph.D. thesis [36].

On a related note, counting points on trinomial curves over the prime fields \mathbb{F}_p in time $(\log(pd))^{O(1)}$ remains a challenging open question. Useful quantitative estimates in this direction were derived in [15] and recently given an alternative proof via quadratic optimization over \mathbb{R} [3].

2 BACKGROUND


2.1 Newton Polygons and Newton Iteration

Definitive sources for p -adic arithmetic and analysis include [25, 28, 29]. We use $\text{ord}_p : \mathbb{C}_p \rightarrow \mathbb{Q}$ for the standard p -adic valuation on \mathbb{C}_p , normalized so that $\text{ord}_p p = 1$. The *most significant* (p -adic) digit of $\sum_{j=s}^{\infty} a_j p^j \in \mathbb{Q}_p$ is a_s , assuming the $a_j \in \{0, \dots, p - 1\}$ and $a_s \neq 0$.

The notion of Newton polygon goes back to 17th century work of Newton on Puiseux series solutions to polynomial equations [32, pp. 126–127]. We will need variants of this notion over \mathbb{C}_p and \mathbb{C} . (See, e.g., [34] for the p -adic case and [2, 22] for the complex case.)

Definition 2.1. Suppose $f(x) := \sum_{i=1}^t c_i x^{a_i} \in \mathbb{Z}[x]$ with $c_i \neq 0$ for all i and $a_1 < \dots < a_t$. We then define the p -adic Newton polygon, $\text{Newt}_p(f)$ (resp. Archimedean Newton polygon, $\text{Newt}_\infty(f)$) to be the convex hull of the set of points $\{(a_i, \text{ord}_p c_i) \mid i \in \{1, \dots, t\}\}$ (resp. the convex hull of $\{(a_i, -\log |c_i|) \mid i \in \{1, \dots, t\}\}$). We call an edge E of a polygon in \mathbb{R}^2 lower if and only if E has an inner normal with positive last coordinate. We also define the horizontal length of a line segment E connecting (r, s) and (u, v) to be $\lambda(E) := |u - r|$. \diamond

Example 2.2. Following the notation of Theorem 1.4, we set $h = 3$ and illustrate $\text{Newt}_p(f_{5,p})$ (for p odd) and $\text{Newt}_\infty(f_{5,1/2})$ below:



Note that the p -adic Newton polygon on the left has exactly 2 lower edges (with horizontal lengths 2 and 3), while the Archimedean Newton polygon on the right has exactly 3 lower edges (with horizontal lengths 1, 1, and 3). \diamond

Theorem 2.3. Following the notation above, the number³ of roots of f in \mathbb{C}_p of valuation v is exactly the horizontal length of the face of $\text{Newt}_p(f)$ with inner normal $(v, 1)$. Furthermore, if $\text{Newt}_\infty(f)$ has a lower edge E with slope v , and no other lower edges with slope in the open interval $(v - \log 3, v + \log 3)$, then the number³ of roots $\zeta \in \mathbb{C}$ of f with $\log |\zeta| \in (v - \log 3, v + \log 3)$ is exactly $\lambda(E)$. \blacksquare

The first portion of Theorem 2.3 goes back to early 20th century work of Hensel, while the second portion is an immediate consequence of [2, Thm. 1.5] (with an important precursor in [22]).

We will also use the following version of Hensel's famous criterion for the rapid convergence of Newton's method over \mathbb{C}_p :

Hensel's Lemma. (See, e.g., [11, Thm. 4.1 & Inequality (5.7)].) Suppose p is prime, $f \in \mathbb{Z}[x]$, $j \geq 1$, $\zeta \in \mathbb{Z}_p$, $\ell = \text{ord}_p f'(\zeta) < \infty$, and $f(\zeta) \equiv 0 \pmod{p^{2\ell+j}}$. Let $\zeta' := \zeta - \frac{f(\zeta)}{f'(\zeta)}$. Then $f(\zeta') \equiv 0 \pmod{p^{2\ell+2j}}$, $\text{ord}_p f'(\zeta') = \ell$, and $\zeta \equiv \zeta' \pmod{p^{\ell+j}}$. \blacksquare

2.2 Separating Roots of Binomials

When $f \in \mathbb{Z}[x]$ is a binomial, all of its roots in \mathbb{C} are multiples of roots of unity that are evenly spaced on a circle. The same turns out to be true over \mathbb{C}_p , but the root spacing then depends more subtly on p and less on the degree. For convenience, we will sometimes write $|\cdot|_\infty$ instead of $|\cdot|$ for the standard norm on \mathbb{C} . In summary, we have the following:

Proposition 2.4. Suppose $f(x) := c_1 + c_2 x^d \in \mathbb{Z}[x]$, $c_1 c_2 \neq 0$, and $|c_1|, |c_2| \leq H$. Also let $p \in \{\infty, 2, 3, 5, \dots\}$ and let \bar{K}_p denote \mathbb{C} or \mathbb{C}_p , according as $p = \infty$ or p is prime. Then for any distinct roots $\zeta_1, \zeta_2 \in \bar{K}_p$ of f , we have that $|\log |\zeta_1 - \zeta_2||_p$ is bounded from above by

$$\begin{cases} \log(d) + \log(H^{1/d}) + \frac{3}{2} \log 3; & \text{for } p = \infty \text{ and } d \geq 3, \\ \log(H^{1/d}); & \text{for } d > p^{\text{ord}_p d}, \text{ and} \\ \log(H^{1/d}) + \frac{\log p}{p-1}; & \text{for } d = p^{\text{ord}_p d} \geq p. \end{cases}$$

Proof: See Appendix of ArXiv version [26]. \blacksquare

2.3 Counting Roots of Binomials Over \mathbb{Q}_p^*

For any ring R we let R^* denote the multiplicatively invertible elements of R . Counting roots of binomials over \mathbb{Q}_p is more involved than counting their roots over \mathbb{R} , but is still quite efficiently doable.

Lemma 2.5. Suppose p is an odd prime and $f(x) := c_1 + c_2 x^d \in \mathbb{Z}[x]$ with $|c_1|, |c_2| \leq H$, $c_1 c_2 \neq 0$, and $\ell := \text{ord}_p d$. Then the number of roots of f in \mathbb{Q}_p is either 0 or $\gcd(d, p-1)$. In particular, f has roots in \mathbb{Q}_p if and only if both of the following conditions hold: (1) $d \mid \text{ord}_p(c_1/c_2)$ and (2) $\left(-\frac{c_1}{c_2} p^{\text{ord}_p(c_2/c_1)}\right)^{p^\ell(p-1)/\gcd(d, p-1)} \equiv 1 \pmod{p^{2\ell+1}}$. \blacksquare

Lemma 2.5 is classical and follows from basic group theory and [1, Cor. 3.2]. The case $p = 2$ is slightly more involved and is stated in the Appendix of the ArXiv version of this paper [26].

Corollary 2.6. Following the notation and assumptions of Lemma 2.5, one can count exactly the number of roots of f in \mathbb{Q}_p in time $O((\log(dpH) \log \log(dpH))^2)$. Furthermore, for any root $\zeta \in \mathbb{Q}_p^*$ there is an $x_0 \in \mathbb{Z} \setminus \{0\}$ that is a root of the mod $p^{2\ell+1}$ reduction of $\frac{c_1}{p^{\text{ord}_p c_1}} + \frac{c_2}{p^{\text{ord}_p c_2}} x^d$, and with $z_0 := p^{\text{ord}_p(c_2/c_1)/d} x_0 \in \mathbb{Q}$ an approximate root of f with associated true root ζ . In particular, the logarithmic height⁴ of z_0 is $O(\log(dH^{1/d}))$.

Proof: See Appendix of ArXiv version [26]. \blacksquare

We now detail how to efficiently find the magical x_0 above to start our p -adic approximations.

