



Computational Section

Root repulsion and faster solving for very sparse polynomials over p -adic fields



J. Maurice Rojas*, Yuyu Zhu

Texas A&M University, TAMU 3368, College Station, TX 77843-3368,
United States of America

ARTICLE INFO

Article history:

Received 22 July 2021

Received in revised form 27

November 2021

Accepted 11 January 2022

Available online 27 April 2022

Communicated by L. Smajlovic

Keywords:

p -adic

Hensel

Newton

Iteration

Trinomial

Approximate

Root counting

ABSTRACT

Text. For any fixed field $K \in \{\mathbb{Q}_2, \mathbb{Q}_3, \mathbb{Q}_5, \dots\}$, we prove that all univariate polynomials f with exactly 3 (resp. 2) monomial terms, degree d , and all coefficients in $\{\pm 1, \dots, \pm H\}$, can be solved over K within deterministic time $\log^{4+o(1)}(dH) \log^3 d$ (resp. $\log^{2+o(1)}(dH)$) in the classical Turing model: Our underlying algorithm correctly counts the number of roots of f in K , and for each such root generates an approximation in \mathbb{Q} with logarithmic height $O(\log^2(dH) \log d)$ that converges at a rate of $O\left((1/p)^{2^i}\right)$ after i steps of Newton iteration. We also prove significant speed-ups in certain settings, a minimal spacing bound of $p^{-O(p \log_p^2(dH) \log d)}$ for distinct roots in \mathbb{C}_p , and even stronger root repulsion when there are nonzero degenerate roots in \mathbb{C}_p : p -adic distance $p^{-O(\log_p(dH))}$. On the other hand, we prove that there is an explicit family of tetranomials with distinct nonzero roots in \mathbb{Z}_p indistinguishable in their first $\Omega(d \log_p H)$ most significant base- p digits. So speed-ups for t -nomials with $t \geq 4$ will require evasion or amortization of such worst-case instances.

* Partially supported by NSF grants CCF-1900881 and CCF-1409020. A much shorter, preliminary version of this work appeared in the proceedings of the conference ISSAC 2021 (July 19–23, virtual event) [50].

* Corresponding author.

E-mail addresses: jmauricerojas@gmail.com (J.M. Rojas), yuyu.zhu1213@gmail.com (Y. Zhu).

Video. For a video summary of this paper, please visit
<https://youtu.be/npfdxLhk04MY>.

© 2022 Elsevier Inc. All rights reserved.

Contents

1. Introduction	656
1.1. Dependence on p	658
1.2. Repulsion, and the separation chasm at four terms	659
1.3. Previous complexity and sparsity results	661
2. Background	661
2.1. Newton polygons and Newton iteration: Archimedean and non-Archimedean	662
2.2. Separating roots of binomials	663
2.3. Characterizing roots of binomials over \mathbb{Q}_p^*	664
2.4. Bit complexity basics and counting roots of binomials	665
2.5. Trees and roots in $\mathbb{Z}/(p^k)$ and \mathbb{Z}_p	666
2.6. Trees and extracting digits of radicals	670
3. Proving Theorem 1.6: trinomial roots never get too close	673
4. Proving Theorem 1.5: tetranomial roots can get too close	676
4.1. The case of prime p	676
4.2. The case $p = \infty$	677
5. Valuation bounds from discriminants and repulsion from degeneracy	677
5.1. Completing the proof of Theorem 1.6: root spacing in the face of degeneracy	681
6. Solving trinomials over \mathbb{Q}_p	682
6.1. Trees and trinomials	682
6.2. Building trees efficiently	686
6.3. The algorithm that proves Theorem 1.1	691
6.4. “Typical” exponents, Las Vegas, and a combined speed-up	694
Acknowledgments	697
References	697

1. Introduction

Solving polynomial equations over the p -adic rational numbers \mathbb{Q}_p underlies many important computational questions in number theory (see, e.g., [25,9,23,51]) and is close to applications in coding theory (see, e.g., [11]). Furthermore, the complexity of solving *structured* equations — 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., [38]). So we will study how sparsity affects the complexity of separating and approximating roots in \mathbb{Q}_p . Unless stated otherwise, all O -constants and Ω -constants are effective and absolute.

Recall that thanks to 17th century work of Descartes, and 20th century work of Lenstra [40] and Poonen [46], it is known that univariate polynomials with exactly t monomial terms have at most $t^{O_K(1)}$ roots in a local field K only when K is \mathbb{R} or a finite algebraic extension of \mathbb{Q}_p for some prime $p \in \mathbb{N}$. (For instance, \mathbb{C} is ruled out because $x^d - 1$ has just 2 monomial terms but d distinct complex roots.) We’ll use $|\cdot|_p$ (resp. $|\cdot|$) for the absolute value on the p -adic complex numbers \mathbb{C}_p 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 $i \geq 0$.

Our first main result is that we can efficiently count the roots of univariate 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. *For any prime p and a trinomial $f \in \mathbb{Z}[x]$ with degree d and all its coefficients in $\{\pm 1, \dots, \pm H\}$, we can find in deterministic time $p^{3+o(1)} \log^{4+o(1)}(dH) \log_p^3 d$ a set $\{\frac{\alpha_1}{\beta_1}, \dots, \frac{\alpha_m}{\beta_m}\} \subset \mathbb{Q}$ of cardinality $m = m(p, f)$ such that:*

- (1) *For all j we have $\alpha_j \neq 0 \implies \log|\alpha_j|, \log|\beta_j| = O(p \log_p^2(dH) \log d)$.*
- (2) *$z_0 := \alpha_j/\beta_j \implies f$ has a root $\zeta_j \in \mathbb{Q}_p$ with sequence of Newton iterates satisfying $|z_i - \zeta_j|_p \leq (1/p)^{2^i} |z_0 - \zeta_j|_p$ for all $i, j \geq 1$.*
- (3) *$m = \#\{\zeta_1, \dots, \zeta_m\}$ and m is exactly the number of roots of f in \mathbb{Q}_p .*

We prove Theorem 1.1 in Section 6.3 via Algorithm 6.12 there. The main idea behind Algorithm 6.12 is conceptually simple: solving for enough of the most significant base- p digits of the roots to guarantee rapid convergence of Newton/Hensel Iteration. However, proving that this can be done efficiently hinges on recent root counts from arithmetic feynomial theory [40,6,12,36] and some delicate root spacing estimates (Theorem 1.6 and Sections 3 and 5 below) that form the technical heart of this paper.

The dependence on p in our complexity bound can be lowered significantly in certain natural settings, e.g., restricting to roots of the form $p^j + O(p^{j+1})$, or making mild assumptions on the gcd of the exponents, or assuming the presence of nonzero degenerate roots in \mathbb{C}_p : See Corollaries 1.4, 1.7, and 6.16 below. An analogue of Theorem 1.1 also holds for $K = \mathbb{R}$ and will be presented in a sequel to this paper [13]. We call a $z_0 \in \mathbb{Q}_p$ satisfying the convergence condition from Theorem 1.1 *an approximate root of f (in the sense of Smale¹)*, with *associated true root ζ* . 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 univariatetrinomial equations, over \mathbb{Q}_p for any *fixed* prime p , in deterministic time polynomial in the input size. ◇

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., [14]). We review the case $(K, t) = (\mathbb{Q}_p, 2)$ with p prime in Corollary 2.8 and Theorem 2.21 of Section 2 below. ◇

¹ This terminology has only been applied over \mathbb{C} so far [57], so we take the opportunity here to extend it to the p -adic rationals. Note that we do not restrict ζ to be non-degenerate.

Despite much work on factoring univariate polynomials over \mathbb{Q}_p (see, e.g., [40,16,30, 10,11]), 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 greatly improves an earlier complexity bound (membership in **NP**, for detecting roots in \mathbb{Q}_p) from [3]. We'll see in Sections 5 and 6 how our speed-up depends on p -adic Diophantine approximation [64,65]. Another key new ingredient in proving Theorem 1.1 is an efficient encoding of roots in $\mathbb{Z}/(p^k)$ from [27,39], with important precursors in [61,11].

1.1. Dependence on p

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 might not have such tame dependence on p . There are 3 barriers (B1–B3 below) we must overcome before achieving such a speed-up:

B1. *Whereas a binomial has at most 3 roots in \mathbb{R} (e.g., $x^3 - x$), a binomial can have as many as $\max\{p, 3\}$ roots in \mathbb{Q}_p (e.g., $x^{\max\{p, 3\}} - x$). Furthermore, trinomials have at most 5, 7, 9, or $3p - 2$ roots in K , according as K is \mathbb{R} , \mathbb{Q}_2 [40], \mathbb{Q}_3 [66], or \mathbb{Q}_p with $p \geq 5$ [6,45], and each bound is sharp. \diamond*

One might think that B1 is the nail in the coffin for dependence $(\log p)^{O(1)}$. However, Hensel's Lemma, and a tree from Section 2.5 below that encodes roots in \mathbb{Z}_p (see also [27]), reveal that the roots of a trinomial in \mathbb{Q}_p can in fact be encoded by a data structure of potentially much smaller size than a naive list of size $\Omega(p)$. (In essence, this means using an explicit collection binomials to encode a union of cosets of \mathbb{F}_p^* .) This harkens back to an intriguing open problem from arithmetic fewnomial theory [20]: Is the zero set of a trinomial in $\mathbb{F}_p[x]$ always expressible as a union of $O(\log p)$ cosets of subgroups of \mathbb{F}_p^* ? Currently, the best bound is $O(\sqrt{p})$ [36] and forms a key ingredient in proving one of our speed-ups: Corollary 1.4 below.

Observe now that the most natural p -adic analogue of a positive real number is a p -adic rational number *with most significant digit 1*, i.e., a number of the form $p^j + O(p^{j+1})$. Restricting to such roots then cuts the root cardinality bounds of B1 down to 2, 6, 4, and 3 (respectively over \mathbb{R} , \mathbb{Q}_2 , \mathbb{Q}_3 , or \mathbb{Q}_p with $p \geq 5$), and yields a significant speed-up. Alternatively, rather than restricting digits of roots, one can observe that trinomials over \mathbb{Z} with many roots in \mathbb{Q}_p have very particular exponents (see, e.g., [6]). This enables another significant speed-up to our main algorithm for “most” choices of exponents. We unite these two speed-ups as follows:

Corollary 1.4. *Following the notation of Theorem 1.1, we can reduce its deterministic time complexity bound by a factor of p in either of following settings: (1) we only seek roots of the form $p^j + O(p^{j+1})$, or (2) we assume that the exponents are $\{0, a_2, a_3\}$ with $\gcd(a_2 a_3 (a_3 - a_2), (p-1)p) \leq 2$. In either case, the stated worst-case height bounds for the approximate roots remain the same.*

We prove Corollary 1.4 in Section 6.4, and leave average-case speed-ups (where one averages over *coefficients*) for future work. It follows from our framework that the speed-ups from Corollary 1.4 continue to hold (modulo a multiple of $C^{O(1)}$) under softer assumptions like (a) restricting to roots with most significant digit in some cardinality C subset of $\{1, \dots, p-1\}$ or (b) assuming $\gcd(a_2 a_3 (a_3 - a_2), (p-1)p) \leq C$. So our assumptions in Corollary 1.4 are more restrictive merely for the sake of simplifying our exposition.

Another barrier to more efficient dependence on p is *explicitly extracting* points from the cosets forming the zero set of a trinomial. For instance:

B2. *Approximating square-roots of p -adic integers not divisible by p , within accuracy 1, implies 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 time $\log^{2+o(1)} p$ (see, e.g., [55, 7, 47]). Furthermore, it remains unknown how to find just a single d th root of a d th power in \mathbb{F}_p^* in time $(\log(p) + \log d)^{O(1)}$, even if randomness is allowed (see, e.g., [1, 18, 21]).* \diamond

We are then led naturally to yet another barrier to efficient dependence on p :

B3. *Even if one only wants to approximate a single root in \mathbb{Q}_p of a trinomial, the minimal currently provable initial accuracy needed to make Newton iteration converge quickly appears to have quasi-linear dependence on p .* \diamond

In particular, our key valuation bounds (see Section 5) currently hinge on estimates for *linear forms in p -adic logarithms* [8, 64, 65], and further improvements to the latter estimates appear to be unknown and difficult.

1.2. Repulsion, and the separation chasm at four terms

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

Theorem 1.5. *Consider the family of tetranomials*

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

with $j \in \mathbb{N}$, $j \geq 3$, $\varepsilon \in \mathbb{Q}$ nonzero, and $d \in \{4, \dots, \lfloor e^h \rfloor\}$ even. Let $H := \max\{\varepsilon^{\pm 2j}\}$. Then $f_{d,\varepsilon,j}$ has distinct nonzero roots ζ_1, ζ_2 in the open unit disk of K (centered at 0) with $|\log|\zeta_1 - \zeta_2||_p = \Omega(d \log H)$ or $|\log|\zeta_1 - \zeta_2|| = \Omega(d \log H)$, according as $(K, \varepsilon) = (\mathbb{Q}_p, p)$ or $(K, \varepsilon) = (\mathbb{R}, 1/2)$. In particular, while the coefficients of $p^{2j} f_{d,p,j}$ all lie in \mathbb{Z} and have $O(\log_p H)$ base- p digits, we need $\Omega(d \log_p H)$ many base- p digits to distinguish the nonzero roots of f in \mathbb{Z}_p .

We prove Theorem 1.5 in Section 4, where we will also see in Remark 4.1 that the basin of attraction for a root of $f_{d,p,j}$ in \mathbb{Q}_p (under the Newton endomorphism $N_{f_{d,p,j}}$) can be exponentially small in $\log d$ as well. The special case $K = \mathbb{R}$ of Theorem 1.5 was derived earlier (in different notation) by Mignotte [42]. (See also [52].) The cases

$K = \mathbb{Q}_p$ with p prime appear to be new, and our proof unifies the Archimedean and non-Archimedean cases via tropical geometry [4]. Approximating roots in \mathbb{Q}_p in average-case time sub-linear in d for tetranomials (where one averages over the coefficients but fixes the exponents) is thus an intriguing open problem.

Mignotte used the tetranomial $f_{d,1/2,j}$ in [42] to show that an earlier root separation bound of Mahler [41], 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 $d \geq 2$, all its coefficients lie in $\{\pm 1, \dots, \pm H\}$, and f is irreducible in $\mathbb{Z}[x]$. Let $\zeta_1, \zeta_2 \in \mathbb{C}$ be distinct roots of f . Then $|\zeta_1 - \zeta_2| > \frac{\sqrt{3}}{(d+1)^{d+\frac{1}{2}} H^{d-1}}$. In particular, $|\log |\zeta_1 - \zeta_2|| = O(d \log(dH))$. ■*

The very last statement is actually a small addendum, making use of the following classic fact: The complex roots of an f as above lie in an open disk, centered at the origin, of radius $2H$ (see, e.g., [48, Ch. 8] or Theorem 2.3 in Section 2.1 below). It is straightforward to prove an analogue of Mahler's bound, of the same asymptotic order for $|\log |\zeta_1 - \zeta_2||$, for roots in \mathbb{C}_p .

Our new algorithmic results are enabled by our third and final main result: Mahler's bound can be dramatically improved for the roots of trinomials in \mathbb{C}_p .

Theorem 1.6. *Suppose p is prime and $f \in \mathbb{Z}[x]$ has exactly 3 monomial terms, degree d , and all its coefficients lie in $\{\pm 1, \dots, \pm H\}$. Let $\zeta_1, \zeta_2 \in \mathbb{C}_p$ be distinct roots of f . Then $\log H \geq \log |\zeta_1 - \zeta_2|_p \geq -O(p \log^2(dH) \log_p d)$. Furthermore, if f has a nonzero degenerate root in \mathbb{C}_p , then the last lower bound can be sharpened to $-O(\log(dH))$.*

The proof of Theorem 1.6 is split across Sections 3 and 5.1, due to the degenerate case being more technically difficult. In particular, Theorem 1.6 provides a p -adic analogue of a separation bound of Koiran for complex roots of trinomials [37], and the proofs of the non-degenerate (a.k.a. square-free) cases over \mathbb{C}_p and \mathbb{C} share much in common. However, the sharper bounds for the degenerate cases involve radically different techniques over \mathbb{C}_p and \mathbb{C} : Over \mathbb{C} , there are refined analogues of Rolle's Theorem that incorporate the multiplicity of roots. Over \mathbb{C}_p , such a refinement is unavailable, so we resort to estimates on the valuation of discriminants of trinomials (see Lemma 5.4).

As to whether our root spacing bounds above are optimal, there are recent examples from [28] showing that $\log |\zeta_1 - \zeta_2|_p = -\Omega(\log \max\{d, H\})$ can occur. However, we are unaware of any examples exhibiting $\log |\zeta_1 - \zeta_2|_p = -\Omega(p^\varepsilon)$ for some $\varepsilon > 0$. Asymptotically optimal separation bounds, over both \mathbb{C}_p and \mathbb{C} , are already known for binomials and we review these bounds in Section 2.2.

The presence of degenerate roots appears to not only increase the repulsion of roots for trinomials but also speed up their approximation:

Corollary 1.7. *Following the notation of Theorem 1.1, if f has a nonzero degenerate root in \mathbb{C}_p , then we can find, in deterministic time $p^{1+o(1)} \left(p^{1/2} + \log^{3+o(1)}(dH) \right)$, or Las Vegas randomized time $p^{1+o(1)} \log^{3+o(1)}(dH)$, a set of approximate roots in the sense of Smale, each in \mathbb{Q} and with logarithmic height $O(\log(dH))$, with distinct associated true roots having union the zero set of f in \mathbb{Q}_p .*

We prove Corollary 1.7 in Remark 6.14 of Section 6.3 below. It is not yet clear whether significantly better bounds for root spacing and root approximation can hold in complete generality: the apparent improvements implied by the presence of degenerate roots could just be a side-effect of our underlying techniques. Curiously, a similar “repulsion from degeneracy” phenomenon also occurs in the (Archimedean) setting of roots in \mathbb{C} : see [37, Proof of Thm. 18].

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 [3]. On the other hand, detecting roots over \mathbb{Q}_p for n -variate $(n+1)$ -nomials is known to be doable in **NP** [3]. 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 [66].

On a related note, counting points on trinomial curves over 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 [33,62] and revisited via real quadratic optimization in [5].

2. Background

Definitive sources for p -adic arithmetic and analysis include [54,53,49]. For algorithmic complexity we note that [44,2] are outstanding references. Let us now collect some basic terminology:

- For any ring R we let R^* denote the multiplicatively invertible elements of R .
- The *logarithmic height* of a rational number a/b with $\gcd(a, b) = 1$ is simply $h(a/b) := \log \max\{|a|, |b|\}$, and we declare $h(0) := 0$.
- Over any algebraically closed field K , we define the *multiplicity of a root* $\zeta \in K$ of $f \in K[x]$ as the greatest m with $(x - \zeta)^m | f$. (We will usually take K to be \mathbb{C}_p or the algebraic closure $\overline{\mathbb{F}}_p$ of \mathbb{F}_p .)
- The *most significant (base- p) 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 $a_s \neq 0$.
- We denote the standard p -adic valuation on \mathbb{C}_p (normalized so that $\text{ord}_p p = 1$) by $\text{ord}_p : \mathbb{C}_p \longrightarrow \mathbb{Q}$.

Recall that the famous *Ultrametric Inequality* states that for any $\alpha, \beta \in \mathbb{C}_p$ we have $\text{ord}_p(\alpha \pm \beta) \geq \min\{\text{ord}_p \alpha, \text{ord}_p \beta\}$. (Equivalently: $|\alpha \pm \beta|_p \leq \max\{|\alpha|_p, |\beta|_p\}$.) We will frequently use (without further mention) this inequality, along with its natural implication $\text{ord}_p \alpha < \text{ord}_p \beta \implies \text{ord}_p(\alpha \pm \beta) = \text{ord}_p \alpha$. We also recall that the metrics $|\cdot|$ and $|\cdot|_p$ are respectively called *Archimedean* and *non-Archimedean* because as $n \rightarrow \infty$ we have $|n| \rightarrow \infty$, while the sequence $|n|_p$ remains inside the bounded set $\{1, 1/p, 1/p^2, \dots\}$.

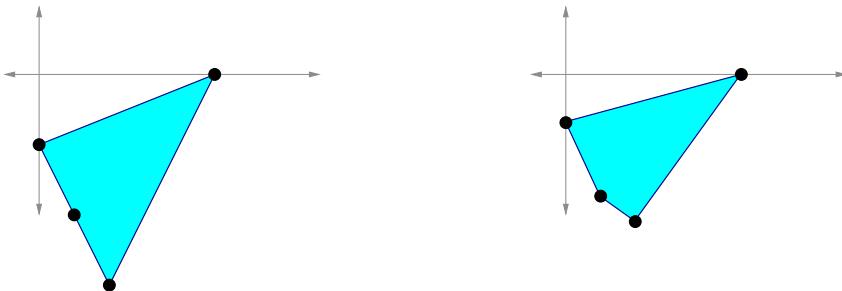
Let us also recall that a polynomial-time *Las Vegas randomized* algorithm is a polynomial-time algorithm that uses polynomially random bits in the input size, errs with probability at worst $1/2$, but always reports if it errs. Such an algorithm can be run k times to boost the success probability to at least $1 - \frac{1}{2^k}$, and this type of randomization is standard in many number-theoretic algorithms such as the fastest current algorithms for factoring polynomials over finite fields or primality checking (see, e.g., [34,19]). In our setting, errors (for a Las Vegas speed-up) consist of reporting too few roots in \mathbb{Q}_p , but such errors can be detected and reported at no extra cost.