2.4 Trees and Roots in $\mathbb{Z}/(p^k)$ and \mathbb{Z}_p

The p -adic analogue of bisecting an isolating interval containing a real root is to approximate the next base- p digit of an approximate root in \mathbb{Q}_p . Shifting from bisecting intervals to extracting digits is crucial since \mathbb{Q}_p is not an ordered field. We will write f' for the derivative of f and $f^{(i)}$ for the i th order derivative of f .

Definition 2.7. [18] For any $f \in \mathbb{Z}[x]$ let \tilde{f} denote the mod p reduction of f . A root $\zeta_0 \in \mathbb{F}_p$ of \tilde{f} is degenerate if and only if $\tilde{f}'(\zeta_0) \equiv 0 \pmod{p}$. For any degenerate root ζ_0 of \tilde{f} (represented as an element of $\{0, \dots, p-1\}$), we then define $s(f, \zeta_0) := \min_{i \geq 0} \{i + \text{ord}_p \frac{f^{(i)}(\zeta_0)}{i!}\}$. Fixing $k \in \mathbb{N}$, for $i \geq 1$, let us inductively define a set $T_{p,k}(f)$ of pairs $(f_{i-1, \mu}, k_{i-1, \mu}) \in \mathbb{Z}[x] \times \mathbb{N}$: We set $(f_{0,0}, k_{0,0}) := (f, k)$. Then for any $i \geq 1$ with $(f_{i-1, \mu}, k_{i-1, \mu}) \in T_{p,k}(f)$, and any degenerate root $\zeta_{i-1} \in \{0, \dots, p-1\}$ of $\tilde{f}_{i-1, \mu}$ with $s_{i-1} := s(f_{i-1, \mu}, \zeta_{i-1}) \in \{2, \dots, k_{i-1, \mu} - 1\}$, we define $\zeta := \mu + p^{i-1} \zeta_{i-1}$, $k_{i, \zeta} := k_{i-1, \mu} - s_{i-1}$, $f_{i, \zeta}(x) := p^{-s(f_{i-1, \mu}, \zeta_{i-1})} f_{i-1, \mu}(\zeta_{i-1} + px) \pmod{p^{k_{i, \zeta}}}$, and then include $(f_{i, \zeta}, k_{i, \zeta})$ in $T_{p,k}(f)$. \diamond

³counting multiplicity

⁴The logarithmic height of a rational number a/b with $\gcd(a, b) = 1$ is simply $\log \max\{|a|, |b|\}$ (and we declare the logarithmic height of 0 to be 0).

Example 2.8. If $f(x) = x^{10} - 10x + 738$ and $p = 3$ then $\tilde{f}(x) = x(x-1)^9 \pmod{3}$, 1 is a degenerate root of \tilde{f} , and one can check that $s(f, 1) = 4$ (no greater than the multiplicity of the factor $x-1$ in \tilde{f}). In particular, $f_{1,1}$ has degree 10 (and 10 monomial terms) but $\tilde{f}_{1,1} = x^3 + 2x^2$. \diamond

The collection of pairs $(f_{i,\zeta}, k_{i,\zeta})$ admits a tree structure that will give us a way to extend Hensel lifting to degenerate roots.

Definition 2.9. [18] Let us identify the elements of $\mathcal{T}_{p,k}(f)$ with nodes of a labelled, rooted, directed tree $\mathcal{T}_{p,k}(f)$ defined inductively as follows⁵: (i) We set $f_{0,0} := f$, $k_{0,0} := k$, and let $(f_{0,0}, k_{0,0})$ be the label of the root node of $\mathcal{T}_{p,k}(f)$.

- (ii) The non-root nodes of $\mathcal{T}_{p,k}(f)$ are uniquely labelled by each $(f_{i,\zeta}, k_{i,\zeta}) \in \mathcal{T}_{p,k}(f)$ with $i \geq 1$.
- (iii) There is an edge from node $(f_{i-1,\mu}, k_{i-1,\mu})$ to node $(f_{i,\zeta}, k_{i,\zeta})$ if and only if there is a degenerate root $\zeta_{i-1} \in \{0, \dots, p-1\}$ of $\tilde{f}_{i-1,\mu}$ with $s(f_{i-1,\mu}, \zeta_{i-1}) \in \{2, \dots, k_{i-1,\mu} - 1\}$ and $\zeta = \mu + p^{i-1}\zeta_{i-1} \in \mathbb{Z}/(p^i)$. \diamond

We call each $f_{i,\zeta}$ with $(f_{i,\zeta}, k_{i,\zeta}) \in \mathcal{T}_{p,k}(f)$ a *nodal polynomial* of $\mathcal{T}_{p,k}(f)$. It is in fact possible to list all the roots of f in $\mathbb{Z}/(p^k)$ from the data contained in $\mathcal{T}_{p,k}(f)$ [12, 18]. We will instead use $\mathcal{T}_{p,k}(f)$, with k chosen via our root separation bounds, to efficiently count the roots of f in \mathbb{Z}_p , and then in \mathbb{Q}_p by rescaling.

Example 2.10. $\mathcal{T}_{p,k}(x^2)$ is a chain of length $\lfloor \frac{k-1}{2} \rfloor$ for any p, k . \diamond

Example 2.11. Let $f(x) = 1 - x^{397}$. Then $\mathcal{T}_{17,k}(f)$, for any $k \geq 1$, consists of a single node, labelled $(1 - x^{397}, k)$, since \tilde{f} has no degenerate roots in \mathbb{F}_{17} . In particular, f has 1 as its only root in \mathbb{Q}_{17} . \diamond

Example 2.12. Let $f(x) = 1 - x^{340}$. Then, when $k \in \{1, 2\}$, the tree $\mathcal{T}_{17,k}(f)$ consists of a single root node, labelled $(1 - x^{340}, k)$. However, when $k \geq 3$, the tree $\mathcal{T}_{17,k}(f)$ has depth 1, and consists of the aforementioned root node and exactly 4 child nodes, labelled $(f_{1,\zeta_0}, k-2)$ where the \tilde{f}_{1,ζ_0} are, respectively, $14x$, $12x + 10$, $5x + 15$, and $3x + 3$. Note that \tilde{f} has exactly 4 roots $\zeta_0 \in \mathbb{F}_{17}$ (1, 4, 13, and 16), each of which is degenerate, and the roots $\zeta_1 \in \mathbb{F}_{17}$ of the \tilde{f}_{1,ζ_0} encode the “next” base-17 digits (0, 2, 14, and 16) of the roots of f in $\mathbb{Z}/(17^2)$. In particular, the roots of f in \mathbb{Q}_{17} are $1 + 0 \cdot 17 + \dots$, $4 + 2 \cdot 17 + \dots$, $13 + 14 \cdot 17 + \dots$, and $16 + 16 \cdot 17 + \dots$ and are all non-degenerate. \diamond

Nodal polynomials thus encode individual base- p digits of roots of f in \mathbb{Z}_p . Their degree also decays in a manner depending on root multiplicity⁶ as follows:

Lemma 2.13. [18, Lem. 2.2 & 3.6] Following the notation of Definition 2.9, suppose $i \geq 1$, ζ_{i-1} has multiplicity m over \mathbb{F}_p , and $(f_{i,\zeta}, k_{i,\zeta}) \in \mathcal{T}_{p,k}(f)$. Then $\mathcal{T}_{p,k}(f)$ has depth $\leq \lfloor (k-1)/2 \rfloor$ and $\deg \tilde{f}_{i,\zeta} \leq s(f_{i-1,\mu}, \zeta_{i-1}) \leq \min\{k_{i-1,\mu} - 1, m\}$. Also, $f_{i,\zeta}(x) = \frac{1}{p^s} f(\zeta_0 + \zeta_1 p + \dots + \zeta_{i-1} p^{i-1} + p^i x)$ where $s := \sum_{j=0}^{i-1} s(f_{j,\zeta_0 + \dots + \zeta_{j-1} p^{j-1}}, \zeta_j)$. \blacksquare

Let $n_p(f)$ denote the number of non-degenerate roots in \mathbb{F}_p of the mod p reduction of f .