2.1. Newton polygons and Newton iteration: Archimedean and non-Archimedean

The notion of Newton polygon goes back to 17th century work of Newton on Puiseux series solutions to polynomial equations [59, pp. 126–127]. We will need variants of this notion over \mathbb{C}_p and \mathbb{C} . (See, e.g., [63] for the p -adic case and [43,4] 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 $\lambda(E) := |u - r|$. ◇

Example 2.2. Consider $g_\varepsilon(x) := x^5 - \varepsilon^{-6}x^2 + 2\varepsilon^{-4}x - \varepsilon^{-2}$. We illustrate $\text{Newt}_p(g_p)$ (for p odd) and $\text{Newt}_\infty(g_{1/2})$ 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). ◇

Theorem 2.3. *Following the notation above, the number of roots of f in \mathbb{C}_p of valuation v (counting multiplicity) 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 number of roots $\zeta \in \mathbb{C}$ of f with $\log |\zeta| \in (v - \log 3, v + \log 3)$, counting multiplicity, is exactly $\lambda(E)$. ■*

The first portion of Theorem 2.3 goes back to work of Dumas around 1906 [26], while the second portion is an immediate consequence of [4, Thm. 1.5] (with an important precursor by Ostrowski around 1940 [43]). The set of slopes of the lower edges of $\text{Newt}_p(f)$ (or of $\text{Newt}_\infty(f)$) is an example of a *tropical variety* [4].

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., [22, 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 $f(\zeta') = 0 \pmod{p^{2\ell+2j}}$, $\text{ord}_p f'(\zeta') = \ell$, and $\zeta = \zeta' \pmod{p^{\ell+2j}}$. ■*

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 much less on the degree. For convenience, we will sometimes write $|\cdot|_\infty$ instead of $|\cdot|$ for the standard norm on \mathbb{C} . Rather than stating lower bounds on $|\zeta_1 - \zeta_2|_p$, we will instead state *upper* bounds on $|\log |\zeta_1 - \zeta_2|_p|$: the latter clearly includes *both* a lower and upper bound on $|\zeta_1 - \zeta_2|_p$. In summary, we have the following:

Proposition 2.4. *Suppose $f(x) := c_1 + c_2 x^d \in \mathbb{Z}[x]$, $d \geq 2$, $c_1 c_2 \neq 0$, and $|c_1|, |c_2| \leq H$. Then for any distinct roots $\zeta_1, \zeta_2 \in \mathbb{C}$ of f , we have $|\log |\zeta_1 - \zeta_2|| \leq \log(d) + \frac{1}{d} \log H$. Also, for any distinct roots $\zeta_1, \zeta_2 \in \mathbb{C}_p$ of f , we have that $|\log |\zeta_1 - \zeta_2|_p|$ is at most $\frac{1}{d} \log H$ or $\frac{\log p}{p-1} + \frac{1}{d} \log H$, according as $d > p^{\text{ord}_p d}$ or $d = p^{\text{ord}_p d} \geq p$.*

Put another way, if one fixes $H \geq 1$ and the prime p , and lets $d \rightarrow \infty$, then the minimal root distance tends to 0 in the Archimedean case. However, in the non-Archimedean case, the minimal root distance is never less than $\frac{1}{H^{1/d} p^{1/(p-1)}} (\geq \frac{1}{2H})$.

Proof of Proposition 2.4. The case $p = \infty$ follows from an estimate for the distance between the vertices of a regular d -gon. In particular, the minimal spacing between distinct complex roots can easily be expressed explicitly as $|c_1/c_2|^{1/d} \sqrt{2(1 - \cos \frac{2\pi}{d})}$, which is clearly bounded from below by $H^{-1/d} \sqrt{2(1 - \cos \frac{2\pi}{d})}$. From the elementary inequality $1 - \cos x \geq x^2 \left(\frac{1}{2} - \frac{\pi^2}{48}\right)$ we easily get $|\frac{1}{2} \log (1 - \cos \frac{2\pi}{d})| \leq \log(d) - \frac{1}{2} \log \left(4\pi^2 - \frac{\pi^2}{6}\right)$ for

all $d \geq 6$. Observing that $|\frac{1}{2} \log(1 - \cos \frac{2\pi}{d})| \leq \log 2$ for $d \in \{2, \dots, 5\}$ we get our stated bound via the Triangle Inequality applied to $|\log(H^{-1/d} \sqrt{2(1 - \cos \frac{2\pi}{d})})|$.

The case of prime p follows easily from the Ultrametric Inequality and classical facts on the spacing of p -adic roots of unity (see, e.g., [49, Cor. 1, Pg. 105, Sec. 4.3 & Thm. Pg. 107, Sec. 4.4]). In particular, when $\gcd(d, p-1) = 1$, distinct d th roots of unity in \mathbb{C}_p are all at unit distance. At the opposite extreme of $d = p^j$ for $j \geq 1$, the set of distances between distinct d th roots of unity is exactly $\left\{ \frac{1}{p^{1/(p-1)}}, \frac{1}{p^{1/(p^1(p-1))}}, \dots, \frac{1}{p^{1/(p^{j-1}(p-1))}} \right\}$. So the minimum distance is $1/p^{1/(p-1)}$ for d a non-trivial p th power. In complete generality, we see that there are distinct d th roots of unity at distance 1 if and only if d is divisible by a prime other than p . Observing that $|x^{1/d}|_p = p^{-\frac{1}{d} \operatorname{ord}_p x}$ and $\operatorname{ord}_p H \leq \log_p H$ for $x \in \mathbb{C}_p^*$ and $H \in \mathbb{N}$, we then see that $|\log|H^{\pm 1/d}|_p| \leq \frac{1}{d} \log H$ and our bound follows from the multiplicativity of norms. ■

2.3. Characterizing roots of binomials over \mathbb{Q}_p^*

Counting roots of binomials over \mathbb{Q}_p is more involved than counting their roots over \mathbb{R} , but is still quite efficiently doable. The first step is reducing the problem to $\mathbb{Z}/(p^k)$ for k linear in the bit-size of the degree of the binomial.

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 $\ell := \operatorname{ord}_p d$. 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:*

$$(1) \quad d \mid \operatorname{ord}_p(c_1/c_2) \quad \text{and}$$

$$(2) \quad \left(-\frac{c_1}{c_2} p^{\operatorname{ord}_p(c_2/c_1)} \right)^{p^{\ell}(p-1)/\gcd(d,p-1)} = 1 \pmod{p^{2\ell+1}}. \quad \blacksquare$$

Lemma 2.5 is classical and follows from basic group theory (the fact that the multiplicative group $(\mathbb{Z}/(p^k))^*$ is cyclic, of order $p^{k-1}(p-1)$, for p odd) and Hensel's Lemma.

Recall that the only roots of unity in \mathbb{Q}_2 are $\{\pm 1\}$ (see, e.g., [49]). The following lemma is then a simple consequence of the multiplicative group $(\mathbb{Z}/(2^k))^*$ being exactly the product $\{\pm 1\} \times \{1, 5, \dots, 5^{2^{k-3}} \pmod{2^k}\}$ (having cardinality 2^{k-1}) when $k \geq 3$ (see, e.g., [7, Thm. 5.6.2, pg. 109]), and Hensel's Lemma.

Lemma 2.6. *Suppose $f(x) := c_1 + c_2 x^d \in \mathbb{Z}[x]$ with $|c_1|, |c_2| \leq H$, and $c_1 c_2 \neq 0$. Then the number of roots of the binomial f in \mathbb{Q}_2 is either 0 or $\gcd(d, 2)$. In particular, if $\ell := \operatorname{ord}_2 d$ and $u := \operatorname{ord}_2(c_2/c_1)$, then f has roots in \mathbb{Q}_2 if and only if both of the following conditions hold: (1) $d \mid u$ and (2) either (i) d is odd or (ii) both $\frac{c_1}{c_2} 2^u = -1 \pmod{8}$ and $\left(-\frac{c_1}{c_2} 2^u \right)^{2^{\ell-1}} = 1 \pmod{2^{2\ell+1}}$. ■*

2.4. Bit complexity basics and counting roots of binomials

The following bit-complexity estimates for finite ring arithmetic will be fundamental for our main algorithmic results, and follow directly from the development of [60, Ch. 4 & 11] (particularly [60, Cor. 11.13, pg. 327]) assuming one uses the recent fast integer multiplication algorithm of Harvey and van der Hoeven [32]. See also [58] for an excellent exposition on most of the bounds below. We use $\log^* x$ to denote the minimal k such that k compositions of \log applied to x yield a real number ≤ 1 .

Theorem 2.7. *For any prime $p \in \mathbb{N}$ and $j, m, n \in \mathbb{N}$, we have the following bit-complexity bounds (in the Turing model) involving $A, a, b, c \in \mathbb{N}$ with $A, a, b \leq 2^n - 1$, $A \geq 2^{n-1}$, $c \leq 2^m - 1$ with $m = O(\log n)$, $r, s \in \{0, \dots, p^j - 1\}$ with $p \nmid r$, and $f, g \in \mathbb{F}_p[x]$ both having degree $\leq d$:*

Operation	Best Current O -bound (as of December 2021)
$a + b$	$O(n)$
$a \cdot b$	$O(n \log n)$
$a \bmod b$	$O(n \log n)$
$A \bmod c$	$O(nm)$
$r \cdot s \bmod p^j$	$O(j \log(p) \log(j \log p))$
$1/r \bmod p^j$	$O(j \log(p) \log^2(j \log p))$
$r^s \bmod p^j$	$O(j^2 \log^2(p) \log(j \log p))$
$f \cdot g$	$O(d \log(p) \log(d \log(p)) 4^{\log^*(d \log p)})$
$\gcd(f, g)$	$O(d \log(p) \log^2(d) \log(\log d) \log \log p)$

We note that the penultimate bound comes directly from [31]. The very last bound is actually a simple combination of the Half-gcd algorithm from [60, Thm. 11.1, Ch. 11] with the fast polynomial multiplication algorithm from [17], and can likely be improved slightly via the techniques of [31].

Corollary 2.8. *Following the notation of Lemmata 2.5 and 2.6, one can count exactly the number of roots of f in \mathbb{Q}_p in time $\log^{2+o(1)}(dH)$. Furthermore, for any root $\zeta \in \mathbb{Q}_p^*$ there is an $x_0 \in \mathbb{Z}/(p^{2\ell+1})$ 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 height of z_0 is $O(\log(pH^{1/d}))$.*

Proof: (Case of odd p). First note that $(\mathbb{Z}/p^{2\ell+1})^*$ is cyclic and Lemma 2.5 tells us that we can reduce deciding the feasibility of $c_1 + c_2 x^d = 0$ over \mathbb{Q}_p^* to checking $d \mid \text{ord}_p(c_1/c_2)$ and $(-c_1/c_2)^r \stackrel{?}{=} 1 \pmod{p^{2\ell+1}}$ with $r = p^\ell(p-1)/\gcd(d, p-1)$.

The p -adic valuation can be computed easily by bisection, ultimately resulting in $O(\log H)$ divisions involving integers with $O(\max\{\log p, \log H\}) = O(\log(pH))$

bits. Checking divisibility by d involves dividing an integer with $O(\log \log H)$ bits by an integer with $O(\log d)$ bits. By Theorem 2.7 these initial steps take time $O(\log(H) \log(pH) \log \log(pH)) + O(m \log m)$, where $m = \max\{\log \log H, \log d\}$. By Theorem 2.7, the r th power can be computed in time $O(\ell^2 \log^2(p) \log(\ell \log p))$. So our overall complexity bound is

$$O(\ell^2 \log^2(p) \log(\ell \log p) + \log(H) \log(pH) \log \log(pH) + \log(d) \log \log d).$$

Since $\ell \leq \log_p d$ our final bound becomes

$$O(\log^2(d) \log(\log d) + \log(H) \log(pH) \log \log(pH)).$$

A simple over-estimate then yields our stated complexity bound. The remainder of the lemma then follows easily from Hensel's Lemma and Proposition 2.4. ■

(Case of $p = 2$) The proof is almost identical to the odd p case, save that we use Lemma 2.6 in place of Lemma 2.5. In particular, the case $\ell = 0$ remains unchanged.

As for the case $\ell \geq 1$, the only change is an extra congruence condition (mod 8) to check whether $\frac{c_1}{c_2} 2^u$ is a square mod $2^{2\ell+1}$ (see, e.g., [7, Ex. 38, pg. 192]). However, this additional complexity is negligible compared to the other steps, so we are done. ■

2.5. Trees and roots in $\mathbb{Z}/(p^k)$ and \mathbb{Z}_p

Recall that for any field K , a root $\zeta \in K$ of f is *degenerate* if and only if $f(\zeta) = f'(\zeta) = 0$. 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 2.9. [39] For any $f \in \mathbb{Z}[x]$ let \tilde{f} denote the mod p reduction of f . Assume \tilde{f} is not identically 0. Then, for any degenerate root $\zeta_0 \in \{0, \dots, p-1\}$ of \tilde{f} , we then define $s(f, \zeta_0) := \min_{i \geq 0} \left\{ i + \text{ord}_p \frac{f^{(i)}(\zeta_0)}{i!} \right\}$. ◇

Example 2.10. 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} in \mathbb{F}_3 , and one can check that $s(f, 1) = 4$. Note that $s(f, 1)$ here is strictly less than 9, which is the multiplicity of the factor $x-1$ of \tilde{f} . ◇

The quantity $s(f, \zeta_0)$, combined with our definition below, will enable us to reduce finding the base- p digits of a root of f in $\mathbb{Z}/(p^k)$ to solving several simpler equations over $\mathbb{Z}/(p)$.

Definition 2.11. [39] Fixing $k \in \mathbb{N}$, for $i \geq 1$, let us inductively define a set $\mathcal{T}_{p,k}(f)$ 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

$(f_{i-1,\mu}, k_{i-1,\mu}) \in \mathcal{T}_{p,k}(f)$, and any degenerate root $\zeta_{i-1} \in \mathbb{F}_p$ 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 $\zeta := \mu + \zeta_{i-1}p^{i-1}$, $k_{i,\zeta} := k_{i-1,\mu} - s_{i-1}$, $f_{i,\zeta}(x) := p^{-s(f_{i-1,\mu}, \zeta_{i-1})} f_{i-1,\mu}(\zeta_{i-1} + px) \pmod{p^{k_{i,\zeta}}}$, and then include append $(f_{i,\zeta}, k_{i,\zeta})$ to $\mathcal{T}_{p,k}(f)$. \diamond

Example 2.12. Continuing Example 2.10, $f_{1,1}$ has degree 10, and 10 monomial terms, but $\tilde{f}_{1,1}(x) = x^3 + 2x^2$ which has roots 0 and 1. Observe in particular that f has roots $1 + 0 \cdot 3$ and $1 + 1 \cdot 3$ in $\mathbb{Z}/(3^2)$, and the degenerate root 1 of \tilde{f} in $\mathbb{Z}/(3)$ can *not* be lifted to either of these roots via the classical Hensel's Lemma. \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 2.13. [39] The set $\mathcal{T}_{p,k}(f)$ naturally admits the structure of a labeled, rooted, directed tree as follows²

- (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 labeled by the $(f_{i,\zeta}, k_{i,\zeta}) \in \mathcal{T}_{p,k}(f)$ 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 $\zeta_{i-1} \in \mathbb{F}_p$ of $\tilde{f}_{i-1,\mu}$ with $s(f_{i-1,\mu}, \zeta_{i-1}) \in \{2, \dots, k_{i-1,\mu} - 1\}$ and $\zeta = \mu + \zeta_{i-1}p^{i-1} \in \mathbb{Z}/(p^i)$. \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 $\mathcal{T}_{p,k}(f)$ [39,27]. We will ultimately use $\mathcal{T}_{p,k}(f)$, with k determined by a root separation/valuation condition (see Corollary 6.6 below), to efficiently *count* the roots of f in \mathbb{Z}_p , and then in \mathbb{Q}_p by rescaling.

Example 2.14. $\mathcal{T}_{p,k}(x^2)$ is a chain of length $\lfloor \frac{k-1}{2} \rfloor$ for any p, k . \diamond

Example 2.15. Let $f(x) = 1 - x^{397}$. Then $\mathcal{T}_{17,k}(f)$, for any $k \geq 1$, consists of a single node, labeled $(1 - x^{397}, k)$, 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 2.16. 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, labeled $(1 - x^{340}, k)$. 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, labeled $(f_{1,\zeta_0}, k-2)$ where the 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

² This definition differs slightly from the original in [39]: the edges are unlabeled here.

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 — originally defined for efficient root counting over $\mathbb{Z}/(p^k)$ — thus encode individual base- p digits of roots of f in \mathbb{Z}_p . Their degree also decays in a manner depending on root multiplicity.

Lemma 2.17. [39, Lem. 2.2 & 3.6] Suppose $f \in \mathbb{Z}[x] \setminus p\mathbb{Z}[x]$ has degree d , $f_{0,0} := f$, $i \geq 1$, $\mu := \zeta_0 + \dots + p^{i-2}\zeta_{i-2}$ is a root of the mod p^{i-1} reduction of f , $\zeta' := \mu + p^{i-1}\zeta_{i-1}$, the pairs $(f_{i-1,\mu}, k_{i-1,\mu})$ and $(f_{i,\zeta'}, k_{i,\zeta'})$ both lie in $\mathcal{T}_{p,k}(f)$, and ζ_{i-1} has multiplicity m as a root of $\tilde{f}_{i-1,\mu}$ in \mathbb{F}_p . Then:

- (1) $\mathcal{T}_{p,k}(f)$ has depth $\leq \lfloor (k-1)/2 \rfloor$ and at most $\lfloor d/2 \rfloor$ nodes at depth $i \geq 1$.
- (2) $\deg \tilde{f}_{i,\zeta'} \leq s(f_{i-1,\mu}, \zeta_{i-1}) \leq \min\{k_{i-1,\mu} - 1, m\}$.
- (3) $f_{i,\zeta'}(x) = p^{-s} f(\zeta_0 + \zeta_1 p + \dots + \zeta_{i-1} p^{i-1} + p^i x)$ where $s := \sum_{j=0}^{i-1} s(f_{j,\zeta_0 + \dots + \zeta_{j-1} p^{j-1}}, \zeta_j) \geq 2i$.
- (4) $f(\zeta_0 + \zeta_1 p + \dots + \zeta_{i-1} p^{i-1}) = 0 \pmod{p^s}$.
- (5) $f'(\zeta_0 + \zeta_1 p + \dots + \zeta_{i-1} p^{i-1}) = 0 \pmod{p^i}$. \blacksquare

Note that Assertion (1) of Lemma 2.17 gives us an upper bound on the depth of $\mathcal{T}_{p,k}(f)$ as a function of k . We will also need to consider lower bounds on k that guarantee that $\mathcal{T}_{p,k}(f)$ has enough depth to be useful for approximating roots in \mathbb{Z}_p .

Let $n_p(f)$ denote the number of non-degenerate roots in \mathbb{F}_p of the mod p reduction of f . We will need to show that the roots of f in \mathbb{Z}_p can be embedded into a collection of series indexed by the non-degenerate roots of the nodal polynomials of $\mathcal{T}_{p,k}(f)$ in \mathbb{F}_p for k sufficiently large.

Lemma 2.18. Suppose $f \in \mathbb{Z}[x]$, $\zeta = \sum_{j=0}^{\infty} \zeta_j p^j \in \mathbb{Z}_p$ is a non-degenerate root of f , and let D be the maximum of $\text{ord}_p(\zeta - \xi)$ over all distinct non-degenerate roots $\zeta, \xi \in \mathbb{Z}_p$ of f (if f has at least 2 non-degenerate roots in \mathbb{Z}_p) or 0 (if f has 1 or fewer non-degenerate roots in \mathbb{Z}_p). Then for all k sufficiently large, $\mathcal{T}_{p,k}(f)$ has a nodal polynomial $f_{j,\zeta'}$ such that $j \leq \lfloor (k-1)/2 \rfloor$ and $\zeta' + p^j \zeta_j = \zeta \pmod{p^{j+1}}$ for some non-degenerate root ζ_j of $\tilde{f}_{j,\zeta'}$. Furthermore, for k sufficiently large we also have that $\mathcal{T}_{p,k}(f)$ has depth $\geq D$, the set $\{(g, j) \in \mathcal{T}_{p,k}(f) \mid n_p(g) > 0\}$ remains fixed and finite, and f has exactly $\sum_{(g,j) \in \mathcal{T}_{p,k}(f)} n_p(g)$ non-degenerate roots in \mathbb{Z}_p .

Proof. First note that $f(\zeta_0 + \dots + \zeta_i p^i) = 0 \pmod{p^{i+1}}$ for all $i \geq 0$. By Definitions 2.9 and 2.13, $s_0 := s(f, \zeta_0) \in \{1, \dots, m\}$, where m is the multiplicity of ζ_0 as a root of \tilde{f} (thanks to Lemma 2.17). Should $m = 1$ then $s_0 = 1$, leaving $f_{0,0} = f$ as our desired nodal polynomial (with ζ_0 a non-degenerate root of $\tilde{f}_{0,0}$) for all $k \geq 1$. Otherwise, $s_0 \geq 2$ (by the definition of $s(\cdot, \cdot)$), in which case $k \geq 1 + s_0 \implies \mathcal{T}_{p,k}(f)$ will have

$f_{1,\zeta_0}(x) = p^{-s_0} f(\zeta_0 + px)$ as a nodal polynomial. However, we need to check if ζ_1 is a non-degenerate root for \tilde{f}_{1,ζ_0} or not.

Proceeding inductively, note that if $i \geq 1$, $\zeta' := \zeta_0 + \zeta_1 p + \cdots + \zeta_{i-1} p^{i-1}$, $s_i := s(f_{i,\zeta'}, \zeta_i)$, and $s' := s_0 + \cdots + s_i$, then $s_i \in \{1, \dots, m\}$ where m is now the multiplicity of ζ_i as a root of $\tilde{f}_{i,\zeta'}$. As before, $m = 1$ implies that $f_{i,\zeta'}$ is our desired nodal polynomial (with ζ_i a non-degenerate root of $\tilde{f}_{i,\zeta'}$) for all $k \geq 1 + s'$. Otherwise, $s_i \geq 2$, in which case $k \geq 1 + s' \implies \mathcal{T}_{p,k}(f)$ will have $f_{i+1,\zeta'+p^i\zeta_i}(x) = p^{-s'} f(\zeta' + p^i \zeta_i + p^{i+1} x)$ as a nodal polynomial, and then we check if ζ_{i+1} is a non-degenerate root for $\tilde{f}_{i+1,\zeta'+p^i\zeta_i}$ or not.

Our induction must end, in finitely many steps, with our desired $f_{j,\zeta'}$. To see why, first observe that nodal polynomials always have integer coefficients and, if $d' := \text{ord}_p f'(\zeta)$, then $d' < \infty$ since ζ is a non-degenerate root and thus $f'(\zeta) = \alpha p^{d'} \pmod{p^{d'+1}}$ for some $\alpha \in \mathbb{Z}_p \setminus p\mathbb{Z}_p$. So if our induction reaches $f_{i,\zeta'}$ with $i \geq d'$, then $\zeta' = \zeta_0 + \cdots + p^{d'-1} \zeta_{d'-1} \implies f'_{d',\zeta'}(\zeta_{d'}) = \alpha p^{2d' - (s_0 + \cdots + s_{d'-1})}$. We thus obtain $2d' \geq s_0 + \cdots + s_{d'-1}$ and, for all $i \in \mathbb{N}$ with $f_{i,\zeta'}$ belonging to a node of $\mathcal{T}_{p,k}(f)$ with a child, the definition of s_i tells us that $s_i \geq 2$. Since $\text{ord}_p f'(\zeta_0 + \cdots + p^i \zeta_i) = d'$ for all $i \geq d'$, we must eventually encounter a $j \geq d'$ with $s_j = 1$, meaning no child for $f_{j,\zeta'}$. So our induction ends with a nodal polynomial $f_{j,\zeta'}$ with no degenerate roots. Moreover, we must have $\tilde{f}_{j,\zeta'}(\zeta_j) = 0 \pmod{p}$ (by definition of ζ and $f_{j,\zeta'}$) and thus ζ_j must be a non-degenerate root of $\tilde{f}_{j,\zeta'}$. Also, our upper bound on j is immediate from Lemma 2.17.

To prove that $\mathcal{T}_{p,k}(f)$ has depth for D for k large enough, note that an f with no non-degenerate roots in \mathbb{Z}_p can *not* have a tree $\mathcal{T}_{p,k}(f)$ having nodal polynomials with non-degenerate roots in \mathbb{F}_p . This is because of the equality

$$f_{i,\zeta'}(x) = p^{-s} f(\zeta_0 + \zeta_1 p + \cdots + \zeta_{i-1} p^{i-1} + p^i x)$$

from Lemma 2.17: $\tilde{f}_{i,\zeta'}$ having a non-degenerate root in \mathbb{F}_p would imply by Hensel's Lemma that f has a root $\zeta \in \mathbb{Z}_p$ with $\text{ord}_p f'(\zeta) < \infty$. So in this case, the stated set of (g, j) is empty for all $k \geq 1$ and the stated sum is 0. In particular, $\mathcal{T}_{p,k}(f)$ always at least has its root node (by definition) and thus $D \geq 0$.

Similarly, an f with just one non-degenerate root in \mathbb{Z}_p can *not* have a tree $\mathcal{T}_{p,k}(f)$ having two distinct nodal polynomials having non-degenerate roots mod p . (Likewise, $\mathcal{T}_{p,k}(f)$ having a single nodal polynomial with two distinct non-degenerate roots mod p is impossible.) So in this case, the stated set of (g, j) has cardinality 1 (with $n_p(g) = 1$ for exactly one pair (g, j)) for all k as specified in the first assertion of our lemma, which we've already proved. So the remaining assertions follow.

So let us now assume f has at least 2 distinct non-degenerate roots in \mathbb{Z}_p . There are clearly no more than $\deg f$ such roots, so our first assertion implies that, for k sufficiently large, *every* non-degenerate root $\zeta \in \mathbb{Z}_p$ of f has an associated node in $\mathcal{T}_{p,k}(f)$ encoding ζ , i.e., $\mathcal{T}_{p,k}(f)$ has depth at least D for k sufficiently large. Clearly then, the set $\{(g, j) \in \mathcal{T}_{p,k}(f) \mid n_p(g) > 0\}$ is finite and will not change as k increases: This is because the set can not lose elements as k increases, and any new element would introduce a new non-degenerate root for f via Hensel's Lemma.

So we now only need to prove that the stated sum counts roots correctly. Toward this end, note by construction that every non-degenerate root $\zeta \in \mathbb{Z}_p$ of f is associated to a unique sequence of the form $(\zeta_0, \dots, \zeta_i) \in \mathbb{F}^{i+1}$ with $\zeta_0, \dots, \zeta_{i-1}$ all degenerate roots for previously defined nodal polynomials, but with ζ_i a *non-degenerate* root of $\tilde{f}_{i,\zeta'}$. So the number of non-degenerate roots of f in \mathbb{Z}_p is no greater than the stated sum.

To conclude, note that Hensel's Lemma (and our earlier observation that nodal polynomials are rescaled shifts of f) implies that each non-degenerate root in \mathbb{F}_p of a nodal polynomial lifts to a unique root of f in \mathbb{Z}_p . Furthermore, since the derivatives of nodal polynomials are rescaled shifts of f' , each such lifted root is a non-degenerate root. So the number of non-degenerate roots of f in \mathbb{Z}_p is at least as large as the stated sum, and we are done. ■

2.6. Trees and extracting digits of radicals

We prove the following crucial lemma in Remark 6.5 of Section 6, but state it now so can apply it in our first algorithm for solving binomials:

Lemma 2.19. *Suppose $f(x) = c_1 + c_2 x^d \in \mathbb{Z}[x]$ with $c_1 c_2 \neq 0 \pmod{p}$ and $\ell := \text{ord}_p d$. Then every non-root nodal polynomial $f_{i,\zeta}$ of $\mathcal{T}_{p,k}(f)$ satisfies $\deg \tilde{f}_{i,\zeta} \leq 2$ or $\deg \tilde{f}_{i,\zeta} \leq 1$, according as $p = 2$ or $p \geq 3$. In particular, $f(\zeta_0) = 0 \pmod{p}$ for some $\zeta_0 \in \{0, \dots, p-1\} \implies s(f, \zeta_0) \leq \ell + 1$. ■*

Remark 2.20. It is a simple exercise to prove, from Lemma 2.19 and Definition 2.13, that $\mathcal{T}_{p,k}(f)$ always has depth ≤ 1 for $f \in \mathbb{Z}[x]$ a binomial with $f(0) \neq 0 \pmod{p}$. The family of examples $x^{p^2} - 1$ (for any $k \geq 4$) shows that this depth can be attained for any prime p . ◇

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 2.21. *Suppose $f \in \mathbb{Z}[x]$ is a binomial of degree d with coefficients of absolute value at most H , $f(0) \neq 0$, $\gamma = \gcd(d, \max\{2, p-1\})$, and $\{\zeta_1, \dots, \zeta_\gamma\}$ is the set of roots of f in \mathbb{Q}_p . Then in time*

$$\left(\frac{p}{\gamma} + \gamma + \log d \right) \log^{1+o(1)}(dp) + \log^{2+o(1)}(pH),$$

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 with associated true root ζ_j .

An algorithm that proves Theorem 2.21 when p is odd is outlined below.

Algorithm 2.22. (Solving binomial equations over \mathbb{Q}_p^* for odd 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 $O\left(\log\left(dH^{1/d}\right)\right)$ such that $\gamma = \gcd(d, p-1)$, z_j is an approximate root 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: Let $\ell := \text{ord}_p d$ and replace f with $f(x) := c'_1 + c'_2 x^d$ where $c'_i := \frac{c_i}{p^{\text{ord}_p c_i}}$ for all i .
- 3: If $\left(-\frac{c'_1}{c'_2}\right)^{p^\ell(p-1)/\gamma} \neq 1 \pmod{p^{2\ell+1}}$ then say "No roots in \mathbb{Q}_p !" and STOP.
- 4: Let $\delta := 1$. If $d \leq -1$ then set $\delta := -1$ and respectively replace d by $|d|$ and $f(x)$ by $x^d f(1/x)$.
- 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 $\tilde{h}(x) := x^\gamma - c'$.
- 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 $x_j - \frac{f(x_j)/p^\ell}{f'(x_j)/p^\ell} \in \mathbb{Z}/(p^2)$.
- 9: **Output** $\left\{(x_1 p^{\text{ord}_p(c_1/c_2)/d})^\delta, \dots, (x_\gamma p^{\text{ord}_p(c_1/c_2)/d})^\delta\right\}$.

Remark 2.23. Step 6 above is designed for simplicity rather than practicality, and can be sped up considerably if one avails to more sophisticated algorithms with complexity quasi-linear in $\gcd(d, p-1) + \log(pd)$: See, e.g., [1, 18, 21]. \diamond

The following algorithm proves the $p = 2$ case of Theorem 2.21.

Algorithm 2.24. (Solving binomial equations over \mathbb{Q}_2^*)**Input.** $c_1, c_2, d \in \mathbb{Z} \setminus \{0\}$ 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}_2 , or $z_1, \dots, z_\gamma \in \mathbb{Q}$ with logarithmic height $O\left(\log\left(dH^{1/d}\right)\right)$ such that $\gamma = \gcd(d, 2)$, 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 $\text{ord}_2 c_1 \neq \text{ord}_2 c_2 \pmod{d}$ then say "No roots in \mathbb{Q}_2 !" and STOP.
- 2: Let $\ell := \text{ord}_2 d$ and replace f with $f(x) := c'_1 + c'_2 x^d$ where $c'_i := \frac{c_i}{2^{\text{ord}_2 c_i}}$ for all i .
- 3: If $c'_1 \neq -c'_2 \pmod{8}$ or $\left(-\frac{c'_1}{c'_2}\right)^{2^{\ell-1}} \neq 1 \pmod{2^{2\ell+1}}$ then say "No roots in \mathbb{Q}_2 !" and STOP.
- 4: Let $\delta := 1$. If $d \leq -1$ then set $\delta := -1$ and respectively replace d by $|d|$ and $f(x)$ by $x^d f(1/x)$.
- 5: Let $x_1 := 1$. If $\gamma = 1$ then GOTO Step 7.
- 6: Let $x_2 := 3$.
- 7: **Output** $\left\{x_1 2^{\text{ord}_2(c_1/c_2)/d}, \dots, x_\gamma 2^{\text{ord}_2(c_1/c_2)/d}\right\}$.

Remark 2.25. Our correctness proof below shows that, for *binomials*, knowing the 2 most significant base- p digits of a root in \mathbb{Q}_p is enough to yield an approximate root in the sense of Smale, *independent of d and H* . Note, however, that each subsequent application of Newton's method to refine an approximation has complexity depending on $\log(dH)$ as well as $\log p$. \diamond

Remark 2.26. We point out that the approximate roots output by our two algorithms above require the use of Newton iteration *applied to f_{1,ζ_0} (instead of f) when $p|d$* . This is clarified in our correctness proof below. \diamond

Proof of Theorem 2.21. It clearly suffices to prove the correctness of Algorithms 2.22 and 2.24, and then analyze their complexity.

Correctness: (Case of odd p) 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 4 allow us to reduce our search for approximate roots to $(\mathbb{Z}/(p^{2\ell+1}))^*$ and assume positive degree d .

Lemma 2.5 implies that Step 3 simply check that the coset of roots of f in \mathbb{C}_p^* intersects \mathbb{Z}_p^* .

Step 5 is the application of an automorphism of \mathbb{F}_p^* so we can reduce the degree of our binomial to γ , which is possibly much smaller than both $p-1$ and d .

Steps 6–7 then clearly find the correct coset of \mathbb{F}_p^* that makes f vanish mod p . In particular, by Hensel's Lemma, Step 9 clearly gives the correct output if $\ell=0$. (Recall that we have replaced each coefficient c_i of f with c'_i .)

If $\ell \geq 1$ then let ζ_0 be any x_j from Step 8. We then have $\deg \tilde{f}_{1,\zeta_0} \leq 1$ thanks to Lemma 2.19. Furthermore, Definition 2.11 tells us that the unique root $\zeta_1 \in \mathbb{F}_p$ of \tilde{f}_{1,ζ_0} is exactly the next base- p digit of a unique root $\zeta \in \mathbb{Z}_p$ of f with $\zeta = \zeta_0$. Also, $\deg \tilde{f}_{1,\zeta_0}$ must be 1 (for otherwise \tilde{f} would not vanish on its coset of roots in \mathbb{F}_p^*) and $s(f, \zeta_0) \geq 2$ since $\ell \geq 1$ forces ζ_0 to be a degenerate root of \tilde{f} . Lemma 2.17 then tells us that Hensel's Lemma — applied to $f_{1,\zeta_0}(x) = p^{-s(f,\zeta_0)} f(\zeta_0 + px)$ and start point $\zeta_1 \in \mathbb{Z}/(p)$ — implies that $\zeta_0 + \zeta_1 p$ yields Newton iterates rapidly converging to a true root $\zeta \in \mathbb{Z}_p$. So Step 8 in fact refines x_1 to the mod p^2 quantity $\zeta_0 + \zeta_1 p$, and thus Steps 7–9 indeed give us suitable approximants in \mathbb{Q} to all the roots of f in \mathbb{Q}_p . So our algorithm is correct.

Note also that the outputs, being integers in $\{0, \dots, p^2 - 1\}$ rescaled by a factor of $p^{\text{ord}_p(c_1/c_2)/d}$ (or possibly the reciprocals of such quantities), clearly each have bit-length $O\left(\log(p) + \frac{\lfloor \log(c_1/c_2) \rfloor}{d \log p} \log p\right) = O\left(\log(p) + \frac{\log H}{d}\right) = O(\log(pH^{1/d})).$ ■

(Case of $p=2$) The proof is almost the same as the Correctness proof for odd p , save that we respectively replace Lemma 2.5 and Algorithm 2.22 by Lemma 2.6 and Algorithm 2.24. In particular, Steps 5–8 of Algorithm 2.22 collapse into Steps 5–6 of Algorithm 2.24.

So we must explain Steps 5–6 here: These steps give us the mod 4 reductions of the γ many roots of f in \mathbb{Z}_2 , since Steps 5 and 6 are executed only after Steps 1 and 3 certify that f indeed has roots in \mathbb{Z}_2 . (Remember that $\gamma \in \{1, 2\}$ for $p=2$.) Furthermore, Hensel's Lemma implies that the root 1 of \tilde{f} lifts to the sole root of f in \mathbb{Z}_2 when $\ell=0$. So the case $\ell=0$ is done.

If $\ell \geq 1$ then there is one more complication: The nodal polynomial $\tilde{f}_{1,1}$ is now quadratic. This is because Lemma 2.19 tells us that $\deg \tilde{f}_{1,1} \leq 2$. Furthermore, $\ell \geq 1$

implies that $\gamma = 2$ (assuming there are roots in \mathbb{Z}_2 and the algorithm hasn't terminated already) and thus f must have exactly 2 roots in \mathbb{Z}_2 . Lemma 2.18 then tells us that $\deg \tilde{f}_{1,1} \leq 1$ would imply f has ≤ 1 root in \mathbb{Z}_2 . Therefore, $\tilde{f}_{1,1}$ must be quadratic.

Furthermore, $\tilde{f}_{1,1}$ must also have 2 distinct roots: This is because $\tilde{f}_{1,1}$ equal to x^2 or $1 + x^2 = (1 + x)^2 \bmod 2$ would imply that no nodal polynomial $\tilde{f}_{i,\zeta}$, for $i \geq 1$, has a non-degenerate root. So, again by Lemma 2.18, we would not attain 2 roots in \mathbb{Z}_2 . (Similarly, it is impossible for $\tilde{f}_{1,1}$ to be irreducible.) Therefore, the mod 4 reductions of the two roots of f in \mathbb{Z}_2 must be 1 and 3. So Steps 5–6 are indeed correct.

Lemma 2.17 then tells us that Hensel's Lemma — applied to $f_{1,1}(x) = 2^{-s(f,1)} f(1+2x)$ and *either* start point 0 or 1 in $\mathbb{Z}/(2)$ — implies that $1 + 0 \cdot 2$ and $1 + 1 \cdot 2$ yield sequences of iterates rapidly converging to true roots in \mathbb{Z}_2 . So Steps 5–7 indeed give us suitable approximants in \mathbb{Q} to all the roots of f in \mathbb{Q}_2 , and our algorithm is correct.

Note also that the outputs, being integers in $\{1, 3\}$ rescaled by a factor of $2^{\text{ord}_2(c_1/c_2)/d}$ (or possibly the reciprocals of such quantities), clearly each have bit-length

$$O\left(\frac{|\log(c_1/c_2)|}{d \log 2} \log 2\right) = O\left(\frac{\log H}{d}\right) = O\left(\log(H^{1/d})\right). \blacksquare$$

Complexity Analysis: (Case of odd p) Via Corollary 2.8, [56], and Theorem 2.7, it is easily checked that Steps 1–5 of Algorithm 2.22 have respective complexity bounds

1. $O(\log(H) \log(pH) \log \log(pH)) + O(\log(d) \log \log d)$
2. $O(\log(d) \log(dp) \log \log(dp))$
3. $O(\log^2(d) \log \log d)$
4. (time negligible compared to the preceding quantities)
5. $O(p^{1/4} \log(p) \log \log(p)) + O(\log^2(p) \log \log(p))$

These add up to time no worse than

$$O(p^{1/4} \log(p) \log \log(p) + \log(H) \log(pH) \log \log(pH) + \log(d) \log(dp) \log \log(dp))$$

so far. 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 $\mathbb{Z}/(p^{2\ell+1})$. Since $\ell \log p \leq \log d$, this takes time no worse than $O\left(\frac{p}{\gamma} \log(p) \log \log(p) + \gamma \log(d) \log \log d\right)$, which is bounded from above by $O\left(\left(\frac{p}{\gamma} + \gamma\right) \log(dp) \log \log(dp)\right)$. Note also that $\frac{p}{\gamma} + \gamma \geq 2\sqrt{p}$ by the Arithmetic-Geometric Inequality. So our final complexity bound is bounded from above by

$$O\left(\left(\frac{p}{\gamma} + \gamma + \log d\right) \log(dp) \log \log(dp) + \log(H) \log(pH) \log \log(pH)\right). \blacksquare$$

(Case of $p = 2$) We simply use the same techniques as for Algorithm 2.22, save for Steps 5–8 there being collapsed into Steps 5–6 in Algorithm 2.24. \blacksquare

3. Proving Theorem 1.6: trinomial roots never get too close

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

Theorem 3.1. [65, Pg. 190] 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 $\alpha_1^{b_1} \cdots \alpha_n^{b_n} \neq 1$ implies that $\alpha_1^{b_1} \cdots \alpha_n^{b_n} - 1$ has p -adic valuation strictly less than

$$\log(2) \log_p(2n) n^{5/2} (256e^2)^{n+1} p \log_p(B) \prod_{i=1}^n \max \left\{ \log |r_i|, \log |s_i|, \frac{1}{16e^2} \right\},$$

where $B := \max\{|b_1|, \dots, |b_n|, 3\}$. In particular, $\log(2)256e^2 < 1312$, $256e^2 < 1892$, and $\frac{1}{16e^2} < 0.0085$. ■

We will prove the square-free case of Theorem 1.6 here, postponing the proof of the non-square-free case to Section 5.1. To prove that two distinct roots $\zeta_1, \zeta_2 \in \mathbb{C}_p$ of a square-free trinomial f can not be too close, we will prove that f' has a root $\tau \in \mathbb{C}_p$ with three special properties: (i) $|f(\tau)|_p$ is not too small, (ii) $|\zeta_1 - \zeta_2|_p \geq p^{-1/(p-1)}|\zeta_1 - \tau|_p$, and (iii) $|\zeta_1 - \tau|_p$ is not too small. Step (i) is where we avail to Yu's Theorem, so let us now quantify our approach.