⁵This definition differs slightly from the original in [18].

⁶Over any field K , we define the multiplicity of a root $\zeta \in K$ of $f \in K[x]$ as the greatest m with $(x - \zeta)^m \mid f$ in $K[x]$.

Lemma 2.14. If $f \in \mathbb{Z}[x]$, D is the maximum of $\text{ord}_p(\zeta_1 - \zeta_2)$ over all $\zeta_1, \zeta_2 \in \mathbb{Z}_p$ with $f(\zeta_1) = f(\zeta_2) = 0 \neq \zeta_1 - \zeta_2$, and $k \geq 1 + D$, then f has exactly $\sum_{(g,j) \in \mathcal{T}_{p,k}(f)} n_p(g)$ non-degenerate roots in \mathbb{Z}_p .

Proof: By Lemma 2.13, $f_{i,\zeta}(x) = \frac{1}{p^s} f(\zeta_0 + \dots + \zeta_{i-1} p^{i-1} + p^i x)$. So by Hensel’s Lemma, any root $\zeta_i \in \mathbb{F}_p$ of $f_{i,\zeta}$ lifts to a unique non-degenerate root $\zeta_i + p\zeta_{i+1} + \dots \in \mathbb{Z}_p$ of $f_{i,\zeta}$. In other words, we obtain $\zeta_0 + \zeta_1 p + \dots \in \mathbb{Z}_p$ as a root of f . Any sequence $(\zeta_0, \dots, \zeta_{i-1}) \in \mathbb{F}_p^i$ defined by a nodal polynomial $f_{i,\zeta}$ thus determines a unique root in \mathbb{Z}_p of f , and we thus see that $\sum_{(g,j) \in \mathcal{T}_{p,k}(f)} n_p(g)$ is a lower bound on the number of non-degenerate roots of f in \mathbb{Z}_p .

To see that we obtain all non-degenerate roots of f in \mathbb{Z}_p this way, note that the mod p^k reduction of any root of f in \mathbb{Z}_p is a root of the mod p^k reduction of f in $\mathbb{Z}/(p^k)$. By the definition of k , the resulting map is an injection since distinct roots in \mathbb{Z}_p must differ somewhere within their $1 + D$ most significant digits. \blacksquare

2.5 Trees and Extracting Digits of Radicals

We prove the following useful lemma in Remark 5.7 of Section 5:

Lemma 2.15. Suppose $f(x) = c_1 + c_2 x^d \in \mathbb{Z}[x]$ with $c_1 c_2 \neq 0$. Then every non-root nodal polynomial $f_{i,\zeta}$ for $\mathcal{T}_{p,k}(f)$ with $\zeta \in (\mathbb{Z}/(p^i))^*$ satisfies $\deg \tilde{f}_{i,\zeta} \leq 2$ or $\deg \tilde{f}_{i,\zeta} \leq 1$, according as $p = 2$ or $p \geq 3$.

With our tree-based encoding of p -adic roots in place, we can now prove that it is easy to find approximate roots in \mathbb{Q}_p for binomials when p is fixed.

Theorem 2.16. Suppose $f \in \mathbb{Z}[x]$ is a binomial of degree d with coefficients of absolute value at most H , $f(0) \neq 0$, $\gamma = \gcd(d, p-1)$, and $\{\zeta_1, \dots, \zeta_\gamma\}$ is the set of roots of f in \mathbb{Q}_p . Then in time

$$O(p \log(dp) \log \log(dp) + [\log(dpH) \log \log(dpH)]^2),$$

we can find, for each $j \in \{1, \dots, \gamma\}$, a $z_0^{(j)} \in \mathbb{Q}$ of logarithmic height $O(\log(dH^{1/d}))$ that is an approximate root of f with associated true root ζ_j .

An algorithm that proves Theorem 2.16, and more, is outlined below. See the ArXiv version of this paper [26] for the case $p = 2$.

Algorithm 2.17. (Solving Binomial Equations Over \mathbb{Q}_p^*)

Input. An odd prime p and $c_1, c_2, d \in \mathbb{Z} \setminus \{0\}$ with $|c_i| \leq H$ for all i .

Output. A true declaration that $f(x) := c_1 + c_2 x^d$ has no roots in \mathbb{Q}_p ,

or $z_1, \dots, z_\gamma \in \mathbb{Q}$ with logarithmic height $O(\log(dH^{1/d}))$ such that $\gamma = \gcd(d, p-1)$, z_j is an approximate root of f with associated true root $\zeta_j \in \mathbb{Q}_p$ for all j , and the ζ_j are pair-wise distinct.

Description.

- 1: If $\text{ord}_p c_1 \neq \text{ord}_p c_2 \pmod{d}$ then say “No roots in $\mathbb{Q}_p!$ ” and STOP.
- 2: Rescale roots so that we may replace c_i by $\frac{c_i}{p^{\text{ord}_p c_i}}$ for all i .
- 3: Invert roots (if needed) so that $d \geq 1$, and let $\ell := \text{ord}_p d$.
- 4: If $\left(-\frac{c_1}{c_2}\right)^{p^\ell(p-1)/\gamma} \not\equiv 1 \pmod{p^{2\ell+1}}$ then say “No roots in $\mathbb{Q}_p!$ ” and STOP.
- 5: Let g be any generator for \mathbb{F}_p^* , $r := (d/\gamma)^{-1} \pmod{p-1}$, $c := (-c_1/c_2)^r \pmod{p}$, and let $\tilde{h}(x) := x^r - c$.
- 6: Find a root $x_1 \in \left\{g^0, \dots, g^{\frac{p-1}{\gamma}-1}\right\}$ of \tilde{h} via brute-force search.
- 7: For all $j \in \{2, \dots, \gamma\}$ let $x_j := x_{j-1} g^{(p-1)/\gamma} \pmod{p}$.
- 8: If $\ell \geq 1$ then, for each $j \in \{1, \dots, \gamma\}$, replace x_j by $x_j - \frac{f(x_j)/p^\ell}{f'(x_j)/p^\ell} \in \mathbb{Z}/(p^{2\ell+1})$.
- 9: Output $\left\{x_1 p^{\text{ord}_p(c_2/c_1)/d}, \dots, x_\gamma p^{\text{ord}_p(c_2/c_1)/d}\right\}$.

Remark 2.18. If one allows randomization then one can considerably speed up Step 6: See, e.g., [5, Ch. 7, Sec. 3]. \diamond

Proof of Theorem 2.16: It clearly suffices to prove the correctness of Algorithm 2.17, and then analyze complexity. In particular, we assume p is odd henceforth. (See [26] for the case $p=2$.)

Correctness: Theorem 2.3 implies that Step 1 merely checks whether the valuations of the roots of f in \mathbb{C}_p^* in fact lie in \mathbb{Z} , which is necessary for f to have roots in \mathbb{Q}_p^* .

Steps 2 and 3 merely involve substitutions of the form $x \leftarrow p^\delta x$ and $x \leftarrow x^{-1}$, and it is elementary to check that the bit length of any resulting approximate roots in \mathbb{Q} remains within $O(\log(dH^{1/d}))$.

Lemma 2.5 implies that Step 4 merely check that the coset of roots of f in \mathbb{C}_p^* intersects \mathbb{Z}_p^* .

Step 5 is merely the application of an automorphism of \mathbb{F}_p^* (that preserves the existence of roots of \tilde{f} in \mathbb{F}_p^*) that enables us to work with a binomial of degree $\gamma \leq d$.

Steps 6–9 then clearly find the correct coset of \mathbb{F}_p^* that makes f vanish. In particular, since \mathbb{F}_p^* is cyclic, Step 10 clearly gives the correct output if $\ell = 0$.