Proposition 3.2. Suppose $f(x) = c_1 + c_2x^{a_2} + c_3x^{a_3} \in \mathbb{Z}[x]$ is a trinomial of degree $d = a_3 > a_2 \geq 1$, with all its coefficients having absolute value at most H , and $\tau \in \mathbb{C}_p$ is a root of f' . Then $\tau^{a_3-a_2} = -\frac{a_2c_2}{a_3c_3}$ and $f(\tau) = c_1 + c_2\tau^{a_2} \left(1 - \frac{a_2}{a_3}\right)$. ■

Lemma 3.3. Following the notation above, assume further that f is square-free. Then $|f(\tau)|_p \geq \exp[-O(p \log_p(d) \log^2(dH))]$.

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 $f(\tau) \neq 0$. Proposition 3.2 we then obtain that $\text{ord}_p f(\tau)$ is

$$\begin{aligned} & \text{ord}_p(c_1 + c_2\tau^{a_2}(1 - a_2/a_3)) \\ &= \text{ord}_p(c_1) + \text{ord}_p(-1) + \text{ord}_p \left(\frac{-(a_3 - a_2)c_2}{a_3c_1} \left(-\frac{a_2c_2}{a_3c_3} \right)^{a_2/(a_3 - a_2)} - 1 \right). \end{aligned} \quad (1)$$

Clearly, $\text{ord}_p c_1 \leq \frac{\log H}{\log p}$ and $\text{ord}_p(-1) = 0$. To bound the third summand on the right-hand side of Equality (1) above, let $T := \frac{-(a_3 - a_2)c_2}{a_3c_1} \left(-\frac{a_2c_2}{a_3c_3} \right)^{a_2/(a_3 - a_2)}$ and observe that $T^{a_3 - a_2} - 1 = \prod_{j=1}^{a_3 - a_2} (T - \omega^j)$ for $\omega \in \mathbb{C}_p$ a primitive $(a_3 - a_2)$ -th root of unity. In particular, $T^{a_3 - a_2} \neq 1$ since $f(\tau\omega^j) \neq 0$ for all $j \in \{1, \dots, a_3 - a_2\}$, thanks to Proposition 3.2 and f not having any degenerate roots. So then $M := \text{ord}_p(T^{a_3 - a_2} - 1) = \sum_{j=1}^{a_3 - a_2} \text{ord}_p(T - \omega^j) < \infty$, with the $(a_3 - a_2)$ -th term of the sum exactly $\text{ord}_p(T - \omega^{a_3 - a_2}) = \text{ord}_p(T - 1)$, i.e., the third summand from Equality (1).

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

$$\text{ord}_p f(\tau) = \text{ord}_p(c_1) + \text{ord}_p(T - \omega^{a_3 - a_2}) < \frac{\log_p(dH)}{1}$$

(by Theorem 2.3) and we obtain our lemma.

On the other hand, should $\text{ord}_p T \geq 0$, we get $\text{ord}_p(T - \omega^j) \geq j \text{ord}_p(\omega) = 0$, for each j . 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 Yu's Theorem (Theorem 3.1 here) upon setting $n = 2$, $\alpha_1 = -\frac{(a_3 - a_2)c_2}{a_3c_1}$, $\alpha_2 = -\frac{a_2c_2}{a_3c_3}$, $b_1 = a_3 - a_2$, and $b_2 = a_2$. In particular, we can assume $|r_i|, |s_i| \leq dH$ for $i \in \{1, 2\}$ and $B = \max\{d, 3\}$, and move the $\log p$ factors in the denominator so that $M < \log(2)256e^2 \log(4)2^{5/2}(256e^2)^2 p \log \max\{d, 3\} \left(\max\left\{ \log_p(dH), \frac{1}{16e^2 \log p} \right\} \right)^2$. For $d = 2$ we get $f(\tau) = \frac{c_1}{4c_3}(4c_1c_3 - c_2^2)$, which is a rational number that this an integer of absolute value at most $H^2 + 4H$ divided by an integer of absolute value at most $4H$. Such a rational number clearly has valuation no greater than $\log_p(H^2 + 4H) = O(\log_p H)$ and thus $|f(\tau)|_p \geq e^{-O(\log H)}$ when $d = 2$. Since $d \geq 2$ for an arbitrary trinomial, and $H \geq 1$, we then obtain $M < 36791093348p \log(d) \log_p^2(dH) = O(p \log(d) \log_p^2(dH))$. In other words, the third summand from (1) is bounded from above by the last O -bound, and thus $\text{ord}_p f(\tau) = O(M)$ since $\frac{\log H}{\log p} = O(M)$. Since $|f(\tau)|_p = e^{-\log(p) \text{ord}_p f(\tau)}$, we are done. ■

The Ultrametric Inequality directly yields the following:

Proposition 3.4. *If $f \in \mathbb{Z}[x]$ and $t \in \mathbb{C}_p$ then $|t|_p \leq 1 \implies |f'(t)|_p \leq 1$. ■*

Below is a rescaled p -adic version of *Rolle's Theorem*, based on [49, 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 $\tau \in \mathbb{C}_p$ with $|\zeta_1 - \tau|_p, |\zeta_2 - \tau|_p \leq c$. ■*

We can now prove part of one of our main results.

Proof of the square-free case of Theorem 1.6. Note that $\zeta_i \neq 0 \implies |\text{ord}_p \zeta_i| \leq \log_p H$ thanks to Theorem 2.3. So then $\text{ord}_p(\zeta_1 - \zeta_2) \geq -\log_p H$ for any pair of distinct roots $\zeta_1, \zeta_2 \in \mathbb{C}_p$ of f and, if $\zeta_1 \zeta_2 = 0$, we also have $\text{ord}_p(\zeta_1 - \zeta_2) \leq \log_p H$. So $\log H \geq \log |\zeta_1 - \zeta_2|_p$ and, if $\zeta_1 \zeta_2 = 0$ then we also have $\log |\zeta_1 - \zeta_2|_p \geq -\log H$. So we may assume $\zeta_1 \zeta_2 \neq 0 \neq f(0)$.

For convenience, let us abbreviate the first (larger) O -bound stated in our theorem by $O(M')$.

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

Suppose $|\zeta_1 - \zeta_2|_p > p^{-2/(p-1)}$ ($= e^{-2\log(p)/(p-1)}$). 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 f' has a root $\tau \in \mathbb{C}_p$ with $|\zeta_i - \tau|_p \leq p^{1/(p-1)} |\zeta_1 - \zeta_2|_p \leq p^{-1/(p-1)}$ for all $i \in \{1, 2\}$. Since f is square-free, Lemma 3.3 implies that $|f(\tau)|_p \geq e^{-O(M')}$. Applying Theorem 3.5 to

$$g(x) := f(x) - \frac{f(\tau) - f(\zeta_1)}{\tau - \zeta_1} x - \frac{\tau f(\zeta_1) - \zeta_1 f(\tau)}{\tau - \zeta_1}$$

(which vanishes at τ and ζ_1), we then see that there is a $\mu \in \mathbb{C}_p$ with $|\mu - \zeta_1|_p \leq 1$ such that $g'(\mu) = 0$, i.e., $f(\tau) - f(\zeta_1) = f'(\mu)(\tau - \zeta_1)$. Note that $|\mu|_p \leq 1$ since $|\mu|_p > 1$ would imply that $|\mu|_p > |\zeta_1|_p$ and thus $|\mu - \zeta_1|_p = |\mu|_p > 1$, giving us a contradiction. As $f(\tau) \neq 0$ we get $f'(\mu) \neq 0$ and $\tau \neq \zeta_1$. From Proposition 3.4 we have $|f'(\mu)|_p \leq 1$, so then $|\tau - \zeta_1|_p = \frac{|f(\tau)|_p}{|f'(\mu)|_p} \geq e^{-O(M')}$. We thus get $|\zeta_1 - \zeta_2|_p \geq p^{-1/(p-1)} |\tau - \zeta_1|_p \geq e^{-O(M') - \frac{\log p}{p-1}} = e^{-O(M')}$. ■

Case 2: (Both roots are large: $|\zeta_1|_p, |\zeta_2|_p > 1$.) Simply observe that $1/\zeta_1$ and $1/\zeta_2$ are roots of the reciprocal polynomial $f^*(x) := x^{\deg f} f(\frac{1}{x})$. In particular, we can apply Case 1 to the trinomial f^* since $\left| \frac{1}{\zeta_1} \right|_p, \left| \frac{1}{\zeta_2} \right|_p < 1$. We then obtain $\left| \frac{1}{\zeta_1} - \frac{1}{\zeta_2} \right|_p \geq e^{-O(M')}$. Hence $|\zeta_1 - \zeta_2|_p = |\zeta_1|_p |\zeta_2|_p \left| \frac{1}{\zeta_1} - \frac{1}{\zeta_2} \right|_p \geq \left| \frac{1}{\zeta_1} - \frac{1}{\zeta_2} \right|_p \geq e^{-O(M')}$. ■

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 1.5: tetranomial roots can get too close

4.1. The case of prime p

Let $g(x) = p^{2j} f(x + p^{j-1}) = p^{2j} (x + p^{j-1})^d - p^{2j} \left(\frac{x + p^{j-1}}{p^j} - \frac{1}{p} \right)^2 = p^{2j} (x + p^{j-1})^d - x^2$. Then g has the same roots as $f_{d,p,j}$, save for a “small” shift by p^{j-1} . Rescaling, we get $G(x) := \frac{g(p^{(j-1)d/2+j}x)}{p^{(j-1)d+2j}} = p^{-(j-1)d-2j} [p^{2j}(p^{(j-1)d/2+j}x + p^{j-1})^d - p^{(j-1)d+2j}x^2] = \sum_{i=0}^d \binom{d}{i} p^{(j-1)(di/2-i)+ij} x^i - x^2 = 1 - x^2 \pmod{p^{d(j-1)/2+1}}$, which is square-free for odd prime p . So if p is odd, then Hensel’s Lemma implies that there are roots $\zeta_1, \zeta_2 \in \mathbb{Z}_p$ of G such that $\zeta_1 \equiv 1 \pmod{p^{d(j-1)/2+1}}$ and $\zeta_2 \equiv -1 \pmod{p^{d(j-1)/2+1}}$.

On the other hand, if $p = 2$, then, as $j > 2$, we have $p^{d(j-1)/2+1} \geq 8$. So, since $G(x) = 1 - x^2 = (3 - x)(5 - x) \pmod{2^3}$, we obtain that G is square-free in $\mathbb{Z}_2[x]$. Hensel’s Lemma then implies that there are roots $\zeta_1, \zeta_2 \in \mathbb{Z}_p$ of G such that $\zeta_1 = 3 \pmod{p^{d(j-1)/2+1}}$ and $\zeta_2 = 5 \pmod{p^{d(j-1)/2+1}}$.

So, whether p is odd or even, we obtain two roots $x_1, x_2 \in \mathbb{Z}_p$ of G with $|x_1|_p = |x_2|_p = 1$. For each $i \in \{1, 2\}$, $y_i = p^{(j-1)d/2+j} x_i$ is then the corresponding root of g . So $\zeta_1 := y_1 + p^{j-1}$ and $\zeta_2 := y_2 + p^{j-1}$ are two roots of f in \mathbb{Z}_p such that $|\zeta_1 - \zeta_2|_p = |(y_1 + p^{j-1}) - (y_2 + p^{j-1})|_p = |y_1 - y_2|_p \leq \max \{ |y_1|_p, |y_2|_p \} = p^{-(j-1)d/2-j} = p^{-\Omega(dj)}$. ■

Remark 4.1. From our proof, we see that $f_{d,p,j}$ has two roots of the form

$$\zeta_i = p^{j-1} + \varepsilon_i p^{(j-1)d/2} + O(p^{1+(j-1)d/2})$$

with $i \in \{1, 2\}$ and $\{\varepsilon_1, \varepsilon_2\}$ equal to $\{\pm 1\}$ or $\{3, 5\}$, according as p is odd or even. In particular, by direct evaluation, it is easily checked that $\text{ord}_p f'_{d,p,j}(\zeta_i) = \text{ord}_p(d) + (j-1)(d-1)$. In other words, we can need as many as $\Omega(d \log H)$ of the most significant base- p digits of a root of a tetranomial in order to use it as a start point for Newton iteration. We will see in Section 5 that $O_p(\log^3(\max\{d, H\}) \log(d))$ base- p digits suffice for trinomials. \diamond

4.2. The case $p = \infty$

Shifting by $\frac{1}{2^{j-1}}$, we get $g(x) := f_{d,\frac{1}{2},j}(x + 2^{1-j}) = (x + 2^{1-j})^d - 2^{2j}x^2 = 2^{d(1-j)} + d2^{(d-1)(1-j)}x + \left(\binom{d}{2}2^{(d-2)(1-j)} - 2^{2j}\right)x^2 + \left(\binom{d}{3}2^{(d-3)(1-j)}x^3 + \dots + x^d\right)$. We will see momentarily that, unlike $\text{Newt}_\infty(f)$ (which has 3 lower edges), $\text{Newt}_\infty(g)$ will have just 2 lower edges. (See the right-hand illustration in Example 2.2.) This will force (via Theorem 2.3) the existence of two distinct roots of small norm for g , thus yielding two nearby roots of f after undoing our earlier shift.

Toward this end, note that the three lowest order terms of g contribute the points $p_0 := (0, d(j-1) \log 2)$, $p_1 := (1, (d-1)(j-1) \log 2 - \log d)$, and $p_2 = \left(2, -\log\left(4^j - \frac{\binom{d}{2}}{2^{(d-2)(j-1)}}\right)\right)$ as potential vertices of $\text{Newt}_\infty(g)$. Observe that $\frac{\binom{d}{2}}{2^{(d-2)(j-1)}} < 0.059$ for all $j \geq 3$ and $d \geq 4$, and thus p_2 is the only point of $\text{Newt}_\infty(f)$ with negative y -coordinate. So p_2 is a vertex of $\text{Newt}_\infty(f)$, and all edges with vertices to the right of p_2 have positive slope. Furthermore, the slopes of the line segments $\overline{p_0p_1}$ and $\overline{p_0p_2}$ are respectively $-(j-1) \log(2) - \log d$ and a number less than $-\frac{1}{2} \log(4^j - 0.059) - \frac{1}{2}d(j-1) \log 2$.

Since $2^{j-1} < \sqrt{4^j - 0.059}$ and $\log d < \frac{1}{2}d(j-1) \log 2$ for all $d \geq 4$ and $j \geq 3$, we thus see that the slope of $\overline{p_0p_2}$ is more negative. So the leftmost lower edge of $\text{Newt}_\infty(g)$ has vertices p_0 and p_2 . It is easily checked that the slope of this edge is less than -10.3 , which is in turn clearly $< -2 \log 3$. So by Theorem 2.3, there are two roots z_1, z_2 of g such that

$$\log |z_i| \leq \frac{1}{2} \left[-\log \left(2^{2j} - \binom{d}{2} 2^{(d-2)(1-j)} \right) - d(j-1) \log 2 \right].$$

These two roots thus satisfy $|z_i| = 2^{-\Omega(dj)}$. Now, for $i \in \{1, 2\}$, $\zeta_i = z_i + 2^{1-j}$ yields roots of $f_{d,\frac{1}{2},j}$ with $|\zeta_1 - \zeta_2| = |z_1 + 2^{1-j} - (z_2 + 2^{1-j})| \leq |z_1| + |z_2| < 2^{-\Omega(dj)}$. \blacksquare

5. Valuation bounds from discriminants and repulsion from degeneracy

While we were able to prove a special case of our bound for the minimal root spacing of trinomials, we will need to examine the roots in \mathbb{C}_p^* more carefully for trinomials that

have degenerate roots in \mathbb{C}_p^* . We will see that the roots appear to repel more strongly in the degenerate case, and a key tool to prove this is the *trinomial discriminant*.

Definition 5.1. [29] Suppose $f(x) = c_1 + c_2x^{a_2} + c_3x^{a_3} \in \mathbb{Z}[x]$ is a trinomial with $a_3 > a_2 \geq 1$, $r := \gcd(a_2, a_3)$, and $\bar{a}_i := \frac{a_i}{r}$ for all i . We then define the *trinomial discriminant* to be

$$\Delta_{\text{tri}}(f) := \bar{a}_3^{\bar{a}_3} c_1^{\bar{a}_3 - \bar{a}_2} c_3^{\bar{a}_2} - \bar{a}_2^{\bar{a}_2} (\bar{a}_3 - \bar{a}_2)^{\bar{a}_3 - \bar{a}_2} (-c_2)^{\bar{a}_3}. \quad \diamond$$

Up to a sign factor, our definition agrees with the definition of the $\{0, a_2, a_3\}$ -discriminant from [29, Ch. 9, pp. 274–275, Prop. 1.8] when $\gcd(a_2, a_3) = 1$. We will also need to recall the following facts:

Lemma 5.2. [3, Lemma 40] *Following the notation of Definition 5.1:*

- (1) *If $c_1c_3 \neq 0$ then $\Delta_{\text{tri}}(f) \neq 0 \iff f$ has no degenerate roots in \mathbb{C}_p . Furthermore, $p \nmid c_1c_3 \gcd(a_2, a_3)$ also implies the equivalence $\Delta_{\text{tri}}(\bar{f}) \neq 0 \pmod{p} \iff \bar{f}$ has nodegenerate roots in $\bar{\mathbb{F}}_p$.*
- (2) *If $\Delta_{\text{tri}}(f) \neq 0$ then $\Delta_{\text{tri}}(f) = \left(\frac{c_3}{c_1}\right)^{\bar{a}_2-1} \prod_{\xi \in \mathbb{C}_p : \bar{f}(\xi)=0} \bar{f}'(\xi) = (-1)^{\bar{a}_3(\bar{a}_3 - \bar{a}_2)} \times \prod_{\xi \in \mathbb{C}_p : \bar{f}(\xi)=0} (\bar{a}_2c_2 + \bar{a}_3c_3\xi^{\bar{a}_3 - \bar{a}_2})$ where $\bar{f} \in \mathbb{Z}[x]$ is the unique polynomial satisfying $f(x) = \bar{f}(x^r)$ identically. ■*

Remark 5.3. The second sentence of Assertion (1) appears not to be well-known but does follow easily from the development of [29, Ch. 9], upon observing that $p \nmid \gcd(a_2, a_3) \implies$ the matrix $\begin{bmatrix} 1 & 1 & 1 \\ 0 & a_2 & a_3 \end{bmatrix}$ has rank 2. Should $p \mid \gcd(a_2, a_3)$ then it is easily checked that every root in \mathbb{F}_p^* of the trinomial \bar{f} above is degenerate. \diamond

Recall that the *classical degree d discriminant* of a polynomial $g(x) = c_0 + \cdots + c_dx^d \in \mathbb{C}_p[x]$ is $\Delta_d(g) := \frac{\text{Res}_{d_1, d-1}(f, f')}{c_d}$ where $\text{Res}_{d_1, d_2}(g_1, g_2)$ denotes the well-known *resultant* of two univariate polynomials, g_1 and g_2 , having respective degrees $\leq d_1$ and $\leq d_2$ (see, e.g., [29, Ch. 12]). We will also need some deeper facts about the discriminants of trinomials, and thereby prove repulsion from degenerate roots along the way:

Lemma 5.4. *Suppose $f(x) = c_1 + c_2x^{a_2} + c_3x^{a_3} \in \mathbb{Z}[x]$ has degree $d = a_3 > a_2 \geq 1$, $c_1c_2c_3 \neq 0$, and $|c_i| \leq H$ for all i . Assume further that f has a degenerate root $\tau \in \mathbb{C}_p$, $r := \gcd(a_2, a_3)$, and $\bar{a}_i := \frac{a_i}{r}$ for all i . Finally, let*

$$\begin{aligned} Q(x) := & (\bar{a}_3 - \bar{a}_2) (1 + 2x + 3x^2 + \cdots + (\bar{a}_2 - 1)x^{\bar{a}_2-2}) \\ & + \bar{a}_2 ((\bar{a}_3 - \bar{a}_2)x^{\bar{a}_2-1} + (\bar{a}_3 - \bar{a}_2 - 1)x^{\bar{a}_2} + \cdots + 1 \cdot x^{\bar{a}_3-2}) \end{aligned}$$

and $q(x) := (\bar{a}_3 - \bar{a}_2) - \bar{a}_3x^{\bar{a}_2} + \bar{a}_2x^{\bar{a}_3}$. Then:

- (1) Any degenerate root $\tau \in \mathbb{C}_p$ of f satisfies $\tau^r \in \mathbb{Q}^*$ and $(\tau^{a_2}, \tau^{a_3}) = \frac{c_1}{a_3 - a_2} \left(-\frac{a_3}{c_2}, \frac{a_2}{c_3} \right)$. Furthermore, if $p \nmid (a_3 - a_2)c_1$, then any degenerate root $\tilde{\tau} \in \overline{\mathbb{F}}_p$ of \tilde{f} satisfies $(c_2\tilde{\tau}^{a_2}, c_3\tilde{\tau}^{a_3}) = \frac{c_1}{a_3 - a_2}(-a_3, a_2)$ and, if $p \nmid c_2c_3$ in addition, then $\tilde{\tau}^r \in \mathbb{F}_p^*$.
- (2) The polynomial q has 1 as its unique degenerate root in \mathbb{C}_p and $q(x) = Q(x)(x - 1)^2$ identically.
- (3) We have $Q(1) = \bar{a}_2\bar{a}_3(\bar{a}_3 - \bar{a}_2)/2$ and, for $\bar{a}_3 \geq 4$, $\Delta_{\bar{a}_3-2}(Q) = \bar{a}_3(\bar{a}_2\bar{a}_3(\bar{a}_3 - \bar{a}_2))^{\bar{a}_3-4}J$, where $J = O(\bar{a}_2^2\bar{a}_3^3(\bar{a}_3 - \bar{a}_2)^2)$ is a nonzero integer.
- (4) For $\bar{a}_3 \geq 4$ we have $\Delta_{\bar{a}_3-2}(Q) = \bar{a}_2^{\bar{a}_3-4} \prod_{\mu \in \mathbb{C}_p : Q(\mu)=0} Q'(\mu)$.
- (5) $|\text{ord}_p(\zeta - \tau)| \leq \log_p \frac{(d-r)d^3H}{8r^4} < 4 \log_p \frac{dH^{1/4}}{r}$ for any non-degenerate root $\zeta \in \mathbb{C}_p$ of f .

Proof of Lemma 5.4. Assertions (1)–(3) are immediate upon applying [3, Lemma 40] to the polynomial \tilde{f} from Lemma 5.2 (which satisfies $f(x) = \tilde{f}(x^r)$). Assertion (4) follows similarly from [29, Product Formula, Pg. 398], which is a product formula for resultants. Assertion (5) will follow routinely upon proving that the roots of Q can't be too close to 1, and that the same holds for the $(1/r)$ -th powers of the roots of Q as well. In particular, we'll soon see that the r th powers of the non-degenerate roots of f are mild rescalings of the roots of Q .

Assertion (5): To simplify matters, we will first reduce to the case $r = 1$. Since the polynomial \tilde{f} from Lemma 5.2 is an instance of the case $r = 1$, and the roots of \tilde{f} are the r th powers of the roots of f , we can perform our reduction by showing that a sufficiently good upper bound on $|\text{ord}_p(\zeta^r - \tau^r)|$ implies our desired upper bound on $|\text{ord}_p(\zeta - \tau)|$. So first note that if $\text{ord}_p \zeta \neq \text{ord}_p \tau$ then $\text{ord}_p(\zeta - \tau) = \min\{\text{ord}_p \zeta, \text{ord}_p \tau\}$. In particular, since $a_3 = r\bar{a}_3$, and a_2 and $a_3 - a_2$ are positive multiples of r , Theorem 2.3 implies:

Any root of f in \mathbb{C}_p must have valuation in the closed interval $\left[\frac{\text{ord}_p(c_2/c_3)}{r}, \frac{\text{ord}_p(c_1/c_2)}{r} \right]$ (2)

or have valuation exactly $\frac{\text{ord}_p(c_1/c_3)}{r\bar{a}_3}$, according as $\text{ord}_p \frac{c_2^2}{c_1c_3} \leq 0$ or not.

So $|\text{ord}_p(\zeta - \tau)| \leq \frac{\log_p H}{r} < \log_p \frac{(d-r)d^3H}{8r^4}$, and the last inequality clearly holds when $\frac{d}{r} \geq 2$. We may thus assume $\text{ord}_p \zeta = \text{ord}_p \tau$.

Now, if $r > 1$, then we can observe that

$$\text{ord}_p(\zeta^r - \tau^r) = r \text{ord}_p(\zeta) + \text{ord}_p \left(1 - \left(\frac{\tau}{\zeta} \right)^r \right). \quad (3)$$

Letting $\omega \in \mathbb{C}_p$ be any primitive r th root of unity, we then obtain $\text{ord}_p \left(1 - \left(\frac{\tau}{\zeta} \right)^r \right) = \sum_{j=0}^{r-1} \text{ord}_p \left(1 - \frac{\tau\omega^j}{\zeta} \right)$. Since each term in the preceding sum is clearly nonnegative we

must then have $\text{ord}_p \left(1 - \frac{\tau}{\zeta}\right) \leq \text{ord}_p \left(1 - \left(\frac{\tau}{\zeta}\right)^r\right)$. So if we have $\text{ord}_p(\zeta^r - \tau^r) \leq M$ for some $M \geq r \text{ord}_p \zeta$ then Equality (3) implies $\left|\text{ord}_p \left(1 - \frac{\tau}{\zeta}\right)\right| \leq M - r \text{ord}_p \zeta$. Fact (2) then implies

$$|\text{ord}_p(\zeta - \tau)| = \left|\text{ord}_p(\zeta) + \text{ord}_p \left(1 - \frac{\tau}{\zeta}\right)\right| \leq M - (r-1) \text{ord}_p \zeta \leq M + \frac{r-1}{r} \log_p H.$$

Since $\frac{1}{r} + \frac{r-1}{r} = 1$, we will clearly establish Assertion (5) if we can prove $\text{ord}_p(\zeta^r - \tau^r) \leq \log_p \frac{(d-r)d^3H^{1/r}}{8r^4}$. Since every root of \bar{f} is the r th power of a root of f (and vice-versa), and since $\deg \bar{f} = \frac{d}{r}$ and $\gcd(\bar{a}_2, \bar{a}_3) = 1$, Fact (2) implies that it suffices to prove the following half of the $r = 1$ case of Assertion (5): $\text{ord}_p(\zeta - \tau) \leq \log_p \frac{(d-1)d^3H}{8}$. (Our stated bound is implied by the preceding bound since $\text{ord}_p \zeta = \text{ord}_p \tau \implies \text{ord}_p(\zeta - \tau) \geq 0$.) We will thus assume $\gcd(a_2, a_3) = 1$ henceforth.

(*The Case $d \in \{2, 3\}$*) Note that $d \geq 2$ because f is a trinomial. The case $d = 2$ is then vacuously true since a quadratic with a degenerate root has no non-degenerate roots.

For $d = 3$, Assertion (2) of our lemma tells us that there is only one non-degenerate root ζ and it is rational. So, evaluating the factorization of f at 0, we must have $\tau^2 \zeta = -\frac{c_1}{c_3}$. Assertion (1) of our lemma tells us that $\tau^3 = \frac{c_1 a_2}{(3-a_2)c_3}$ and thus $\frac{\zeta}{\tau} = -\frac{3-a_2}{a_2}$. So we obtain $\text{ord}_p(\tau - \zeta) = \text{ord}_p(\tau) + \text{ord}_p(1 - \frac{\zeta}{\tau}) = \frac{\text{ord}_p((c_2 a_2)/(3c_3))}{3-a_2} + \text{ord}_p \left(\frac{3-a_2}{a_2}\right)$, where the last equality follows from Theorem 2.3 applied to f' . Since $|c_2 a_2| \leq 2H$ and $3 - a_2 \leq 2$, it easily follows that $\text{ord}_p(\tau - \zeta) \leq \log_p(4H) < \log_p \frac{(d-1)d^3H}{8}$. Our assertion thus holds when $d \leq 3$. ■

(*The Case $d \geq 4$*) We will first prove an upper bound on $\text{ord}_p(1 - \mu)$ for all roots $\mu \in \mathbb{C}_p \setminus \{1\}$ of q . Observe that Assertion (2) and the classical theory of discriminants [29, Ch. 12] imply that Q has exactly $a_3 - 2$ distinct roots in \mathbb{C}_p^* and $\Delta_{a_3-2}(Q) \neq 0$. The first half of Assertion (3) then tells us that $\prod_{\mu \in \mathbb{C}_p : Q(\mu)=0} (1 - \mu) = \frac{Q(1)}{a_2} = \frac{a_3(a_3 - a_2)}{2}$,

since the leading coefficient of Q is a_2 . So then

$$\sum_{\mu \in \mathbb{C}_p : Q(\mu)=0} \text{ord}_p(1 - \mu) = \text{ord}_p \left(\frac{a_3(a_3 - a_2)}{2} \right) \leq \log_p \left(\frac{a_3(a_3 - a_2)}{2} \right) \leq \log_p \binom{d}{2}. \quad (4)$$

Thanks to Theorem 2.3, $\text{ord}_p a_2 = 0$ (i.e., the leading coefficient of Q not being divisible by p) implies that all the roots $\mu \in \mathbb{C}_p$ of Q have nonnegative valuation. So then $\text{ord}_p(1 - \mu) \geq 0$ and, thanks to Bound (4), we obtain $\text{ord}_p(1 - \mu) \leq \log_p \binom{d}{2} < \log_p \frac{(d-1)d^3 \cdot d}{8}$. (Note that the coefficients of q have absolute value at most $d = a_3$.) So we may assume $\sigma := \text{ord}_p a_2 > 0$ henceforth.

Since $\gcd(a_2, a_3) = 1$ we must have $\text{ord}_p a_3 = 0 = \text{ord}_p(a_3 - a_2)$. Theorem 2.3 applied to q then tells us that Q has exactly $a_3 - a_2$ roots in \mathbb{C}_p of p -adic valuation $-\frac{\sigma}{a_3 - a_2}$, and

exactly $a_2 - 2$ roots $\mu \in \mathbb{C}_p$ of p -adic valuation 0, since $q(x) = Q(x)(x-1)^2$. In particular, $\text{ord}_p(1 - \mu) = -\frac{\sigma}{a_3 - a_2} \geq -\log_p(d-1)$ on the set of roots with negative valuation, and $\text{ord}_p(1 - \mu) \geq 0$ at the roots $\mu \in \mathbb{C}_p$ with $\text{ord}_p \mu = 0$.

Equality (4) then implies that each of the $a_3 - 2$ roots $\mu \in \mathbb{C}_p$ of Q with $\text{ord}_p \mu = 0$ must satisfy $\text{ord}_p(1 - \mu) = (a_3 - a_2) \frac{\sigma}{a_3 - a_2} + \text{ord}_p\left(\frac{a_3(a_3 - a_2)}{2}\right) = \text{ord}_p\left(\frac{a_2 a_3 (a_3 - a_2)}{2}\right) \leq \log_p\left(\frac{a_3 a_2 (a_3 - a_2)}{2}\right)$. By the Arithmetic Geometric Inequality, $a_2(a_3 - a_2) \leq a_3^2/4$, so we arrive at $\text{ord}_p(1 - \mu) \leq \log_p(d^3/8) < \log_p((d-1)d^3 \cdot d/8)$ and we have proved Assertion (5) in the special case $f(x) = q(x)$.

A direct computation via Assertion (1) of our lemma then yields $f(x) = \frac{c_1}{(a_3 - a_2)\tau^2} q(x/\tau)$ identically. So the roots of f are simply scalings of the roots of q by a factor τ . Since $f'(\tau) = 0$, Theorem 2.3 implies that $\text{ord}_p \tau = \frac{\text{ord}_p(a_2 c_2) - \text{ord}_p(a_3 c_3)}{a_3 - a_2}$, which clearly lies in the closed interval $[-\log_p(dH), \log_p((d-1)H)]$. So then $\text{ord}_p(\tau - \zeta) = \text{ord}_p \tau + \text{ord}_p(1 - \mu)$ for some root $\mu \in \mathbb{C}_p$ of Q . In other words, $\text{ord}(\tau - \zeta) \leq \log_p((d-1)Hd^3/8) = \log_p((d-1)d^3 H/8)$. ■

Assertion (1) of Lemma 5.4 tells us that degenerate roots in \mathbb{C}_p^* of trinomials satisfy binomial equations with well-bounded coefficients. Our earlier Algorithms 2.22 and 2.24 (for solving binomial equations) thus imply that degenerate roots of trinomials are easy to approximate. Our final step in proving Theorem 1.6 will be estimating the spacing of non-degenerate roots in \mathbb{C}_p for trinomials having degenerate roots in \mathbb{C}_p .

5.1. Completing the proof of Theorem 1.6: root spacing in the face of degeneracy

First note that we may assume $\zeta_1 \zeta_2 \neq 0 \neq f(0)$, since this initial reduction to nonzero roots (from the proof of the square-free case in Section 3) does not require f to be square-free. Note also that Proposition 2.4 and Assertion (5) of Lemma 5.4 tells us that our sharper lower bound holds if at least one ζ_i is a degenerate root. So we may assume that ζ_1 and ζ_2 are both non-degenerate roots. Furthermore, letting $r := \text{gcd}(a_2, a_3)$, we can reduce to special case $r = 1$ via the same argument as from the proof of Assertion (5) of Lemma 5.4. So we will also assume $\text{gcd}(a_2, a_3) = 1$.

Our proof then follows almost exactly the format of the square-free case, with just two small changes: (a) We replace f by the polynomial $F(x) := \frac{f(x)}{(x-\tau)^2}$, where $\tau \in \mathbb{Q}$ is the unique degenerate root of f . (That f has exactly one degenerate root, and it has multiplicity 2, follows from Assertions (1) and (2) of Lemma 5.4.) (b) We replace Lemma 3.3 by a direct proof that $|F(\tau)|_p \geq e^{-O(\log(dH))}$.

To prove the last bound, observe that $F(\tau) = \frac{c_1}{(a_3 - a_2)\tau^2} Q(1)$. Since $\text{ord}_p \tau = \frac{\text{ord}_p(a_2 c_2 / (a_3 c_3))}{a_3 - a_2}$, Assertion (3) of Lemma 5.4 then tells us that

$$\text{ord}_p F(\tau) \leq \log_p(H) + \log_p(dH) + \log_p O(a_2^2 a_3^3 (a_3 - a_2)^2) = O(\log_p(dH)). \quad \blacksquare$$

6. Solving trinomials over \mathbb{Q}_p

Unlike the binomial case (see Remark 2.20), the tree $\mathcal{T}_{p,k}(f)$ can have depth $\Omega(\log_p(dH))$ or greater for a trinomial $f \in \mathbb{Z}[x]$ with $p \nmid f(0)$ and k sufficiently large [28]. However, Lemma 6.1 below will show that the structure of $\mathcal{T}_{p,k}(f)$ is still simple: *No* subtree of $\mathcal{T}_{p,k}(f)$ emanating from a vertex of depth ≥ 1 has more than 2 vertices of out-degree more than 2. Corollary 6.6 below will establish how large k must be so that $\mathcal{T}_{p,k}(f)$ is deep enough to encode (via Lemma 2.18) all the non-degenerate roots of f in \mathbb{Z}_p , and do so with sufficient accuracy for Newton iteration to converge quickly. Our estimates on k will enable us to approximate all the roots of f in \mathbb{Q}_p in time $p^{3+o(1)} \log^{4+o(1)}(dH) \log_p^3 d$. Mild assumptions on the exponents of f can also guarantee that the root node of $\mathcal{T}_{p,k}(f)$ has $O(\sqrt{p})$ or even fewer children, and the presence of degenerate roots in \mathbb{Q}_p^* for f enables even tighter estimates for k . Each of these restrictions leads to speed-ups we will describe.

6.1. Trees and trinomials

Lemma 6.1. *Suppose $f(x) = c_1 + c_2 x^{a_2} + c_3 x^{a_3} \in \mathbb{Z}[x]$ is a trinomial of degree $d = a_3 > a_2 \geq 1$, with all its coefficients having absolute value at most H . Then every non-root nodal polynomial $f_{i,\zeta}$ of $\mathcal{T}_{p,k}(f)$ with $\zeta \neq 0 \pmod{p}$ 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 6.2. 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^* :

$$0 + 1 \cdot 2 + 0 \cdot 2^2 + 1 \cdot 2^3 + O(2^4), \quad 0 + 1 \cdot 2 + 1 \cdot 2^2 + O(2^5),$$

$1 + 0 \cdot 2 + 0 \cdot 2^2 + \dots$, $1 + 0 \cdot 2 + 1 \cdot 2^2 + O(2^5)$, $1 + 1 \cdot 2 + 0 \cdot 2^2 + 2^3 + O(2^4)$, and $1 + 1 \cdot 2 + 1 \cdot 2^2 + O(2^3)$. This is because $\tilde{f}_{2,2} = \tilde{f}_{2,1} = \tilde{f}_{2,1+2} = x^2 + x$ and each of these (terminal) nodal polynomials has exactly 2 non-degenerate roots in \mathbb{F}_2 , each of which lifts to a unique root in \mathbb{Z}_2 . Note that $\tilde{f}_{1,1}(x) = x^4 + x^2$ has degree 4, and corresponds to the unique depth 1 vertex with 2 children. \diamond

Example 6.3. Composing Example 2.12 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 6.1 we will need a powerful result of Lenstra [40] 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 6.4. [40, 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 $(k+1, m, \nu(x-1))$ in the notation of [40], and the parameter r there is set to 1 in our application here.

Proof of Lemma 6.1. 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 relabeling the ζ in the subscripts of the $f_{i,\zeta}$) the tree $\mathcal{T}_{p,k}(f)$. So to study \tilde{f}_{1,ζ_0} with $\zeta_0 \in \{1, \dots, p-1\}$, it suffices to study $\tilde{f}_{1,1}$.

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 6.4 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

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

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 $\deg f_{1,1} \leq r$.

By Lemma 2.17, 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 $(*)$ 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 $(*)$ 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 [40, Prop. 2.4], we have $\text{ord}_p d_2(r) \leq \frac{2 \log r}{\log p}$. Note that, for any fixed p , the quantity $r - \frac{2 \log r}{\log p}$ is an increasing function of r for $r \geq \frac{2}{\log p}$. Furthermore, $\left\lceil 7 - \frac{2 \log 7}{\log p} \right\rceil \geq 2$ for all $p \geq 2$ and $\left\lceil 5 - \frac{2 \log 5}{\log p} \right\rceil \geq 3$ 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 6.5. Lemma 2.19, on binomials, can now be proved by modifying the proof above slightly: We replace Inequality $(*)$ by $r - \text{ord}_p d_1(r) \leq 1$, replace $d_2(r)$ with $d_1(r)$, and let $\varepsilon_p = 1$ for all p . In particular, the definition of $s(f, \zeta_0)$ tells us that

$$s(f, \zeta_0) \leq 1 + \text{ord}_p f'(\zeta_0) = 1 + \text{ord}_p d = 1 + \ell. \quad \diamond$$

It seems harder to get an upper bound on $s(f, \zeta_0)$ for trinomials than binomials. Nevertheless, we can derive a bound quadratic in $\log d$ and linear in $\log H$, and thereby estimate how large k must be for our tree $\mathcal{T}_{p,k}(f)$ to be deep enough for trinomial root approximation.

Corollary 6.6. *Suppose $f(x) = c_1 + c_2 x^{a_2} + c_3 x^{a_3} \in \mathbb{Z}[x]$ has degree d , $0 < a_2 < a_3$, $p \nmid c_1$, $c_2 c_3 \neq 0$, and $|c_i| \leq H$ for all i . Let $r := \gcd(a_2, a_3)$, define S_0 to be the maximum of $s(f, \zeta_0)$ (see Definition 2.9) for any $\zeta_0 \in \{1, \dots, p-1\}$ satisfying $f(\zeta_0) = f'(\zeta_0) = 0 \bmod p$, and set $S_0 := 0$ should there be no such ζ_0 . Also let D be the maximum of $\text{ord}_p(\zeta - \xi)$ over all distinct non-degenerate roots $\zeta, \xi \in \mathbb{Z}_p$ of f (if f has at least 2 non-degenerate roots in \mathbb{Z}_p) or 0 (if f has 1 or fewer non-degenerate roots in \mathbb{Z}_p); and define M_p to be 4, 3, or 2, according as p is 2, 3, or ≥ 5 . Then:*

1. $k \geq 1 + S_0 \min\{1, D\} + M_p \max\{D - 1, 0\} \implies$ the depth of $\mathcal{T}_{p,k}(f)$ is at least D .
2. $a_2 = 1 \implies S_0 < \log_p(p^2 d^2 H^2)$.
3. $d \geq 3 \implies S_0 = O(p \log(\frac{d}{r}) \log_p(\frac{dH}{r}))$.
4. f has a degenerate root in $\mathbb{C}_p \implies S_0 \leq \log_p(p^2 d r)$.
5. The lower bound for k from Assertion (1) can be attained for $k = O(p \log_p^2(dH) \log d)$ or $k = O(\log_p(dH))$, according as f has no degenerate roots in \mathbb{C}_p , or at least one such root.

Remark 6.7. Note that $d \geq 2$ for any trinomial, and $d = 2$ implies $a_2 = 1$ above. One should also remember that Theorem 1.6 provides an explicit upper bound for D . \diamond

Proof of Corollary 6.6. Assertion (1): $\mathcal{T}_{p,k}(f)$ always includes a root node by definition, so the case $D = 0$ is trivial and we assume $D \geq 1$.