If $\ell > 0$ then, because $\deg \tilde{f}_{1, \zeta_0} \leq 1$ thanks to Lemma 2.15, Definition 2.7 tells us that the unique root of f_{1, ζ_0} in $\mathbb{Z}/(p^{2\ell+1-s})$ (with $s := s(f, \zeta_0)$) determines the next $2\ell + 1 - s$ base- p digits of a unique root of f in \mathbb{Z}_p . (See also the proof of [18, Lemma 1.5].) So Steps 7–9 indeed give us suitable approximants in \mathbb{Q} to all the roots in \mathbb{Z}_p of (our renormalized) f in \mathbb{Z}_p . So our algorithm is correct.

Note also that the outputs, being numbers in $\mathbb{Z}/(p^{2\ell+1})$, rescaled by a factor of $p^{\text{ord}_p(c_2/c_1)/d}$, clearly each have bit-length $O(\text{ord}_p(d) \log(p) + \frac{|\log(c_2/c_1)|}{d \log p} \log p) = O(\log(d) + \frac{\log H}{d}) = O(\log(dH^{1/d}))$. \blacksquare

Complexity Analysis: Via Corollary 2.6, [30], and some additional elementary bit complexity estimates for modular arithmetic [33], it is clear that, *save for Steps 6–9*, Algorithm 2.17 has complexity

$$O(p^{1/4} \log(p) \log \log p + [\log(dpH) \log \log(dpH)]^2),$$

provided we use Harvey and van der Jooen's recent fast multiplication algorithm [14]. Steps 6–7 (whose complexity dominates the complexity of Steps 6–9), involve $\frac{p-1}{\gamma} - 1$ multiplications in \mathbb{F}_p and $\gamma - 1$ multiplications in $\mathbb{Z}/(p^{2\ell+1})$. This takes time no worse than $O(p \log(dp) \log \log(dp))$, so we are done. \blacksquare

3 PROVING THEOREM 1.5

Let us first recall the following version of *Yu's Theorem*:

Theorem 3.1. [35, Cor. 1] *Suppose p is any prime, $n \geq 2$, $\alpha_1, \dots, \alpha_n \in \mathbb{Q}$ with $\alpha_i = r_i/s_i$ a reduced fraction for each i , and $b_1, \dots, b_n \in \mathbb{Z}$ are not all zero. Then $\alpha_1^{b_1} \dots \alpha_n^{b_n} \neq 1$ implies that $\alpha_1^{b_1} \dots \alpha_n^{b_n} - 1$ has p -adic valuation bounded from above by*

$$11145 \left(\frac{24(n+1)^2}{\log p} \right)^{n+2} (p-1) \left(\prod_{i=1}^n \log A_i \right) \log(4B)$$

$$\times \max \left\{ \log(2^{12} \cdot 3n(n+1) \log A_n), \frac{\log p}{n} \right\},$$

where $B := \max\{|b_1|, \dots, |b_n|, 3\}$, and A_1, \dots, A_n are any real numbers such that $A_1 \leq \dots \leq A_n$ and, for each j , $A_j \geq \max\{|r_j|, |s_j|, p\}$. \blacksquare

To prove that two distinct roots $\zeta_1, \zeta_2 \in \mathbb{C}_p$ of a trinomial f can not be too close, we will build a special point $m \in \mathbb{C}$ with four special properties: (i) $f'(m) = 0$, (ii) $|f(m)|_p$ is not too small,

(iii) $|\zeta_1 - \zeta_2|_p \geq p^{-1/(p-1)} |\zeta_1 - m|_p$, and (iv) $|\zeta_1 - m|$ is not too small. So let us quantify this approach toward proving Theorem 1.5.

Proposition 3.2. *Let $f(x) = c_1 + c_2x^{a_2} + c_3x^{a_3} \in \mathbb{Z}[x]$ with $a_3 > a_2 \geq 1$, $c_1 \neq 0$, and suppose $m \in \mathbb{C}_p$ is a root of f' . Then $m^{a_3-a_2} = -\frac{a_2c_2}{a_3c_3}$ and $f(m) = c_1 + c_2m^{a_2} \left(1 - \frac{a_2}{a_3}\right)$. \blacksquare*

Lemma 3.3. *Following the notation above, assume further that f is square-free. Then*

$$|f(m)|_p \geq \exp \left(-O \left(\frac{p}{\log^2 p} \log(d) \log^2(dH+p) \log \log(dH+p) \right) \right).$$

Proof: First note that if f is square-free then f has no repeated factors, and thus no degenerate roots in \mathbb{C}_p . So $f(m) \neq 0$. By Proposition 3.2 we then obtain $\text{ord}_p f(m) = \text{ord}_p(c_1 + c_2m^{a_2}(1 - a_2/a_3))$

$$= \text{ord}_p(c_1) + \text{ord}_p(-1) + \text{ord}_p \left(\frac{-c_2(a_3-a_2)}{c_1a_3} \left(-\frac{a_2c_2}{a_3c_3} \right)^{a_2/(a_3-a_2)} - 1 \right). \quad (1)$$

Clearly, $\text{ord}_p c_1 \leq \frac{\log H}{\log p}$ and $\text{ord}_p(-1) = 0$. To bound the third

summand above, let $T := \frac{-c_2(a_3-a_2)}{c_1a_3} \left(-\frac{a_2c_2}{a_3c_3} \right)^{a_2/(a_3-a_2)}$ and observe that $T^{a_3-a_2} - 1 = \prod_{j=1}^{a_3-a_2} (T - \zeta^j)$ for $\zeta \in \mathbb{C}_p$ a primitive $(a_3 - a_2)$ -th root of unity. In particular, $T^{a_3-a_2} \neq 1$ since $f(m\zeta^j) \neq 0$ for all $j \in \{1, \dots, a_3 - a_2\}$, thanks to Proposition 3.2 and f not having any degenerate roots. So then

$$M := \text{ord}_p(T^{a_3-a_2} - 1) = \sum_{j=1}^{a_3-a_2} \text{ord}_p(T - \zeta^j) < \infty,$$

with the $(a_3 - a_2)$ -th term of the sum exactly $\text{ord}_p(T - \zeta^{a_3-a_2}) = \text{ord}_p(T - 1)$, i.e., the third summand from (1).

Suppose $\text{ord}_p T < 0$. Then for each $i \in \{1, \dots, a_3 - a_2\}$, the Ultrametric Inequality gives us $\text{ord}_p(T - \zeta^j) = \text{ord}_p T < 0$, since roots of unity always have p -adic valuation 0. We must then have

$$\text{ord}_p f(m) = \text{ord}_p(c_1) + \text{ord}_p(T - \zeta^{a_3-a_2}) < \frac{\log(dH)}{\log p}$$

and we obtain our lemma.

On the other hand, should $\text{ord}_p T \geq 0$, we get $\text{ord}_p(T - \zeta^j) \geq j \text{ord}_p(\zeta) = 0$, for each j , by the Ultrametric Inequality. So $M \geq \text{ord}_p(T - 1)$ and we'll be done if we find a sufficiently good upper bound on M .

By luck, M is boundable directly from Theorem 3.1, upon setting $n = 2$, $\alpha_1 = -c_2 \frac{a_3-a_2}{c_1a_3}$, $\alpha_2 = -\frac{a_2c_2}{a_3c_3}$, $b_1 = a_3 - a_2$, and $b_2 = a_2$. In particular, we can use $A_i = \max\{dH, p\}$ for $i \in \{1, 2\}$ and $B = \max\{d, 3\}$, yielding $\log A_1, \log A_2, \log B = O(\max\{\log(dH), \log p\})$, so that $M \leq Cp \log^2 \max\{dH, p\} \log(4 \max\{d, 3\})$

$$\times \max \left\{ \log(18 \cdot 2^{12} \log \max\{dH, p\}), \frac{\log p}{2} \right\} \Big/ \log^4 p$$

for $C = 11145 \cdot 216^4$. So then

$$\begin{aligned} M &= O \left(\frac{p}{\log^4 p} \log^2(dH+p) (\log d) (\log \log(dH+p) + \log p) \right) \\ &= O \left(\frac{p}{\log^4 p} \log(d) \log^2(dH+p) \log \log(dH+p) \log p \right) \\ &= O \left(\frac{p}{\log^3 p} \log(d) \log^2(dH+p) \log \log(dH+p) \right). \end{aligned}$$

In other words, the third summand from (1) is bounded from above by the last O -bound, and thus $\text{ord}_p f(m) = O(M)$ since $\frac{\log H}{\log p} = O(M)$. Since $|f(m)|_p = e^{-\log(p) \text{ord}_p f(m)}$, we are done. \blacksquare

Thanks to the Ultrametric Inequality we can effectively bound $|f'(x)|_p$ for $|x|_p \leq 1$ with $f'(x) \neq 0$.