Our lower bound on k then follows easily from Lemma 6.1: Since f has distinct non-degenerate roots $\zeta, \xi \in \mathbb{Z}_p$ with $\text{ord}(\zeta - \xi) \geq 1$ by assumption, this means that $\zeta = \xi \bmod p$ and thus \tilde{f} must have a degenerate root $\zeta'_0 \in \{1, \dots, p-1\}$ (since $p \nmid c_1$). Having $k \geq 1 + S_0$ then simply allows the root node to have maximally many child nodes (and thus depth ≥ 1), thanks to Definition 2.9. Furthermore, thanks to Lemma 6.1, the summand $M_p \max\{D - 1, 0\}$ simply guarantees that $\mathcal{T}_{p,k}$ has depth D and that $\mathcal{T}_{p,k}(f)$ has maximally many nodes at depth $\leq D$. (Note that for any nodal polynomial $f_{i,\zeta'}$ with $i \geq 1$, we have that $s(f_{i,\zeta'}, \zeta_i)$ is bounded from above by 4, 3, or 2, according as p is 2, 3, or ≥ 5 , thanks to Lemma 2.17.) In particular, we see that any k satisfying our lower bound yields a k satisfying all the assumptions of Lemma 2.18. \blacksquare

Assertion (2): Immediate from $s(f, \zeta_0) \leq 2 + \text{ord}_p \frac{f''(\zeta_0)}{2}$ (thanks to the definition of $s(\cdot, \cdot)$ as a minimum), $f''(\zeta_0) = d(d-1)c_3\zeta_0^{d-2}$, and $\text{ord}_p \zeta_0 = 0$. \blacksquare

Note. We now temporarily assume that $\gcd(a_2, a_3) = 1$, to simplify the proofs of Assertions (3) and (4), and show later how to reduce the case $\gcd(a_2, a_3) > 1$ to the case $\gcd(a_2, a_3) = 1$. \diamond

Assertion (3): First note that we must have $p \nmid c_2$ or $p \nmid c_3$ in order for \tilde{f} to have a root in \mathbb{F}_p^* .

Since $f'(\zeta_0) = a_2 c_2 \zeta_0^{a_2-1} + a_3 c_3 \zeta_0^{a_3-1} = 0 \pmod{p}$, and $\gcd(a_2, a_3) = 1$, we see that $p|a_2 \implies \text{ord}_p c_3 = \text{ord}_p a_2 > 0$ and $p \nmid a_3 c_2$. In which case, $\text{ord}_p f'(\zeta_0) = \text{ord}_p(a_2 c_2) + \text{ord}_p\left(1 - \frac{-a_3 c_3}{a_2 c_2} \zeta_0^{a_3-a_2}\right)$, and then we can bound $\text{ord}_p f'(\zeta_0)$ from above by the $n = 2$ case of Yu's Theorem if the second valuation is *not* ∞ . Should this valuation be ∞ , then we can instead apply the $n = 2$ case of Yu's Theorem to $\text{ord}_p f''(\zeta_0) = \text{ord}_p(a_2(a_2-1)c_2) + \text{ord}_p\left(1 - \frac{-a_3(a_3-1)c_3}{a_2(a_2-1)c_2} \zeta_0^{a_3-a_2}\right)$, since $\frac{a_3-1}{a_2-1} \neq 1$. So we obtain a bound of $S_0 < 2 + 2\text{ord}_p(r) + \log_p\left(\frac{d}{r}\left(\frac{d}{r}-1\right)H\right) + \log(2)\log(4)2^{57/2}e^6p\log\left(\frac{d}{r}\left(\frac{d}{r}-1\right)H\right)$ directly from Theorem 3.1, and the fact that $s(f, \zeta_0) \leq \min\{1 + \text{ord}_p f'(\zeta_0), 2 + \text{ord}_p f''(\zeta_0)\}$.

Similarly, $p|a_3 \implies \text{ord}_p c_2 = \text{ord}_p a_3 > 0$ and $p \nmid a_2 c_3$. In which case, $\text{ord}_p f'(\zeta_0) = \text{ord}_p(a_3 c_3) + \text{ord}_p\left(1 - \frac{-a_2 c_2}{a_3 c_3} \zeta_0^{a_2-a_3}\right)$, and we proceed in the same way as the last paragraph.

So let us now assume $p \nmid a_2 a_3$. Then $f'(\zeta_0) = 0 \pmod{p} \implies p \nmid c_2 c_3$, since $\text{ord}_p \zeta_0 = 0$ and p can not divide both c_2 and c_3 . So then we again attain the same bound as in the last two paragraphs. ■

Assertion (4): Note that $p \nmid c_1$ implies that any degenerate root $\tau \in \mathbb{C}_p$ of f must be nonzero. Lemma 5.4 then tells us that τ is the only degenerate root of f in \mathbb{C}_p and $\tau \in \mathbb{Q}_p^*$. Moreover, from the proof of Lemma 5.4, we have $f(\tau x) = \frac{c_1}{(a_3-a_2)\tau^2}q(x)$ identically and $\text{ord}_p \tau = \frac{\text{ord}_p(a_2 c_2) - \text{ord}_p(a_3 c_3)}{a_3-a_2}$. (Recall that $q(x) = (a_3-a_2) - a_3 x^{a_2} + a_2 x^{a_3}$ has 1 as its unique degenerate root in \mathbb{C}_p .)

Now, we must have $p \nmid c_2$ or $p \nmid c_3$ in order for there to be any roots at all for \tilde{f} .

Sub-Case $p \nmid c_2$. If τ has negative valuation, then we must have $p|c_3$ by Theorem 2.3. Also, $f'(\zeta_0) = \zeta_0^{a_2-1}(c_2 a_2 + c_3 a_3 \zeta_0^{a_3-a_2}) = 0 \pmod{p} \implies p|a_2$ since $p \nmid c_2$. Since $\text{ord}_p \tau = \frac{\text{ord}_p(a_2 c_2) - \text{ord}_p(a_3 c_3)}{a_3-a_2} < 0$ by assumption, we must have $\text{ord}_p(a_3 c_3) > \text{ord}_p(a_2 c_2)$ and thus $\text{ord}_p f'(\zeta_0) = \text{ord}_p(c_2 a_2) = \text{ord}_p(a_2)$. In other words, $\text{ord}_p \tau < 0 \implies S_0 \leq 1 + \text{ord}_p(a_2)$.

So let us now assume $\text{ord}_p \tau = 0$. Then by our identity $f(\tau x) = \frac{c_1}{(a_3-a_2)\tau^2}q(x)$, and the fact that $\tau \in \mathbb{Q}^*$ (via Assertion (1) of Lemma 5.4), the vector of coefficient valuations for f and the vector of coefficient valuations for q differ by a multiple of $(1, 1, 1)$. So our assumptions that $p \nmid c_1 c_2$ and $\gcd(a_2, a_3) = 1$ imply that $p \nmid (a_3-a_2)a_3$. So then, a_3-a_2 is invertible mod p and, by the rescaling between f and q , we have that \tilde{f} and \tilde{q} share the same value of S_0 (as well as the same number of degenerate roots in $\{1, \dots, p-1\}$). So let us now work with q instead, and assume for the remainder of this sub-case that ζ_0 is a degenerate root of $\tilde{q} \pmod{p}$.

If $p|a_2$ then $\text{ord}_p q'(\zeta_0) = \text{ord}_p(a_2) + \text{ord}_p(-1 + \zeta_0^{a_3-a_2})$ (since $p \nmid a_3$). Also, $\text{ord}_p q''(\zeta_0) = \text{ord}_p(a_2) + \text{ord}_p(-a_2 + a_3 \zeta_0^{a_3-a_2} - (-1 + \zeta_0^{a_3-a_2}))$. Since $p|a_2$ and $p \nmid a_3$, we see that $\text{ord}_p(-1 + \zeta_0^{a_3-a_2}) > 0$ implies that $\text{ord}_p q''(\zeta_0) = \text{ord}_p a_2$. On the other hand, if $\text{ord}_p(-1 + \zeta_0^{a_3-a_2}) = 0$, then $\text{ord}_p q'(\zeta_0) = \text{ord}_p a_2$ from our earlier formula for $\text{ord}_p q'(\zeta_0)$. So by the definition of $s(\cdot, \cdot)$, we obtain $S_0 \leq 2 + \text{ord}_p a_2$.

To conclude, $p \nmid a_2$, combined with our earlier conclusion that $p \nmid (a_3 - a_2)a_3$, implies that $\zeta_0 = 1$, thanks to Assertion (1) of Lemma 5.4. In which case, $q'(1) = 0$ but $q''(1) = a_2a_3((a_3 - 1) - (a_2 - 1)) = a_2a_3(a_3 - a_2)$ and thus $S_0 \leq 2$.

Sub-Case $p \nmid c_3$. Here, we must have $\text{ord}_p \tau = 0$ and thus $\text{ord}_p(a_2c_2) = \text{ord}_p(a_3c_3)$ by our earlier formula for $\text{ord}_p \tau$. In particular, we must have $\text{ord}_p(a_2c_2) = \text{ord}_p a_3$ since $p \nmid c_3$. Note also that $p|a_2$ thus implies $p|a_3$, which would contradict $\gcd(a_2, a_3) = 1$. So we must also have $p \nmid a_2$ and thus $\text{ord}_p c_2 = \text{ord}_p a_3$. Since we already proved the Sub-Case $p \nmid c_2$, let us now assume $p|c_2$ (and thus $p|a_3$).

By our identity $f(\tau x) = \frac{c_1}{(a_3 - a_2)\tau^2}q(x)$, and the fact that $\tau \in \mathbb{Q}^*$ (via Assertion (1) of Lemma 5.4), the vector of coefficient valuations for f and the vector of coefficient valuations for q differ by a multiple of $(1, 1, 1)$. So our assumptions that $p \nmid c_1c_3$ and $\gcd(a_2, a_3) = 1$ imply that $p \nmid (a_3 - a_2)a_2$. So then, $a_3 - a_2$ is invertible mod p and, by the rescaling between f and q , we have that \tilde{f} and \tilde{q} share the same value of S_0 (as well as the same number of degenerate roots in $\{1, \dots, p-1\}$). So let us now work with q instead, and assume now that ζ_0 is a degenerate root of \tilde{q} mod p .

Observe then that $\text{ord}_p q'(\zeta_0) = \text{ord}_p(a_3) + \text{ord}_p(-1 + \zeta_0^{a_3 - a_2})$ (since $p \nmid a_2$). Also, $\text{ord}_p q''(\zeta_0) = \text{ord}_p(a_3) + \text{ord}_p(-a_2 + a_3\zeta_0^{a_3 - a_2} - (-1 + \zeta_0^{a_3 - a_2}))$. Since $p|a_3$ and $p \nmid a_2$, we see that $\text{ord}_p(-1 + \zeta_0^{a_3 - a_2}) > 0$ implies that $\text{ord}_p q''(\zeta_0) = \text{ord}_p a_3$. On the other hand, if $\text{ord}_p(-1 + \zeta_0^{a_3 - a_2}) = 0$, then $\text{ord}_p q'(\zeta_0) = \text{ord}_p a_3$ from our earlier formula for $\text{ord}_p q'(\zeta_0)$. So by the definition of $s(\cdot, \cdot)$, we obtain $S_0 \leq 2 + \text{ord}_p a_3$. ■

Extending to $\gcd(a_2, a_3) > 1$. To complete our proofs of Assertions (3) and (4) let us assume $r := \gcd(a_2, a_3) > 1$ and recall that \bar{f} is the unique polynomial in $\mathbb{Z}[x]$ satisfying $f(x) = \bar{f}(x^r)$ identically. Clearly then, $\deg \bar{f} = \frac{\deg f}{r}$ and any root $\tau \in \mathbb{C}_p$ of f induces a root τ^r of \bar{f} . Furthermore, \bar{f} having a degenerate root $\zeta_0 \in \{1, \dots, p-1\}$ clearly implies that the mod p reduction of \bar{f} has μ_0 as a degenerate root, where $\mu_0 \in \{1, \dots, p-1\}$ is the mod p reduction of ζ_0^r . The Chain Rule then implies $\text{ord}_p f'(\zeta_0) = \text{ord}_p(r) + \text{ord}_p \bar{f}'(\mu_0) \leq \log_p(r) + \text{ord}_p \bar{f}'(\mu_0)$.

Should $f'(\zeta_0)$ vanish identically, then Assertion (1) of Lemma 5.4 easily implies that all the degenerate roots of f have multiplicity 2 and thus $f''(\zeta_0)$ can not vanish. In which case, via the Chain Rule again, $\text{ord}_p f''(\zeta_0) = 2\text{ord}_p(r) + \text{ord}_p \bar{f}''(\mu_0) \leq 2\log_p(r) + \text{ord}_p \bar{f}''(\mu_0)$. So our general formula follows immediately from the case $r = 1$, which we've already proved. ■

Assertion (5): Immediate from Assertions (3) and (4), and Theorem 1.6. ■

6.2. Building trees efficiently

It is easy to see that the only degenerate root the quadratic trinomial $c_1 + c_2x + x^2 \in \mathbb{Z}[x]$ can have mod p is exactly $-c_2/2$ when $p \geq 3$. (For $p = 2$ it is clear that the only monic degenerate quadratics are $x^2 + 1$ and x^2 , with respective degenerate roots 1 and 0.) It will be useful to have a similar statement for trinomials with $(p, d) \in \{2, 3\} \times \{3, 4\}$.

Proposition 6.8. Suppose $f(x) = c_0 + c_1x + c_2x^2 + c_3x^3 + c_4x^4 \in \mathbb{Z}[x]$ has degree $d \geq 2$, and $|c_i| \leq H$ for all i . Then:

0. The discriminant of f can be evaluated in time $O(\log(\max\{p, H\}) \log \log \max\{p, H\})$.
1. When $p \leq 3$ we can find all the degenerate roots of f in \mathbb{F}_p (or correctly declare there none) in time $O(\log H)$. In particular, f has at most 1 (resp. 2) degenerate root(s) in \mathbb{F}_p , according as $d \leq 3$ or $d = 4$.
2. For any prime p we can find all the non-degenerate roots of f (or correctly declare there are none) in deterministic time $O(p^{1/2} \log^2 p)$.

Proof. Assertion (0) follows from the definitions of the quartic, cubic, and quadratic discriminants (see, e.g., [29, Ch. 12]), Theorem 2.7, and the fact that evaluating $\Delta_d(f)$ reduces to evaluating a 7×7 , 5×5 , or 3×3 determinant in the coefficients of f (followed by division by the leading coefficient of f), after reducing the coefficients mod p .

For Assertion (1), first note that $p \leq 3$ implies that we can reduce the coefficients of f and f' mod p in time $O(\log H)$ thanks to Theorem 2.7. We can then simply use brute-force (over a search space with at most 3 elements!) to find all the degenerate roots of f in time $O(1)$. In particular, since any degenerate root must have multiplicity ≥ 2 , the only way f can have more than 1 degenerate root is for $d = 4$, in which case there can be no more than 2 degenerate roots. For instance, $x^4 + x^2 + 1$ (resp. $x^4 + x^2$) has degenerate roots $\{\pm 1\} \in \mathbb{F}_3$ (resp. $\{0, 1\} \in \mathbb{F}_2$).

Assertion (2) follows immediately from Shoup's deterministic algorithm for factoring arbitrary univariate polynomials over a finite field [55], upon specializing to degree ≤ 4 . ■

Lemma 6.9. For any trinomial $f(x) = c_1 + c_2x^{a_2} + c_3x^{a_3} \in \mathbb{Z}[x]$ of degree d , with $p \nmid c_1$, $0 < a_2 < a_3$, and $|c_i| \leq H$ for all i , let ν denote the number of degenerate roots of \tilde{f} in \mathbb{F}_p^* and let \mathcal{D} denote the depth of $\mathcal{T}_{p,k}(f)$. Then $\mathcal{T}_{p,k}(f)$ has $\leq 1 + (2\mathcal{D} - 1)\nu$ nodes; and we can compute the mod p reductions of all the nodal polynomials $f_{i,\zeta}$ of $\mathcal{T}_{p,k}(f)$, as well as all the values of the $s(f_{i-1,\mu}, \zeta_{i-1})$, in deterministic time

$$(p \log(d) + k\nu\mathcal{D} \log(p) \log(d) \log H)^{1+o(1)}.$$

Proof. By Lemma 6.1, all non-root nodal polynomials have mod p reduction of degree no greater than 4. Thus, the root node of $\mathcal{T}_{p,k}(f)$ has $\leq \nu$ ($\leq p - 1$) children, and any node at depth ≥ 1 has no more than 2 children (since a polynomial of degree ≤ 4 has ≤ 2 degenerate roots). Lemma 2.17 also tells us that $\deg \tilde{f}_{i,\mu+\zeta_{i-1}p^{i-1}}$ is at most the multiplicity of $\zeta_{i-1} \in \mathbb{F}_p^*$ as a root of $\tilde{f}_{i-1,\mu}$. So any node v that has an ancestor at level ≥ 1 with 2 children can have no more than 1 child. Thus, there can be no more than 2ν nodes at depth $i \geq 2$. It is then clear that $\mathcal{T}_{p,k}(f)$ has at most $1 + (2\mathcal{D} - 1)\nu$ nodes.

We now check whether \tilde{f} has any degenerate roots in \mathbb{F}_p : By assumption, they must lie in \mathbb{F}_p^* . Also, should $p|c_3$, \tilde{f} would be a binomial and thus have degenerate roots in \mathbb{F}_p^*

only if $p|a_2$; in which case any root of \tilde{f} in \mathbb{F}_p^* is degenerate. We can then decide if there are degenerate roots simply by checking whether $(-c_1/c_2)^{(p-1)/\gcd(a_2,p-1)} = 1 \pmod{p}$, which can be done in time $O(\log(dH) \log(\log(dH)) + \log^2(p) \log \log p)$ via Theorem 2.7. Should there be any degenerate roots, there will then be exactly $\gcd(a_2, p-1)$ many, and we can then find them in time no worse than $O((p + \log d) \log(dp) \log(\log(dp)) + \log(H) \log(pH) \log \log(pH))$ via brute-force (much like our earlier complexity analysis of Steps 5–7 of Algorithm 2.22).

So let us assume $p \nmid c_3$. Note that $p \nmid \gcd(a_2, a_3) \implies$ every root of \tilde{f} in \mathbb{F}_p^* is degenerate, in which case we can simply find all these roots first by reducing the coefficients (resp. exponents) of $\tilde{f} \pmod{p}$ (resp. $\pmod{p-1}$) in time

$$O(\log(\max\{d, p\}) \log(\log \max\{d, p\}) + \log(\max\{p, H\}) \log \log \max\{p, H\})$$

and then applying brute-force search in time $O(p \log^2(p) \log \log p)$. So let us assume $p \nmid \gcd(a_2, a_3)$. Observe then that \tilde{f} has degenerate roots in $\mathbb{F}_p^* \iff \Delta_{\text{tri}}(\tilde{f}) = 0 \pmod{p}$, thanks to Assertion (1) of Lemma 5.2. In particular, by Theorem 2.7, $\Delta_{\text{tri}}(\tilde{f})$ can be computed \pmod{p} in time $O(\log(\max\{d, p\}) \log(\log \max\{d, p\}) + \log(\max\{H, p\}) \log \log \max\{H, p\})$ (to reduce the exponents of $\Delta_{\text{tri}}(\tilde{f}) \pmod{p-1}$ and the power bases \pmod{p}) plus $O(\log^2(p) \log \log p)$ to compute the monomials of $\Delta_{\text{tri}}(\tilde{f})$. If $\Delta_{\text{tri}}(\tilde{f}) \neq 0 \pmod{p}$ then we know \tilde{f} has no degenerate roots and then $\mathcal{T}_{p,k}(f)$ is simply a single root node. Otherwise, let $r' := \gcd(a_2, a_3, p-1)$ and apply the Extended Euclidean Algorithm (in time $O(\log(p) \log^2 \log p)$ via Theorem 2.7) to $a_2 \pmod{p-1}$ and $a_3 \pmod{p-1}$ to find $\alpha, \beta \in \mathbb{Z}$ with logarithmic height $O(\log p)$ such that $\alpha(a_2 \pmod{p-1}) + \beta(a_3 \pmod{p-1}) = r'$. Assertion (1) of Lemma 5.4 then tells us that the degenerate roots of \tilde{f} in \mathbb{F}_p^* are exactly the roots of $g(x) := x^{r'} - (-1)^\alpha \left(\frac{c_1}{a_3 - a_2}\right)^{\alpha+\beta} \left(\frac{a_3}{c_2}\right)^\alpha \left(\frac{a_2}{c_3}\right)^\beta$ in \mathbb{F}_p^* . Lemmata 2.5 and 2.6 and Theorem 2.7 then easily imply that deciding whether g has any roots in \mathbb{F}_p^* takes time $O(\log^2(p) \log \log p)$, and there are exactly r' many degenerate roots in \mathbb{F}_p^* if so. Just as in the last paragraph, we can then apply brute-force to g in time

$$O((p + \log d) \log(dp) \log(\log(dp)) + \log(H) \log(pH) \log \log(pH))$$

to find all the degenerate roots of \tilde{f} in \mathbb{F}_p^* .