Proposition 3.4. *Following the notation and assumptions of Lemma 3.3 we have $|x|_p \leq 1 \implies |f'(x)|_p \leq 1$. \blacksquare*

The following is a rescaled version of a p -adic version of *Rolle's Theorem* that appeared in [25, Sec. 2.4, Thm., Pg. 316]

Theorem 3.5. *Let $f \in \mathbb{C}_p[x]$ have two distinct roots $\zeta_1, \zeta_2 \in \mathbb{C}_p$ with $|\zeta_1 - \zeta_2|_p = cp^{1/(p-1)}$ for some $c > 0$. Then f' has a root $m \in \mathbb{C}_p$ with $|\zeta_1 - m|_p, |\zeta_2 - m|_p \leq c$. ■*

We can now prove one of our main results.

Proof of Theorem 1.5: For convenience, let us abbreviate the stated O -bound by $O(M)$. Note that if one of the ζ_i is 0 then the other root ζ_j has valuation satisfying $\text{ord}_p \zeta_j \leq \frac{\log H}{\log p}$, thanks to Theorem 2.3. So then

$$|\log |\zeta_i - \zeta_j|_p| = |\log |\zeta_j|_p| = \left| \log e^{-\log(p) \text{ord}_p \zeta_j} \right| = |-\log(p) \text{ord}_p \zeta_j| \leq \log H = O(M).$$

So we may assume $\zeta_1 \zeta_2 \neq 0$ and $f(0) \neq 0$.

Case 1: (Both roots are small: $|\zeta_1|_p, |\zeta_2|_p \leq 1$.)

Suppose $|\zeta_1 - \zeta_2|_p > p^{-2/(p-1)}$. Then $|\zeta_1 - \zeta_2|_p > e^{-2 \log(p)/(p-1)}$. Since $2 \log(p)/(p-1) = O(M)$ we are done.

Now assume that $|\zeta_1 - \zeta_2|_p \leq p^{-2/(p-1)}$. Then by Theorem 3.5 there is an $m \in \mathbb{C}_p$ such that $f'(m) = 0$ and

$$|\zeta_i - m|_p \leq p^{1/(p-1)} |\zeta_1 - \zeta_2|_p \leq p^{-1/(p-1)}$$

for all $i \in \{1, 2\}$. Note that the Ultrametric Inequality implies that $|m|_p \leq p^{-1/(p-1)}$.

Since f is square-free, Lemma 3.3 implies that $|f(m)|_p \geq e^{-O(M)}$.

Applying Theorem 3.5 to $g(x) := f(x) - \frac{f(m) - f(\zeta_1)}{m - \zeta_1} x - \frac{mf(\zeta_1) - \zeta_1 f(m)}{m - \zeta_1}$ (which vanishes at m and ζ_1), we then see that there is a $\zeta \in \mathbb{C}_p$ with $|\zeta - \zeta_1|_p \leq 1$ (and thus $|\zeta|_p \leq 1$) such that $g'(\zeta) = 0$, i.e., $f(m) = f(\zeta) - f(\zeta_1) = f'(\zeta)(m - \zeta_1)$. As $f(m) \neq 0$ we get $f'(\zeta) \neq 0$ and $m \neq \zeta_1$. From Proposition 3.4 we have $|f'(\zeta)|_p \leq 1$, so then $|m - \zeta_1|_p = \frac{|f(m)|_p}{|f'(\zeta)|_p} \geq e^{-O(M)}$. We thus get $|\zeta_1 - \zeta_2|_p \geq p^{-1/(p-1)} |m - \zeta_1|_p \geq e^{-O(M)} \frac{\log p}{p-1} = e^{-O(M)}$. ■

Case 2: (Both roots are large: $|\zeta_1|_p, |\zeta_2|_p > 1$.)

See the Appendix of the ArXiv version of this paper [26]. ■

Case 3: (Only one root has norm > 1 .)

Without loss of generality, we may assume that $|\zeta_1|_p \leq 1 < |\zeta_2|_p$. We then simply note that, as $|\zeta_1|_p \neq |\zeta_2|_p$, we have $|\zeta_1 - \zeta_2|_p = \max \{|\zeta_1|_p, |\zeta_2|_p\} > 1$ and we are done. ■

4 PROVING THEOREM 1.4

4.1 The Case of Prime p

Let $g(x) = p^{2h} f(x + p^{h-1}) = p^{2h} (x + p^{h-1})^d - p^{2h} \left(\frac{x + p^{h-1}}{p^h} - \frac{1}{p} \right)^2 = p^{2h} (x + p^{h-1})^d - x^2$. Then g has the same roots as $f_{d,p}$, save for a "small" shift by p^{h-1} . Rescaling, we get $G(x) := \frac{g(p^{(h-1)d/2+h}x)}{p^{(h-1)d+2h}}$ $= p^{-(h-1)d-2h} \left[p^{2h} (p^{(h-1)d/2+h}x + p^{h-1})^d - p^{(h-1)d+2h} x^2 \right]$ $= \sum_{i=0}^d \binom{d}{i} p^{(h-1)(di/2-i)+ih} x^i - x^2 = 1 - x^2 \pmod{p^{d(h-1)/2+1}}$, which is square-free for odd prime p . (The case of $p = 2$ is in the Appendix of the ArXiv version of this paper [26].) Hensel's Lemma then implies that there are roots $\zeta_1, \zeta_2 \in \mathbb{Z}_p$ of G such that $\zeta_1 \equiv 1 \pmod{p^{d(h-1)/2+1}}$ and $\zeta_2 \equiv -1 \pmod{p^{d(h-1)/2+1}}$.

So $|\zeta_1|_p = |\zeta_2|_p = 1$. For each $i \in \{1, 2\}$, $y_i = p^{(h-1)d/2+h} \zeta_i$ is the corresponding root of G , and thus of g . Then $x_1 = y_1 + p^{h-1}$ and $x_2 = y_2 + p^{h-1}$ are two roots of f in \mathbb{Z}_p such that $|x_1 - x_2|_p = \left| (y_1 + p^{h-1}) - (y_2 + p^{h-1}) \right|_p = |y_1 - y_2|_p \leq \max \{ |y_1|_p, |y_2|_p \} = p^{-(h-1)d/2-h} = p^{-\Omega(dh)}$. ■

4.2 The Case $p = \infty$

While this case was derived earlier in [21], we provide a new proof based entirely on the Archimedean Newton polygon in the Appendix to the ArXiv version of this paper [26].

5 SOLVING TRINOMIALS OVER \mathbb{Q}_p

Unlike the binomial case, the tree $\mathcal{T}_{p,k}(f)$ can have high depth for large k and f an arbitrary trinomial. However, Lemma 5.2 below will show that the structure of $\mathcal{T}_{p,k}(f)$ is still simple: Depth no greater than $\lfloor (k-1)/2 \rfloor$, and all but possibly one path in $\mathcal{T}_{p,k}(f)$ having more than 2 vertices of out-degree ≥ 2 . We will prove an upper bound on k that is large enough to count all roots in \mathbb{Z}_p (via Lemma 2.14), but still small enough for us to approximate all these roots in time $p^{5+o(1)} \log^{16+o(1)}(dH)$.