Assuming \tilde{f} has degenerate roots in \mathbb{F}_p^* , let us now see how to compute the child nodes of the root node in $\mathcal{T}_{p,k}(f)$: First note that the coefficient of x^i in the monomial term expansion of $c(\mu+px)^a \pmod{p^j}$ (for $i \leq j$) is simply $c \binom{a}{i} \mu^{a-i} p^i \pmod{p^j}$. Also, Lemma 2.17 tells us that $f_{i,\zeta}(x) = p^{-s} f(\mu + p^i x) \pmod{p^j}$ for suitable (s, μ, j) . Putting this together, this means we can compute $s(f, \zeta_0)$ and \tilde{f}_{1,ζ_0} (for all degenerate roots $\zeta_0 \in \mathbb{F}_p^*$ of \tilde{f}) by evaluating $\zeta_0^{a_2}$ and $\zeta_0^{a_3} \pmod{p^k}$, $\binom{a_2}{i}$ and $\binom{a_3}{i}$ for $i \in \{0, 1, 2\}$ if $p \geq 5$, and $O(1)$ additional ring operations in $\mathbb{Z}/(p^k)$. (We instead take $i \in \{0, 1, 2, 3\}$ or $\{0, 1, 2, 3, 4\}$ according as p is 3 or 2.) Via Recursive Squaring (a.k.a. the Binary Method [7, pp. 102–103]), Theorem 2.7 tells us that we can compute the a_2 nd and a_3 rd powers of all the degenerate roots $\zeta_0 \in \mathbb{F}_p^*$ in time $O(v \cdot \log(d) \cdot k \log(p) \log(k \log p))$, and the remaining

operations are negligible in comparison. In particular, each $s(f, \zeta_0)$ can be computed by bisection and the resulting complexity is also negligible compared to the preceding O -estimate.

So in summary, all computations necessary to find all child nodes of the root node take time no greater than

$$O((p \log(p) + \log d) \log(dp) \log \log(dp) + p \log^2(p) \log(\log p) + \log(H) \log(dpH) \log \log(dpH) + \nu k \log(d) \log(p) \log(k \log p)).$$

Having computed all the mod p reductions of the nodal polynomials \tilde{f}_{1,ζ_0} at depth 1, we then proceed inductively, performing almost the same calculations as in the last two paragraphs. The only difference, assuming $p \geq 5$, is then applying the quadratic discriminant (instead of the trinomial discriminant) to detect and find the *sole* degenerate root of $f_{i-1,\mu}$ (for $i \in \{2, \dots, k-1\}$), should there be one. (Should $p \in \{2, 3\}$ then we simply apply Proposition 6.8 instead, and possibly have two degenerate roots in the worst case when $p = 2$.) This eliminates the need for brute-force search, and gives us an improved complexity bound of $O(k \log(p) \log(k \log p) \log(d) + \log H)$ to compute the children (no more than two) of each $f_{i-1,\mu}$.

Summing all the resulting complexity estimates over all $O(\nu \mathcal{D})$ children, and overestimating slightly, we obtain our stated bound. ■

Corollary 6.10. *Following the notation of Lemma 6.9, we have the following improved complexity bounds for computing the mod p reductions of all the nodal polynomials of $\mathcal{T}_{p,k}(f)$ and their respective $s(\cdot, \cdot)$ values:*

1. *If we only wish to construct the sub-tree of $\mathcal{T}_{p,k}(f)$ corresponding to $\zeta_0 = 1$, and correctly declare whether 1 is a degenerate root of f :*

$$\text{Deterministic time } \mathcal{D}(k \log^2(p) \log(d) \log H)^{1+o(1)}.$$

2. *If the exponents are $\{0, a_2, a_3\}$ with $\gcd(a_2 a_3 (a_3 - a_2), (p-1)p) \leq 2$:*

$$\text{Deterministic time } p^{\frac{1}{2}+o(1)} + \mathcal{D}(k^2 \log(p) \log(d) \log H)^{1+o(1)}$$

$$\text{or Las Vegas randomized time } \mathcal{D}(k \log^2(p) \log(d) \log H)^{1+o(1)}.$$

Remark 6.11. While we state a randomized speed-up in Assertion (2) above, any asymptotic gains are unfortunately overwhelmed by the upper bounds on k and \mathcal{D} for the non-degenerate case from Corollary 6.6 and Theorem 1.6. Nevertheless, we state our bounds in a refined way above, should better bounds on k and \mathcal{D} become available in the future. ◇

Proof of Corollary 6.10. In what follows, we keep in mind the template of the proof of Lemma 6.9, and simply point out the key changes resulting in speed-ups.

Assertion (1): Here there is no need to search for roots of \tilde{f} : We merely evaluate \tilde{f} and \tilde{f}' at 1 to see if 1 is a degenerate root. This amounts to time

$$O(\log(\max\{d, p\}) \log(\log \max\{d, p\}) + \log(\max\{p, H\}) \log(\log \max\{p, H\}))$$

to reduce exponents mod $p-1$ and coefficients mod p , and then time $O(\log^2(p) \log \log(p))$ for the evaluation. At this point, we also know if 1 fails to be a degenerate root of \tilde{f} .

We then need time $O(\max\{k \log p, \log H\} \log \max\{k \log p, \log H\})$ to reduce the coefficients of \tilde{f} mod p^k , and then time $O(\log(d)k \log(p) \log(k \log p))$ to compute $s(f, 1)$ and the child node of the root node. For the remaining descendants, Lemma 6.1 tells us that there are at most 2 children, and any subsequent siblings can have no further offspring with more than one child. Also, as observed earlier, we can find the degenerate roots of the mod p reduction of any non-root nodal polynomial in time $O(\log H)$. So the remaining child nodes take time $\mathcal{D} - 1$ times $O(k \log(d) \log(p) \log(k \log p) + \log(H))$ to compute. ■

Assertion (2): The gcd assumption on the exponents implies there can be at most 2 degenerate roots for \tilde{f} in \mathbb{F}_p (and they are nonzero since we originally assumed $p \nmid c_1$ in Lemma 6.9): This follows from basic group theory if $p|c_3$ and via Lemma 5.4 if $p \nmid c_3$.

If $p|c_3$ then we can decide whether \tilde{f} has a degenerate root in \mathbb{F}_p^* by computing $g_1 := \gcd(\tilde{f}, x^{p-1} - 1)$ and checking whether $\deg g_1 \geq 1$ or not: If $\deg g_1 = 1$ then we can easily find the unique root of g_1 using ≤ 2 arithmetic operations in \mathbb{F}_p . If $\deg g_1 = 2$ then we can find the roots either in deterministic time $O(p^{1/2} \log^2 p)$ via Shoup's fast deterministic factoring algorithm [55], or Las Vegas time $\log^{2+o(1)} p$ via the fast randomized factorization algorithm of Kedlaya-Umans [34]. Furthermore, g_1 can be computed efficiently by first computing $x^{a_2} \bmod x^{p-1} - 1$ via Recursive Squaring (a.k.a. the Binary Method [7, pp. 102–103]), and then computing the rest of $\tilde{f} \bmod x^{p-1} - 1$. This entails $O(\log d)$ reductions (of exponents) mod $p-1$, along with 3 arithmetic operations in \mathbb{F}_p , meaning additional deterministic time $O(\log(d) \log(\max\{d, p\}) \log \log \max\{d, p\})$ via Theorem 2.7.

If $p \nmid c_3$ then we can decide whether \tilde{f} has a degenerate root in \mathbb{F}_p^* by first checking $\Delta_{\text{tri}}(\tilde{f}) \stackrel{?}{=} 0 \bmod p$, which takes time

$$O(\log(\max\{d, p\}) \log(\log \max\{d, p\}) + \log(\max\{H, p\}) \log \log \max\{H, p\})$$

(as already observed in our last proof). If this discriminant indeed vanishes mod p then we compute $g_2 := \gcd(\tilde{f}, \tilde{f}') = \gcd(\tilde{f}, \tilde{f}'/x^{a_2-1})$. Like g_1 , the polynomial g_2 has degree ≤ 2 , and it can be computed efficiently, along with its roots (if any) in deterministic time

$$O(p^{1/2} \log^2(p) + \log(d) \log(\max\{d, p\}) \log \log \max\{d, p\}),$$

or Las Vegas time

$$O(\log^{2+o(1)}(p) + \log(d) \log(\max\{d, p\}) \log \log \max\{d, p\}).$$

We then proceed as in the proof of Assertion (1), with at worst twice as many children. ■

6.3. The algorithm that proves Theorem 1.1

Recall that a *terminal* node of a tree is a node with no children.

Algorithm 6.12. (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 $O(p \log^2(dH) \log d)$ 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 $\text{ord}_p \frac{c_2^2}{c_1 c_3} \geq 0$ and $\text{ord}_p c_1 \neq \text{ord}_p c_3 \pmod{a_3}$ or

$\text{ord}_p \frac{c_2^2}{c_1 c_3} < 0$ and $\text{ord}_p c_1 \neq \text{ord}_p c_2 \pmod{a_2}$ 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: Decide, via gcd-free bases, $\Delta_{\text{tri}}(f) \stackrel{?}{=} 0$. If so, set $\delta := 1$. Otherwise, set $\delta := 0$.

4: If $\delta = 1$ then set $r' := \text{gcd}(a_2, a_3, p - 1)$ and, via the Extended Euclidean Algorithm, find $\alpha, \beta \in \mathbb{Z}$ with logarithmic height $O(\log p)$ such that $\alpha(a_2 \pmod{p-1}) + \beta(a_3 \pmod{p-1}) = r'$. Then, via Algorithm 2.22 (or its $p = 2$ version, Algorithm 2.24), output the 2 most

significant base- p digits of the roots of $g(x) := x^{r'} - (-1)^\alpha \left(\frac{c_1}{a_3 - a_2} \right)^{\alpha+\beta} \left(\frac{a_3}{c_2} \right)^\alpha \left(\frac{a_2}{c_3} \right)^\beta$ in \mathbb{Z}_p .

5: Set k to be the lower bound from Corollary 6.6 (employing the stated upper bound on S_0 , and the upper bound on D from Theorem 1.6, should S_0 or D not be known), and compute the mod p reductions $\tilde{f}_{i,\zeta}$ of all the nodal polynomials of $\mathcal{T}_{p,k}(f)$.

6: By computing $\deg \text{gcd}(\tilde{f}_{i,\zeta}, x^p - x)$ for the non-root nodal polynomials of $\mathcal{T}_{p,k}(f)$, and brute-force search over \mathbb{F}_p^* for \tilde{f} , determine which nodal polynomials have non-degenerate roots.

7: Output every non-degenerate root $\zeta_0 \in \mathbb{F}_p$ of \tilde{f} . Also output, for each non-root nodal polynomial $f_{i,\zeta}$ found in Step 6, the set $\{\zeta + p^i \zeta_i \mid \zeta_i \in \mathbb{F}_p \text{ and } \tilde{f}_{i,\zeta}(\zeta_i) = 0 \neq \tilde{f}'_{i,\zeta}(\zeta_i)\}$.

8: If $p \mid c_3$ 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 rescaling of the reciprocal polynomial f^* .

Remark 6.13. We point out that some of the approximate roots output by our algorithm above require the use of Newton iteration *applied to $f_{i,\zeta}$ (instead of f)*. This is clarified in our correctness proof below. ◇

Proof of Theorem 1.1. First note that the root 0 is trivially detected by checking whether the constant term c_1 is 0. So we may assume $c_1 \neq 0$ and focus on roots in \mathbb{Q}_p^* . Note also that the rescalings from Steps 2 and 8 (which are simply replacements of f with $p^{j_1} f(p^{j_2} x)$ for suitable $j_1, j_2 \in \mathbb{Z}$) result in a possible increase in the bit-sizes our

outputs, but this increase is $O(\log H)$ thanks to Theorem 2.3. So we focus on roots in \mathbb{Z}_p of valuation 0, and assume $p \nmid c_1$ and $\text{ord}_p(c_2) \text{ord}_p(c_3) = 0$.

Condition (1) (the logarithmic height bound for our approximate roots) then clearly holds thanks to Step 5 of our algorithm, the definition of $\mathcal{T}_{p,k}(f)$, Lemma 2.18, Theorem 1.6, and Corollary 6.6.

Condition (2) (on the convergence of the Newton iterates) follows easily from the definition of $f_{i,\mu}$. In particular, Lemma 2.17 tells us that $f_{i,\mu}(x) = p^{-s} f(\mu + p^i x) \bmod p^j$ for suitable (s, μ, j) , and thus a non-degenerate root $\zeta_i \in \mathbb{F}_p$ of $\tilde{f}_{i,\zeta}$ yields a root $\mu + p^i \zeta_i$ of $f \bmod p^{i+1}$. Moreover, by Hensel's Lemma, $z_0 := \zeta_i$ is an approximate root of $f_{i,\mu}$, meaning that the sequence $(\mu + p^i z_n)_{n \in \mathbb{N}}$ derived from the iterates $(z_n)_{n \in \mathbb{N}}$ coming from applying Newton iteration to $(f_{i,\mu}, z_0)$ satisfies $|\xi - (\mu + p^i z_n)|_p \leq \left(\frac{1}{p}\right)^{2^{n-1}} |\xi - (\mu + p^i z_0)|_p$, where $\xi \in \mathbb{Z}_p$ is some true (non-degenerate) root of f . From Lemma 2.18 (and our choice of k via Corollary 6.6) we know that *all* the non-degenerate roots of f can be recovered this way, and uniquely so.

Condition (3) on correctly counting the roots of f in \mathbb{Q}_p follows immediately from Steps 3–8. In particular, first note that Step 4 actually outputs approximations of all the degenerate roots of f in \mathbb{Z}_p . This is because, as we already saw in the third paragraph of the proof of Lemma 6.9, the binomial g vanishes exactly on the degenerate roots of f in \mathbb{F}_p^* and $\deg g = r'$ is exactly the number of these roots. Lemmata 2.5 and 2.6 and Theorem 2.7 then easily imply that deciding whether g has any roots in \mathbb{F}_p^* takes time $O(\log^2(p) \log \log p)$. So Step 4 correctly counts the degenerate roots in \mathbb{Q}_p thanks to our earlier work on Algorithms 2.22 and 2.24. Also, Corollary 6.6 and Lemma 2.18 tell us that the outputs from Step 7 are a collection of approximate roots that, en masse, converge to the set of non-degenerate roots of f in \mathbb{Z}_p of valuation 0, with no overlap. Step 8 then accounts for the remaining degenerate and non-degenerate roots in \mathbb{Q}_p .

The time complexity estimates from our theorem will follow from our complexity analysis of Algorithm 6.12 below. First, however, let us prove correctness for our algorithm.

Correctness: Via Theorem 2.3, Step 1 guarantees that f has roots of integral valuation, which is a necessary condition for their to be roots in \mathbb{Q}_p . Steps 2 and 8 involve substitutions that only negligibly affect the heights of the coefficients, similar to the binomial case (where the underlying rescalings are stated in finer detail).

Step 3 correctly detects degenerate roots in \mathbb{C}_p^* thanks to Lemma 5.2. As observed above, Steps 4–7 correctly count the number of non-degenerate roots of f in \mathbb{Z}_p of valuation 0. In particular, Step 4 is accomplished via Lemmata 5.2 and 5.4, and the characterization of degenerate roots from the latter lemma implies that we can use the Extended Euclidean Algorithm to find a binomial efficiently encoding the degenerate roots of f in \mathbb{Q}_p (as already detailed in the third paragraph of the proof of Lemma 6.9). ■

Complexity Analysis: Steps 1, 2, and 8 involve basic field arithmetic that will be dominated by Steps 3–7. So we will focus on Steps 3–7 only.

Step 3 can be accomplished in time $O(\log^2(dH))$ via [3, Thm. 39]. Note in particular that detecting vanishing for $\Delta_{\text{tri}}(f)$ is much easier than computing its valuation.

Step 4 takes time $O((p + \log(dH)) \log(dpH) \log \log(dpH))$ thanks to Theorem 2.21.

Letting ν and \mathcal{D} respectively denote the number of degenerate roots of \tilde{f} in \mathbb{F}_p^* and the depth of $\mathcal{T}_{p,k}(f)$, Step 5 takes time $O(\nu p^2 \log^4(dH) \log_p^3(d) \log(p \log(dH)))$ or

$$O((p + \log d) \log(dp) \log \log(dp) + p \log^2(p) \log \log(p) + \nu \log^2(dH) \log(d) \log_p \log(dH) + \log(H) \log(dpH) \log \log(dpH)),$$

according as $\delta = 0$ or $\delta = 1$. This follows immediately from an elementary calculation, upon substituting the corresponding value of k from Corollary 6.6 into Lemma 6.9, using the fact that the depth \mathcal{D} is bounded from above by one of our two bounds from Theorem 1.6.

The brute-force portion of Step 6 clearly takes time $O(p \log^2(p) \log \log p)$ via Theorem 2.7. Lemma 6.9 tells us that $\mathcal{T}_{p,k}(f)$ has $O(\nu \mathcal{D})$ nodes, and Lemma 6.1 tells us that each non-root nodal polynomial has mod p reduction with degree ≤ 4 . So the remaining multi-node gcd computation takes time $O(\nu \mathcal{D} \cdot \log(p) \log \log p)$ via Theorem 2.7. So the overall time for Step 6 is $O(p [\nu \log^2(dH) \log_p(d) + \log^2 p] \log \log p)$ or $O([p \log^2(p) + \nu \log(dH)] \log \log p)$, according as δ is 0 or 1, thanks to Theorem 1.6.

As for Step 7, we already know the non-degenerate roots in \mathbb{F}_p of \tilde{f} from Step 6. For the remaining nodes, observe that Lemma 6.1 tells us that the mod p reductions of the non-root nodal polynomials have degree at most 4. Also, the root has ν children, each yielding a tree that is a chain with (at worst) one bifurcation. Furthermore, note that the presence of a non-degenerate root in \mathbb{F}_p for $\tilde{f}_{i,\zeta}$ implies that $\tilde{f}_{i,\zeta}$ can have at most 1 degenerate root in \mathbb{F}_p , meaning that its child will have degree at most 2 by Lemma 2.17. Finally, note that once a quadratic $\tilde{f}_{i,\zeta}$ has a non-degenerate root in \mathbb{F}_p , it can no longer have any children. In other words, we have shown that there can be at most $O(\nu)$ nodes having $\tilde{f}_{i,\zeta}$ possessing a non-degenerate root. Applying Shoup's deterministic factoring algorithm [55] to the non-root nodal polynomials, we then see that finding the non-degenerate roots for our entire tree takes time $O(\nu \cdot p^{1/2} \log^2 p)$.

In summary, we see that Step 5 dominates our overall complexity when $\delta = 0$, yielding a bound of $O(\nu p^2 \log^4(dH) \log_p^3(d) \log(p \log(dH)))$. When $\delta = 1$, Steps 4, 5, and 7 dominate together, yielding an overall complexity bound of Noting that $\nu \leq p - 1$, we

$$O((p + \log(dH)) \log(dpH) \log \log(dpH) + p \log^2(p) \log \log(p) + \nu [p^{1/2} \log^2(p) + \log^2(dH) \log(d) \log_p \log(dH)]).$$

are done after an elementary calculation. ■

Remark 6.14. A consequence of our proof is that it also contains a proof of the deterministic complexity bound of Corollary 1.7, since we included above the case where f has a degenerate root. To get the Las Vegas randomized bound, we simply replace the fast deterministic factoring algorithm from [55] in Step 7 with the fast randomized factoring algorithm from [34]. ◇

6.4. “Typical” exponents, Las Vegas, and a combined speed-up

For our final speed-ups we will make use of the fact that trinomials can only vanish on a small number of cosets in \mathbb{F}_q^* : Building on earlier results from [15,12,35], Kelley and Owen proved [36, Thm. 1.2] that $c_1 + c_2x^{a_2} + c_3x^{a_3} \in \mathbb{F}_q[x]$, with q a prime power, vanishes at no more than $\left\lfloor \frac{1}{2} + \sqrt{\frac{q-1}{r'}} \right\rfloor$ cosets of the size r' subgroup of \mathbb{F}_q^* (and nowhere else), where $r' = \gcd(a_2, a_3, q-1)$. In particular, this bound is optimal for \mathbb{F}_q an even degree extension of a prime field. For q prime, there is even computational evidence (for all $q \leq 292837$) that the number of such cosets might in fact no greater than $2 \log q$ [20].