We begin with a central bound, derived via Theorem 3.1:

Theorem 5.1. *If $f(x) = c_1 + c_2 x^{a_2} + c_3 x^{a_3} \in \mathbb{Z}[x]$ is a trinomial of degree $d = a_3 > a_2 \geq 1$, with coefficients of absolute value at most H , then $\sum_{\zeta \in \mathbb{Z}_p : f(\zeta) = \text{ord}_p \zeta = 0 \neq f'(\zeta)} \text{ord}_p f'(\zeta) = O(p^2 \log^8(dH))$.*

Proof: See Appendix of ArXiv version [26]. ■


Lemma 5.2. *Following the notation and assumptions of Theorem 5.1, every non-root nodal polynomial $f_{i,\zeta}$ of $\mathcal{T}_{p,k}(f)$ with $\zeta \in (\mathbb{Z}/(p^i))^*$ satisfies $\deg \tilde{f}_{i,\zeta} \leq 4$, $\deg \tilde{f}_{i,\zeta} \leq 3$, or $\deg \tilde{f}_{i,\zeta} \leq 2$, according as $p = 2$, $p = 3$, or $p \geq 5$.*


Example 5.3. *Recalling Example 2.8, which had*

$$f(x) = x^{10} - 10x + 738,$$

observe that $\mathcal{T}_{3,7}(f)$ is a chain of length 2. In particular, $\tilde{f}_{1,1}(x) = x^2(x-1)$, 0 is a degenerate root of $\tilde{f}_{1,1}$, and $s(\tilde{f}_{1,1}, 0) = 2$. We can then easily calculate that $\tilde{f}_{2,1}(x) = 2(x-1)(x-2) \pmod{3}$.

There are a total of 4 non-degenerate roots in \mathbb{F}_3 for the nodal polynomials: 1 for $\tilde{f}_{0,0}$, 1 for $\tilde{f}_{1,1}$, and 2 for $\tilde{f}_{2,1}$. These non-degenerate roots in \mathbb{F}_3 then lift to the following roots of f in \mathbb{Z}_3 : $0 + O(3^1)$, $1 + 1 \cdot 3 + O(3^2)$, $1 + 0 \cdot 3 + 1 \cdot 3^2 + O(3^3)$, and $1 + 0 \cdot 3 + 2 \cdot 3^2 + O(3^3)$. A quick calculation via Maple's `rootp` command tells us that these are all the 3-adic rational roots of f . ◊

Example 5.4. *One can check that for $f(x) := x^{10} + 11x^2 - 12$, the tree $\mathcal{T}_{2,8}(f)$ is isomorphic to . In particular, this f has exactly 6 roots in \mathbb{Q}_2^* : $\tilde{f}_{2,2} = \tilde{f}_{2,1} = \tilde{f}_{2,3} = x^2 + x$ and each of these (terminal) nodal polynomials has exactly 2 non-degenerate roots in \mathbb{F}_2 . Remembering the earlier digits encoded in $\mathcal{T}_{2,8}(f)$, these 6 roots then each lift to a unique root of f in \mathbb{Z}_2 . Note that $\tilde{f}_{1,1}(x) = x^4 + x^2$ has degree 4. ◊*

Example 5.5. *Composing Example 2.8 with x^2 , let us take $f(x) := x^{20} - 10x^2 + 738$. One then sees that the tree $\mathcal{T}_{3,7}(f)$ is isomorphic to . In particular, this f has exactly 8 roots in \mathbb{Q}_3^* , each arising as a Hensel lift of a non-degenerate root in \mathbb{F}_3 of some nodal polynomial:*

$\tilde{f}_{1,0}$, $\tilde{f}_{1,1}$, $\tilde{f}_{2,1}$, $\tilde{f}_{1,2}$, and $\tilde{f}_{2,8}$ respectively contribute 2, 1, 2, 1, and 2 roots. Note that $\tilde{f}_{1,2}(x) = x^3 + 2x^2 + x$ has degree 3. \diamond

To prove Lemma 5.2 we will need a powerful result of Lenstra [19] on the Newton polygons of shifted sparse polynomials. First, let us define $d_m(r)$ to be the least common multiple of all integers that can be written as the product of at most m pairwise distinct positive integers that are at most r , and set $d_m(r) := 1$ if $mr = 0$.

Theorem 5.6. [19, Sec. 3] *Suppose $f \in \mathbb{Q}[x]$ is a t -nomial, $g(x) = f(1 + px)$, and r is the largest nonnegative integer such that $r - \text{ord}_p d_{t-1}(r) \leq \max_{0 \leq j \leq t-1} \{j - \text{ord}_p(j!)\}$. Then any lower edge of $\text{Newt}_p(g)$ with inner normal $(v, 1)$ with $v \geq 1$ lies in the strip $[0, r] \times \mathbb{R}$. ■*

We point out that the vector of parameters (t, r, v) from our statement above would be $(k+1, m, v(x-1))$ in the notation of [19], and the parameter r there is set to 1 in our application here.

Proof of Lemma 5.2: First note that replacing x by cx , for any $c \in \{1, \dots, p-1\}$, preserves the number of roots of f in \mathbb{Z}_p and (up to relabelling the ζ in the subscripts of the $f_{i,\zeta}$) the tree $\mathcal{T}_{p,k}(f)$. So to study f_{1,ζ_0} with $\zeta_0 \in \{1, \dots, p-1\}$, it suffices to study $f_{1,1}$.

Note that the lower hull of any Newton polygon can be identified with a piecewise linear convex function on an interval. In particular, $f_{1,1}(x) = p^{-s(f,1)} f(1 + px)$ and thus the lower hull of $\text{Newt}_p(f_{1,1})$ can be identified with the sum of the lower hull of $\text{Newt}_p(f(1+x))$ and the function $x - s(f,1)$. Note also that by the definition of Newt_p , the minimal y -coordinate of a point of $\text{Newt}_p(f(1+px))$ is exactly $s(f,1)$. Theorem 5.6 then tells us that all lower edges of $\text{Newt}_p(f_{1,1})$ of non-positive slope lie in the strip $[0, r] \times \mathbb{R}$, where r is the largest nonnegative integer such that

$$(\star) \quad r - \text{ord}_p d_2(r) \leq \varepsilon_p$$

where $\varepsilon_2 = 1$ and $\varepsilon_p = 2$ for all $p \geq 3$. In particular, the definition of $\text{Newt}_p(f_{1,1})$ tells us that p divides the coefficient of x^j in $f_{1,1}$ for all $j \geq r+1$ and thus $\deg \tilde{f}_{1,1} \leq r$. By Lemma 2.13, all other non-root nodal polynomials $f_{i,\zeta}$ with $\zeta \neq 0 \pmod p$ satisfy $\deg \tilde{f}_{i,\zeta} \leq \deg \tilde{f}_{1,1}$. So it suffices to prove that r satisfies the stated bounds of our lemma. This is easily verified by first observing that $d_2(0) = d_2(1) = 1$ and $d_2(2) = 2$. So Inequality (\star) certainly holds for $r \in \{0, 1, 2\}$, regardless of p . Observing that $d_2(3) = 6$ and $d_2(4) = 24$, we then see that Inequality (\star) holds at $r = 4$ (resp. $r = 3$) when $p = 2$ (resp. $p = 3$).

So it is enough to show that:

- (i) $r - \text{ord}_2 d_2(r) \geq 2$ for $r \geq 5$,
- (ii) $r - \text{ord}_3 d_2(r) \geq 3$ for $r \geq 4$, and
- (iii) $r - \text{ord}_p d_2(r) \geq 3$ for $r \geq 3$ and $p \geq 5$.

From [19, Prop. 2.4], we have $\text{ord}_p d_2(r) \leq 2 \frac{\log r}{\log p}$. Note that, for any

fixed p , the quantity $r - 2 \frac{\log r}{\log p}$ is an increasing function of r for $r \geq \frac{2}{\log p}$. Furthermore, $\left\lceil 7 - 2 \frac{\log 7}{\log p} \right\rceil \geq 2$ for all $p \geq 2$ and $\left\lceil 5 - 2 \frac{\log 5}{\log p} \right\rceil \geq 3$ for all $p \geq 3$. Noting that $d_2(5) = 120$ and $d_2(6) = 360$, it is then easily checked that (i)–(iii) all hold. ■

Remark 5.7. *The proof of Lemma 2.15 is simply the variation of the proof above where we replace Inequality (\star) by $r - \text{ord}_p d_1(r) \leq 1$. \diamond*

Lemma 5.8. *For any trinomial $f \in \mathbb{Z}[x]$ with $\tilde{f}(0) \neq 0 \pmod p$, we can compute the mod p reductions of all the nodal polynomials of $\mathcal{T}_{p,k}(f)$ in time $p^{1+o(1)} k^{2+o(1)}$.*

Proof: By Lemma 2.13, $\mathcal{T}_{p,k}(f)$ has depth $\leq \left\lfloor \frac{k-1}{2} \right\rfloor$. By Lemma 5.2, all non-root nodal polynomials have mod p reduction of degree no greater than 4. Thus, the root of $\mathcal{T}_{p,k}(f)$ has at most $p-1$ children (since $\tilde{f}(0) \neq 0$), and any node at depth ≥ 2 has no more than 2 children (since a degree 4 polynomial has at most 2 degenerate roots). Lemma 2.13 also tells us that $\deg \tilde{f}_{i,\mu+p^{i-1}\zeta_{i-1}}$ is at most the multiplicity of ζ_{i-1} as a root of $\tilde{f}_{i-1,\mu}$. So any node v that has an ancestor at level ≥ 2 with 2 children can have no more than 1 child. Thus, there can be no more than $2(p-1)$ nodes at depth $i \geq 3$. It is then clear that $\mathcal{T}_{p,k}(f)$ has at most $1 + \left(2 \left\lfloor \frac{k-1}{2} \right\rfloor - 1\right)(p-1)$ nodes.

Now, note that the coefficient of x^i in the monomial term expansion of $c(\mu + px)^a \pmod{p^j}$ is simply $c \binom{a}{i} \mu^{a-i} p^i \pmod{p^j}$. Since f is a trinomial, and Lemma 2.13 tell us that $f_{i,\zeta}(x) = p^{-s} f(\mu + px) \pmod{p^j}$ for suitable (s, μ, j) , we can then clearly compute the coefficients of x^0, \dots, x^4 of any non-root nodal polynomial mod p^k using $O(\log p)$ multiplications and $O(1)$ additions. This takes time $O(k \log^2(p) \log(k \log p))$ via fast modular arithmetic [33], provided we use Harvey and van der Jooen's recent fast multiplication algorithm [14]. Summing over all non-root nodal polynomials, and noting that the cost of reduction mod p is negligible compared to the complexity of our earlier steps, we are done. ■

We can now outline the algorithm that proves Theorem 1.1.

Algorithm 5.9. (Solving Trinomial Equations Over \mathbb{Q}_p^*)

Input. A prime p and $c_1, c_2, c_3, a_2, a_3 \in \mathbb{Z} \setminus \{0\}$ with $|c_i| \leq H$ for all i and $1 \leq a_2 < a_3 =: d$.

Output. A true declaration that $f(x) := c_1 + c_2 x^{a_2} + c_3 x^{a_3}$ has no roots in \mathbb{Q}_p , or $z_1, \dots, z_m \in \mathbb{Q}$ with logarithmic height $O(p^2 \log^8(dH))$ such that m is the number of roots of f in \mathbb{Q}_p , z_j is an approximate root of f with associated true root $\zeta_j \in \mathbb{Q}_p$ for all j , and $\#\{\zeta_j\} = m$.

Description.

- 1: If $\text{ord}_p c_1 \neq \text{ord}_p c_2 \pmod{a_2}$ and $\text{ord}_p c_2 \neq \text{ord}_p c_3 \pmod{a_3 - a_2}$ then say "No roots in \mathbb{Q}_p !" and STOP.
- 2: Rescale and invert roots if necessary, so that we may assume $p \nmid c_1 c_2$ and $\text{ord}_p c_3 \geq 0$.
- 3: Compute the mod p reductions of all the nodal polynomials of $\mathcal{T}_{p,k}(f)$, for $k := 2D + 1$ where $D \geq \max_{\zeta \in \mathbb{Z}_p : \text{ord}_p \zeta = 0} \text{ord}_p f'(\zeta)$.
- 4: Use Hensel Lifting to find the first $2D + 1$ base- p digits of all the non-degenerate roots of f in \mathbb{Z}_p of valuation 0.
- 5: Via Algorithm 2.17 (or its $p=2$ version) find the first $O(\log(dH))$ base- p digits of all the degenerate roots of f .
- 6: If $p|c_3$ then rescale and invert roots to compute approximants for the remaining roots of f in \mathbb{Q}_p by computing roots of valuation 0 for a rescaled version of f with coefficients reversed.

Proof of Theorem 1.1: First note that 0 can not be a root since $f(0) \neq 0$ by assumption. So we can focus on roots in \mathbb{Q}_p^* . The height bound for our approximate roots from Assertion (1) follows directly from Step 3. Assertion (2) follows easily from Theorem 5.1: Steps 3 and 4 (which use Hensel's Lemma) imply a decay rate of $O(p^{-(2D+2^i)})$ for the p -adic distance of the i th Newton iterate to a true root. So the decay rate is no worse than $O((p^{-1/(2D+1)})^{2^i})$, and thus Assertion (2) holds with $\mu = p^{1/O(p^2 \log^8(dH))}$.

Assertion (3) on correctly counting the roots of f in \mathbb{Q}_p follows immediately from Steps 3–5.

So all that remains is to prove correctness (including elaborating Step 5) and to do a sufficiently good complexity analysis.

Correctness: Thanks to Theorem 2.3, Step 1 merely guarantees that f has roots of integral valuation, which is a necessary condition for their to be roots in \mathbb{Q}_p . Step 2 merely involves simple substitutions that only negligibly affect the heights of the coefficients, similar to the binomial case.⁷ Steps 3 and 4 correctly count the number of non-degenerate roots of f in \mathbb{Z}_p of valuation 0, thanks to Lemma 2.14.

For Step 5, since 0 is not a root, we can rearrange the equations $f(\zeta) = \zeta f'(\zeta) = 0$ to obtain that $\zeta \in \mathbb{Q}_p^*$ is a degenerate root of f if and only if $[c_1, c_2\zeta^{a_2}, c_3\zeta^{a_3}]^T$ is a right null-vector for $B := \begin{bmatrix} 1 & 1 & 1 \\ 0 & a_2 & a_3 \end{bmatrix}$. Since $[a_3 - a_2, -a_3, a_2]^T$ generates the right null-space of B we must have $(a_3 - a_2)c_2\zeta^{a_2} = -c_1a_3$ and $-a_3c_3\zeta^{a_3 - a_2} = c_2a_2$. Via an application of the Extended Euclidean Algorithm, we can find $R, S \in \mathbb{Z}$ with $Ra_2 + S(a_3 - a_2) = \gcd(a_2, a_3)$ and the logarithmic heights of R and S of order $O(\log d)$. So by multiplying and dividing suitable powers of our binomial equations, we get that ζ must satisfy the single equation $((a_3 - a_2)c_2)^R (-a_3c_3)^S \zeta^{\gcd(a_2, a_3)} = (-c_1a_3)^R (c_2a_2)^S$. The latter equation can be solved easily, within our overall time bound, via Algorithm 2.17. Note in particular that while the coefficient heights look much larger, any root ζ ultimately satisfies the *original* pair of binomials, thus implying ζ must have low logarithmic height.

Step 6 merely takes care of the remaining roots, at negligible affect to the coefficient heights.