It is easy to see that, for any fixed prime p , $\gcd(a_2a_3(a_3 - a_2), (p-1)p) \leq 2$ for a positive density subset of $(a_2, a_3) \in \mathbb{N}^2$. (Simply pick a_2 and a_3 to avoid certain arithmetic progressions depending on p and the divisors of $p-1$.) So one can argue that a large fraction of trinomials over \mathbb{Z} have $O(\sqrt{p})$ roots in \mathbb{F}_p and, via Lemma 2.18, $O(\sqrt{p})$ roots in \mathbb{Q}_p . Apropos of this paucity of roots for “most” exponents, let us recall a useful trick that will allow us to significantly reduce the degree of a large fraction of trinomials over \mathbb{F}_p : Via a fast algorithm for the *Shortest Lattice Vector Problem* in \mathbb{Z}^2 (see, e.g., [24]), one can prove the following result:

Lemma 6.15. [12, Special Case of Lemma 1.11] *Given any prime p , and $a_2, a_3 \in \mathbb{N}$ with $0 < a_2 < a_3 < p-1$ and $r' := \gcd(a_2, a_3, (p-1)p)$, one can find within $\log^{O(1)} p$ bit operations an integer e such that for all $i \in \{2, 3\}$, $ea_i = m_i \bmod p-1$ and $|m_i| \leq r' \sqrt{2(p-1)}$. ■*

Proof of Corollary 1.4. We follow the template of the proof of Theorem 1.1, save for some key differences. The first difference is that, under our assumptions, we can compute the tree $\mathcal{T}_{p,k}(f)$ faster via Corollary 6.10 instead of Lemma 6.9. We then need to compute the non-degenerate roots of all the nodal polynomials, so the next key difference is that we can use the degree reduction of Lemma 6.15 to speed up this up at the root node. (The remaining nodes receive no further speed-up unless randomization is used.)

So we merely need to recompute our complexity bounds. Recall that \mathcal{D} denotes the depth of the tree $\mathcal{T}_{p,k}(f)$, and ν is the number of children of the root node (which for k sufficiently large, is the number of degenerate roots of \tilde{f}). We note the changes to the complexity of Algorithm 6.12 below, in both the *restricted root* case (where we only seek root of the form $p^j + O(p^{j+1})$) and the *small gcd* case (where we assume $\gcd(a_2a_3(a_3 - a_2), (p-1)p) \leq 2$):

A. Step 4 can be sped up to deterministic time

$$O(\log^2(p) \log \log(p) + \log(\max\{d, p\}) \log(\log \max\{d, p\}) + \log(\max\{p, H\}) \log(\log \max\{p, H\}))$$

in the restricted root case; or deterministic time

$$O(p^{1/2} \log^2(p) + \log(\max\{d, p\}) \log(\log \max\{d, p\}) + \log(\max\{p, H\}) \log(\log \max\{p, H\})),$$

or Las Vegas randomized time

$$O\left(\log^{2+o(1)}(p) + \log(\max\{d, p\}) \log \log \max\{d, p\} + \log(\max\{H, p\}) \log \log \max\{H, p\}\right)$$

in the small gcd case.

B. Step 5 can be sped up to deterministic time

$$O((p + \log d) \log(dp) \log(\log(dp)) + p \log^2(p) \log \log(p) + \mathcal{D}[k \log(p) \log(k \log p) \log(d) + \log H] + \log(H) \log(dpH) \log \log(dpH)),$$

in both cases. If f has a degenerate root in \mathbb{C}_p^* then we can further speed up both cases to Las Vegas randomized time

$$O(\log^2(dH) \log(\log(dH)) + \log^2(p) \log(\log p) + \log(dpH) \log \log(dpH)).$$

C. We replace Step 6 of Algorithm 6.12 with the following:

6': By computing $\deg \gcd(\tilde{f}_{i,\zeta}, x^p - x)$ for the non-root nodal polynomials of $\mathcal{T}_{p,k}(f)$, and factoring a degree-reduced version of \tilde{f} (if needed), determine which nodal polynomials have non-degenerate roots in \mathbb{F}_p .

This modified step takes deterministic time $O(\mathcal{D} \log(p) \log \log p)$ in the restricted root case; or deterministic time $O(p \log^2(p) + \mathcal{D} \log(p) \log \log p)$ or Las Vegas randomized time $O(p^{3/4} \log^{1+o(1)}(p) + \mathcal{D} \log(p) \log \log p)$ in the small gcd case.

D. Step 7 can be sped up to deterministic time $p^{1/2} \log^{2+o(1)} p$ or Las Vegas randomized time $\log^{2+o(1)} p$, in both cases.

We now explain Changes A–D.

A. In the restricted root case, we merely need to evaluate f and f' at 1, so our first bound is clear.

In the small gcd case, the number of degenerate roots is at most 2 thanks to our gcd assumption and Lemma 5.4. So instead of employing Algorithms 2.22 or 2.24, we simply find the degenerate roots by factoring, using either the fast deterministic algorithm from [55] or the fast Las Vegas randomized factorization algorithm from [34].

B. The complexity bounds follows by applying Corollary 6.10 instead of Lemma 6.9, ultimately yielding $O(p^2 \log^4(dH) \log_p^3(d) \log(p \log(dH)))$ via Corollary 6.6 and Theorem 1.6. As noted in Remark 6.11, our current bounds for k and \mathcal{D} obstruct any Las Vegas speed-up for Step 5 (in the non-degenerate case).

C. The deterministic speed-ups follow from the complexity analysis of Algorithm 6.12, in the proof of Theorem 1.1, simply by setting $\nu = 2$ in the bound there. Note also that in the restricted root case, there is no need to search for any roots of \tilde{f} since we only care about most significant digit 1: We merely need to evaluate \tilde{f} and \tilde{f}' at 1.

To get our Las Vegas speed-up, we replace the brute-force search for degenerate roots of \tilde{f} with a targeted factorization: First build a degree-reduced version of \tilde{f} via Lemma 6.15 to apply the automorphism of \mathbb{F}_p^* defined by $x \mapsto x^e$ to replace \tilde{f} by $\tilde{g}(x) := \tilde{f}(x^e)$, and compute $e' := 1/e \bmod p-1$, in deterministic time $\log^{O(1)} p$. This reduces $\deg \tilde{f}$ to $\deg \tilde{g} \leq 2\sqrt{2(p-1)}$. To find the roots of \tilde{f} in \mathbb{F}_p^* we can then find the roots of \tilde{g} in \mathbb{F}_p^* by using the Kedlaya-Umans factorization algorithm [34], take the e' th powers mod p of these roots, and then identify which of these roots of \tilde{f} is a degenerate root found earlier. This takes time $(2\sqrt{2(p-1)})^{1.5} \log^{1+o(1)}(p) + \log^{O(1)} p = p^{3/4} \log^{1+o(1)} p$.

Since $\nu \leq 2$ in both cases, the remaining multinodal gcd computation takes additional deterministic time $O(\mathcal{D} \log(p) \log \log p)$.

D. Since we already found the non-degenerate roots of \tilde{f} in \mathbb{F}_p in Step 6', we merely need to speed up finding the non-degenerate roots in \mathbb{F}_p of the remaining nodal polynomials: We already observed in the proof of Theorem 1.1 that there are $O(\nu)$ nodes having a $\tilde{f}_{i,\zeta}$ possessing a non-degenerate root. But $\nu \leq 2$ in both cases, so we only need to worry about $O(1)$ nodes. So our proof of Theorem 1.1 already implies a deterministic speed-up to $O(p^{1/2} \log^2 p)$ (for $O(1)$ applications of Shoup's deterministic factoring algorithm [55]), in both cases.

However, if we replace Shoup's algorithm with the fast randomized factorization algorithm from [34], then we can speed Step 7 up to Las Vegas randomized time $\log^{2+o(1)} p$ in both cases.

To conclude, we see that Step 5 dominates the deterministic complexity in both cases (restricted root and small gcd), and wipes out any Las Vegas speed-up unless better bounds for k and \mathcal{D} are available. Summing our complexity estimates, we obtain our desired bounds. ■

An immediate consequence of our last proof — if we can apply the sharper bounds for \mathcal{D} and k from the degenerate cases of Theorem 1.6 and Corollary 6.6 — is the following combined speed-up:

Corollary 6.16. *If the trinomial $f \in \mathbb{Z}[x]$ has a nonzero degenerate root in \mathbb{C}_p then we can speed up the Las Vegas complexity bound of Corollary 1.7 to*

$$\log^{2+o(1)}(p) + \log^{2+o(1)}(dH) \log_p d \text{ (in the restricted root case)}$$

or

$$p^{3/4} \log^{1+o(1)}(p) + \log^{2+o(1)}(dH) \log_p d \text{ (in the small gcd case).} \blacksquare$$

Acknowledgments

We thank Erich Bach and Bjorn Poonen for informative discussions on Hensel's Lemma. Special thanks to Elliott Fairchild, Josh Goldstein, and David Zureick-Brown for inspirational conversations during the pandemic. We also thank the anonymous referees for helpful suggestions that improved our paper, particularly the one referee who pointed out that barrier B1 might be amenable to a good data structure.

References

- [1] Leonard Adleman, Kenneth Manders, Gary Miller, On taking roots in finite fields, in: 18th Annual Symposium on Foundations of Computer Science, Providence, R.I., 1977, pp. 175–178, IEEE, 1977.
- [2] Sanjeev Arora, Boaz Barak, Computational Complexity. A Modern Approach, Cambridge University Press, Cambridge, 2009.
- [3] Martín Avendaño, Ashraf Ibrahim, J. Maurice Rojas, Korben Rusek, Faster p -adic feasibility for certain multivariate sparse polynomials, *J. Symb. Comput.* 47 (4) (2012) 454–479.
- [4] Martín Avendaño, Roman Kogan, Mounir Nisse, J. Maurice Rojas, Metric estimates and membership complexity for Archimedean amoebae and tropical hypersurfaces, *J. Complex.* 46 (2018) 45–65.
- [5] Martín Avendaño, Jorge Martín-Morales, Bivariate trinomials over finite fields, *Houst. J. Math.* (2022), in press.
- [6] Martín Avendaño, Teresa Krick, Sharp bounds for the number of roots of univariate fewnomials, *J. Number Theory* 131 (7) (2011) 1209–1228.
- [7] Eric Bach, Jeffrey Shallit, Algorithmic Number Theory, Volume 1: Efficient Algorithms, MIT Press, Cambridge, Massachusetts, 1996.
- [8] A. Baker, Logarithmic forms and the abc -conjecture, in: Number Theory, Eger, 1996, De Gruyter, Berlin, 1998, pp. 37–44.
- [9] Jennifer Balakrishnan, Netan Dogra, J. Steffen Müller, Jan Tuitman, Jan Vonk, Explicit Chabauty–Kim for the split Cartan modular curve of level 13, *Ann. Math.* (2) 189 (3) (2019) 885–944.
- [10] Jens-Dietrich Bauch, Enric Nart, Hayden D. Stainsby, Complexity of OM factorizations of polynomials over local fields, *LMS J. Comput. Math.* 16 (2013) 139–171.
- [11] Jérémie Berthomieu, Grégoire Lecerf, Guillaume Quintin, Polynomial root finding over local rings and application to error correcting codes, *Appl. Algebra Eng. Commun. Comput.* 24 (2013) 413–443.
- [12] Jingguo Bi, Qi Cheng, J. Maurice Rojas, Sub-linear root detection, and new hardness results, for sparse polynomials over finite fields, in: Proceedings of the 38th International Symposium on Symbolic and Algebraic Computation, ISSAC '13, Association for Computing Machinery, New York, NY, USA, 2013, pp. 61–68.
- [13] Erick Boniface, Weixun Deng, J. Maurice Rojas, Near-optimal root spacing and faster solving for very sparse polynomials over \mathbb{R} , arXiv, <https://arxiv.org/abs/2202.06115>, 2022.
- [14] J.M. Borwein, P.B. Borwein, On the complexity of familiar functions and numbers, *SIAM Rev.* 30 (4) (1988) 589–601.
- [15] Ran Canetti, John Friedlander, Sergei Konyagin, Michael Larsen, Daniel Lieman, Igor Shparlinski, On the statistical properties of Diffie-Hellman distributions, *Isr. J. Math.* 120 (1) (Dec 2000) 23–46.
- [16] David G. Cantor, Daniel M. Gordon, Factoring polynomials over p -adic fields, in: Wieb Bosma (Ed.), Algorithmic Number Theory, Springer Berlin Heidelberg, Berlin, Heidelberg, 2000, pp. 185–208.
- [17] David G. Cantor, Erich Kaltofen, On fast multiplication of polynomials over arbitrary algebras, *Acta Inform.* 28 (7) (1991) 693–701.
- [18] Zhengjun Cao, Qian Sha, Xiao Fan, Adleman-Manders-Miller root extraction method revisited, in: Information Security and Cryptology, in: Lecture Notes in Comput. Sci., vol. 7537, Springer, Heidelberg, 2012, pp. 77–85.
- [19] Qi Cheng, Primality proving via one round in ECPP and one iteration in AKS, *J. Cryptol.* 20 (3) (July 2007) 375–387.
- [20] Qi Cheng, Shuhong Gao, J. Maurice Rojas, Daqing Wan, Sparse univariate polynomials with many roots over finite fields, *Finite Fields Appl.* 46 (2017) 235–246.
- [21] Gook Hwa Cho, Soonhak Kwon, Hyang-Sook Lee, A refinement of Müller's cube root algorithm, *Finite Fields Appl.* 67 (2020) 101708.

- [22] Keith Conrad, Notes on Hensel's Lemma, downloadable from kconrad.math.uconn.edu/blurbs/gradnumthy/hensel.pdf, 2021.
- [23] Edgar Costa, David Harvey, Kiran S. Kedlaya, Zeta functions of nondegenerate hypersurfaces in toric varieties via controlled reduction in p -adic cohomology, in: Proceedings of the Thirteenth Algorithmic Number Theory Symposium, in: Open Book Ser., vol. 2, Math. Sci. Publ., Berkeley, CA, 2019, pp. 221–238.
- [24] Daniel Dadush, Chris Peikert, Santosh Vempala, Enumerative lattice algorithms in any norm via M-ellipsoid coverings, in: 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, IEEE Computer Soc., Los Alamitos, CA, 2011, pp. 580–589.
- [25] Anindya De, Piyush P. Kurur, Chandan Saha, Ramprasad Saptharishi, Fast integer multiplication using modular arithmetic, SIAM J. Comput. 42 (2) (2013) 685–699.
- [26] Gustave Dumas, Sur quelques cas d'irréductibilité des polynômes à coefficients rationnels, J. Math. Pures Appl. 2 (6) (1906) 191–258.
- [27] Ashish Dwivedi, Rajat Mittal, Nitin Saxena, Counting basic-irreducible factors mod p^k in deterministic poly-time and p -adic applications, in: 34th Computational Complexity Conference, in: LIPIcs. Leibniz Int. Proc. Inform., vol. 137, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019, p. 15, 29.
- [28] Elliot Fairchild, Joshua Goldstein, J. Maurice Rojas, Trinomials with Tightly Packed p -Adic Roots, Texas A&M University, 2022, in preparation.
- [29] Israel M. Gel'fand, Misha M. Kapranov, Andrei V. Zelevinsky, Discriminants, Resultants, and Multidimensional Determinants. Mathematics: Theory & Applications, Birkhäuser Boston, Inc., Boston, MA, 1994.
- [30] Jordi Guàrdia, Enric Nart, Sebastian Pauli, Single-factor lifting and factorization of polynomials over local fields, J. Symb. Comput. 47 (11) (2012) 1318–1346.
- [31] David Harvey, Joris van der Hoeven, Polynomial multiplication over finite fields in time $O(n \log n)$, HAL preprint, <https://hal.archives-ouvertes.fr/hal-02070816>, 2019.
- [32] David Harvey, Joris van der Hoeven, Integer multiplication in time $O(n \log n)$, Ann. Math. (2) 193 (2) (2021) 563–617.
- [33] Loo-Keng Hua, H.S. Vandiver, On the number of solutions of some trinomial equations in a finite field, Proc. Natl. Acad. Sci. USA 35 (1949) 477–481.
- [34] Kiran Kedlaya, Christopher Umans, Fast polynomial factorization and modular composition, in: Peter Bro Miltersen, Rüdiger Reischuk, Georg Schnitger, Dieter van Melkebeek (Eds.), Computational Complexity of Discrete Problems, Dagstuhl Seminar Proceedings, Number 08381, Dagstuhl, Germany, 2008, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2008.
- [35] Zander Kelley, Roots of sparse polynomials over a finite field, LMS J. Comput. Math. 19 (A) (2016) 196–204.
- [36] Zander Kelley, Sean W. Owen, Estimating the number of roots of trinomials over finite fields, J. Symb. Comput. 79 (2017) 108–118, SI: MEGA 2015.
- [37] Pascal Koiran, Root separation for trinomials, J. Symb. Comput. 95 (2019) 151–161.
- [38] Pascal Koiran, Natacha Portier, Sébastien Tavenas, A Wronskian approach to the real τ -conjecture, J. Symb. Comput. 68 (part 2) (2015) 195–214.
- [39] Leann Kopp, Natalie Randall, J. Maurice Rojas, Yuyu Zhu, Randomized polynomial-time root counting in prime power rings, Math. Comput. 89 (321) (January 2020) 373–385.
- [40] Hendrik W. Lenstra, On the factorization of lacunary polynomials, Number Theory Prog. 1 (1999) 277–291.
- [41] Kurt Mahler, An inequality for the discriminant of a polynomial, Mich. Math. J. 11 (3) (1964) 257–262.
- [42] Maurice Mignotte, On the distance between the roots of a polynomial, Appl. Algebra Eng. Commun. Comput. 6 (327–332) (1995) 11.
- [43] Alexandre Ostrowski, Recherches sur la méthode de Graeffe et les zéros des polynomes et des séries de Laurent, Acta Math. 72 (1940) 99–155.
- [44] Christos H. Papadimitriou, Computational Complexity, Addison-Wesley, 1995.
- [45] Kaitlyn Phillipson, J. Maurice Rojas, Fewnomial systems with many roots, and an adelic tau conjecture, in: Tropical and Non-Archimedean Geometry, in: Contemp. Math., vol. 605, Amer. Math. Soc., Providence, RI, 2013, pp. 45–71.
- [46] Bjorn Poonen, Zeros of sparse polynomials over local fields of characteristic p , Math. Res. Lett. 5 (3) (1998) 273–279.

- [47] Bjorn Poonen, Using zeta functions to factor polynomials over finite fields, in: *Arithmetic Geometry: Computation and Applications*, in: *Contemp. Math.*, vol. 722, Amer. Math. Soc., Providence, RI, 2019, pp. 141–147.
- [48] Q.I. Rahman, G. Schmeisser, *Analytic Theory of Polynomials*, London Mathematical Society Monographs. New Series, vol. 26, The Clarendon Press, Oxford University Press, Oxford, 2002.
- [49] Alain M. Robert, *A Course in p -adic Analysis*, Springer-Verlag, New York, 2000.
- [50] J. Maurice Rojas, Yuyu Zhu, 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, Association for Computing Machinery, New York, NY, USA, 2021, pp. 337–344.
- [51] Jeremy Rouse, Andrew V. Sutherland, David Zureick-Brown, ℓ -adic images of Galois for elliptic curves over \mathbb{Q} , arXiv:2106.11141, 2021.
- [52] Michael Sagraloff, A near-optimal algorithm for computing real roots of sparse polynomials, in: *39th International Symposium on Symbolic and Algebraic Computation*, ISSAC 2014, 2014, pp. 359–366.
- [53] W.H. Schikhof, *Ultrametric Calculus*, Cambridge Studies in Advanced Mathematics, vol. 4, Cambridge University Press, Cambridge, 2006, An introduction to p -adic analysis, reprint of the 1984 original, MR0791759.
- [54] J.-P. Serre, *A Course in Arithmetic*, Graduate Texts in Mathematics, vol. 7, Springer-Verlag, New York-Heidelberg, 1973, translated from the French.
- [55] Victor Shoup, On the deterministic complexity of factoring polynomials over finite fields, *Inf. Process. Lett.* 33 (5) (1990) 261–267.
- [56] Igor Shparlinski, On finding primitive roots in finite fields, *Theor. Comput. Sci.* 157 (2) (1996) 273–275.
- [57] 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, Springer, New York, 1986, pp. 185–196.
- [58] Andrew V. Sutherland, Lecture notes for Math 18.783 (elliptic curves), Lecture #3, February 24, 2021.
- [59] H.W. Turnbull (Ed.), *The Correspondence of Isaac Newton*, Vol. II: 1676–1687, Cambridge University Press, New York, 1960, published for the Royal Society.
- [60] Joachim von zur Gathen, Jürgen Gerhard, *Modern Computer Algebra*, third edition, Cambridge University Press, Cambridge, 2013.
- [61] Joachim von zur Gathen, Silke Hartlieb, Factoring modular polynomials, *J. Symb. Comput.* 26 (1998) 583–606.
- [62] André Weil, Numbers of solutions of equations in finite fields, *Bull. Am. Math. Soc.* 55 (5) (May 1949) 497–508.
- [63] Edwin Weiss, *Algebraic Number Theory*, International Series in Pure and Applied Mathematics, McGraw-Hill, 1963.
- [64] Kunrui Yu, Linear forms in p -adic logarithms III, *Compos. Math.* 3 (241–276) (1994).
- [65] Kunrui Yu, p -adic logarithmic forms and group varieties. III, *Forum Math.* 19 (2) (2007) 187–280.
- [66] Yuyu Zhu, *Trees, Point Counting Beyond Fields, and Root Separation*, PhD thesis, doctoral dissertation, TAMU 3368, College Station, TX 77843-3368, Texas A&M University, May 2020.