Note that we do need to renormalize the roots at the end, due to the various rescalings, but this adds a negligible summand of $O(\log H)$ to the logarithmic heights of the roots. So we are done. ■

Complexity Analysis: Steps 3–4 dominate the overall complexity: Theorem 5.1 tells us that we can take $D = O(p^2 \log^8(dH))$, and thus Lemma 5.8 (combined with the known upper bounds on the number of p -adic rational roots of a trinomial [4, 19]) implies that the complexity of Steps 3–4 is no worse than

$$O\left(p^5 \log^3(p) \log^{16}(dH) \log(p \log(dH))\right),$$

assuming we employ brute-force search to find the roots in \mathbb{F}_p of the mod p reductions of the nodal polynomials. ■

6 ACKNOWLEDGEMENTS

We thank Erich Bach and Bjorn Poonen for informative discussions on Hensel’s Lemma. We also thank the anonymous referees for helpful suggestions that improved our paper.

REFERENCES

- [1] Martín Avendaño, Ashraf Ibrahim, J. Maurice Rojas, , and Korben Rusek. Faster p -adic feasibility for certain multivariate sparse polynomials. *Journal of Symbolic Computation*, 47(4):454–479, 2012.
- [2] Martín Avendaño, Roman Kogan, Mounir Nisse, and J. Maurice Rojas. Metric estimates and membership complexity for Archimedean amoebae and tropical hypersurfaces. *Journal of Complexity*, 46:45–65, 2018.
- [3] Martín Avendaño and Jorge Martín-Morales. Bivariate trinomials over finite fields. *Houston Journal of Mathematics*, to appear, 2021.
- [4] Martín Avendaño and Teresa Krick. Sharp bounds for the number of roots of univariate fewnomials. *Journal of Number Theory*, 131(7):1209 – 1228, 2011.
- [5] Eric Bach and Jeffrey Shallit. *Algorithmic number theory, volume 1: efficient algorithms*. MIT Press, Cambridge, Massachusetts, 1996.

- [6] A. Baker. Logarithmic forms and the abc -conjecture. In *Number theory (Eger, 1996)*, pages 37–44. de Gruyter, Berlin, 1998.
- [7] Jens-Dietrich Bauch, Enric Nart, and Hayden D. Stainsby. Complexity of OM factorizations of polynomials over local fields. *LMS Journal of Computation and Mathematics*, 16:139–171, 2013.
- [8] Jérémy Berthomieu, Grégoire Lecerf, and Guillaume Quintin. Polynomial root finding over local rings and application to error correcting codes. *Appl. Algebra Eng. Commun. Comput.*, 24:413–443, 2013.
- [9] J. M. Borwein and P. B. Borwein. On the complexity of familiar functions and numbers. *SIAM Rev.*, 30(4):589–601, 1988.
- [10] David G. Cantor and Daniel M. Gordon. Factoring polynomials over p -adic fields. In Wieb Bosma, editor, *Algorithmic Number Theory*, pages 185–208, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
- [11] Keith Conrad. Notes on Hensel’s Lemma. *Downloadable from* <https://kconrad.math.uconn.edu/blurbs/gradnumthy/hensel.pdf>, 2021.
- [12] Ashish Dwivedi, Rajat Mittal, and Nitin Saxena. Counting basic-irreducible factors mod p^k in deterministic poly-time and p -adic applications. In *34th Computational Complexity Conference*, volume 137 of *LIPICs. Leibniz Int. Proc. Inform.*, pages Art. No. 15, 29. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019.
- [13] Jordi Guàrdia, Enric Nart, and Sebastian Pauli. Single-factor lifting and factorization of polynomials over local fields. *Journal of Symbolic Computation*, 47(11):1318 – 1346, 2012.
- [14] David Harvey and Joris van der Hoeven. Integer multiplication in time $O(n \log n)$. *HAL preprint*, <https://hal.archives-ouvertes.fr/hal-02070778>, 2019.
- [15] Loo-Keng Hua and H. S. Vandiver. On the number of solutions of some trinomial equations in a finite field. *Proc. Nat. Acad. Sci. U.S.A.*, 35:477–481, 1949.
- [16] Pascal Koiran. Root separation for trinomials. *J. Symbolic Comput.*, 95:151–161, 2019.
- [17] Pascal Koiran, Natacha Portier, and Sébastien Tavenas. A Wronskian approach to the real r -conjecture. *J. Symbolic Comput.*, 68(part 2):195–214, 2015.
- [18] Leann Kopp, Natalie Randall, J. Maurice Rojas, and Yuyu Zhu. Randomized Polynomial-Time Root Counting in Prime Power Rings. *Mathematics of Computation*, 89(321):373–385, January 2020.
- [19] Hendrik W Lenstra. On the factorization of lacunary polynomials. *Number Theory in Progress*, 1:277–291, 1999.
- [20] Kurt Mahler. An inequality for the discriminant of a polynomial. *The Michigan Mathematical Journal*, 11(3):257–262, 1964.
- [21] Maurice Mignotte. On the distance between the roots of a polynomial. *Appl. Algebra Eng. Commun. Comput.*, 6:327–332, 11 1995.
- [22] Alexandre Ostrowski. Recherches sur la méthode de Graeffe et les zéros des polynômes et des séries de Laurent. *Acta Math.*, 72:99–155, 1940.
- [23] Bjorn Poonen. Zeros of sparse polynomials over local fields of characteristic p . *Math. Res. Lett.*, 5(3):273–279, 1998.
- [24] Bjorn Poonen. Using zeta functions to factor polynomials over finite fields. In *Arithmetic geometry: computation and applications*, volume 722 of *Contemp. Math.*, pages 141–147. Amer. Math. Soc., Providence, RI, 2019.
- [25] Alain M. Robert. *A Course in p -adic Analysis*. Springer-Verlag New York, 2000.
- [26] J. Maurice Rojas and Yuyu Zhu. A complexity chasm for solving sparse polynomial equations over p -adic fields (extended version with appendix). *ArXiv*, <https://arxiv.org/abs/2003.00314>, 2021.
- [27] Michael Sagraloff. A near-optimal algorithm for computing real roots of sparse polynomials. In *ISSAC 2014 (39th International Symposium on Symbolic and Algebraic Computation)*, pages 359–366, 2014.
- [28] W. H. Schikhof. *Ultrametric calculus*, volume 4 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2006. An introduction to p -adic analysis, Reprint of the 1984 original [MR0791759].
- [29] J.-P. Serre. *A course in arithmetic*. Springer-Verlag, New York-Heidelberg, 1973. Translated from the French, Graduate Texts in Mathematics, No. 7.
- [30] Igor Shparlinski. On finding primitive roots in finite fields. *Theoret. Comput. Sci.*, 157(2):273–275, 1996.
- [31] Steve Smale. Newton’s method estimates from data at one point. In *The merging of disciplines: new directions in pure, applied, and computational mathematics (Laramie, Wyo., 1985)*, pages 185–196. Springer, New York, 1986.
- [32] H. W. Turnbull, editor. *The correspondence of Isaac Newton, Vol. II: 1676–1687*. Cambridge University Press, New York, 1960. Published for the Royal Society.
- [33] Joachim von zur Gathen and Jürgen Gerhard. *Modern computer algebra*. Cambridge University Press, Cambridge, third edition, 2013.
- [34] Edwin Weiss. *Algebraic number theory*. International series in pure and applied mathematics. McGraw-Hill, 1963.
- [35] Kunrui Yu. Linear forms in p -adic logarithms III. *Compositio Mathematica*, 3(241-276), 1994.
- [36] Yuyu Zhu. *Trees, Point Counting Beyond Fields, and Root Separation*. PhD thesis, Texas A&M doctoral dissertation, TAMU 3368, College Station, TX 77843-3368, 5 2020.

⁷In particular, we may restrict to trees $\mathcal{T}_{p,k}(f)$ where $\deg \tilde{f}_{i,\zeta} \leq 4$ for all $i \geq 1$.