SUB-LINEAR POINT COUNTING FOR VARIABLE SEPARATED CURVES OVER PRIME POWER RINGS

CALEB ROBELLE, J. MAURICE ROJAS, AND YUYU ZHU

Abstract. Let $k, p \in \mathbb{N}$ with p prime and let $f \in \mathbb{Z}[x_1, x_2]$ be a bivariate polynomial with degree d and all coefficients of absolute value at most p^k. Suppose also that f is variable separated, i.e., $f = g_1 + g_2$ for $g_i \in \mathbb{Z}[x_i]$. We give the first algorithm, with complexity sub-linear in p, to count the number of roots of f over $\mathbb{Z}/\langle p^k \rangle$ for arbitrary k: Our Las Vegas randomized algorithm works in time $(dk \log p)^{O(1)}\sqrt{p}$, and admits a quantum version for smooth curves working in time $(d \log p)^{O(1)}k$. Save for some subtleties concerning non-isolated singularities (which we fully explain below), our techniques generalize to counting roots of polynomials in $\mathbb{Z}[x_1, \ldots, x_n]$ over $\mathbb{Z}/\langle p^k \rangle$.

Our techniques are a first step toward efficient point counting for varieties over Galois rings (which is relevant to error correcting codes over higher-dimensional varieties), and also imply new speed-ups for computing Igusa zeta functions of curves. The latter zeta functions are fundamental in arithmetic geometry.

Current affiliation and address of authors:

(Robelle):
University of Maryland, Baltimore County
1000 Hilltop Circle
Baltimore, MD 21250

(Rojas & Zhu):
Texas A&M University, Department of Mathematics
TAMU 3368
College Station, TX 77845

emails: `carobel1@umbc.edu` , `rojas@math.tamu.edu` , `zhuyuyu@math.tamu.edu`

C.B. was partially supported by NSF grant DMS-1757872.
J.M.R. and Y.Z. were partially supported by NSF grants CCF-1900881 and DMS-1757872.

1. Introduction

Counting points on algebraic curves over finite fields is a seemingly simple problem that nevertheless helped form the core of arithmetic geometry in the 20th century and now forms an important part of cryptography [Mil86, Kob87, GG16] and coding theory [vdG01]. Efficient algorithms for this problem continue to be a lively part of computational number theory: The barest list of references would have to include [Sch85, Pil90, AH01, Ked01, CDV06, LW08, Wan08, CL08, Har15].[1] Here, we consider algorithms for the natural extension of this problem to prime power rings, and find the first efficient algorithms for a broad class of (not necessarily smooth) curves: See Theorem 1.1 below. It will be useful to first discuss some motivation before covering further background.

1.1. A Connection to Error Correcting Codes.

Suppose $k, p \in \mathbb{N}$ with p prime, $\mathbb{F}_p$ is the field with p elements, and $r \in \mathbb{Z}[x_1]$ is a univariate polynomial of degree m that is irreducible mod p. We call a quotient ring R of the form $\mathbb{Z}[x_1]/\langle p^k, r(x_1)\rangle$ a *Galois ring*. Note that such an R is finite, and can be the prime power ring $\mathbb{Z}/\langle p^k\rangle$ (for $m=1$) or the field $\mathbb{F}_q$ (for $k=1$ and $q=p^m$), to name a few examples.

Since numerous error correcting codes and cryptosystems are based on arithmetic over $\mathbb{F}_q$ or $\mathbb{F}_q[x_1]$, it has been observed (see, e.g., [GCM91, GSS00, BLQ13, CH15]) that one can generalize and improve these constructions by using arithmetic over R or $R[x_1]$ instead. For instance, Guruswami and Sudan's famous list-decoding method for error correcting codes [GS99] involves finding the roots in $\mathbb{F}_q[x_1]$ of a polynomial in $\mathbb{F}_q[x_1, x_2]$ as a key step, and has a natural generalization to Galois rings (see, e.g., [HKC$^+$94, Sud97, BW10] and [BLQ13, Sec. 4]). Furthermore, counting solutions to equations like $f(x_1, \ldots, x_n)=0$ over Galois rings determines the weights of codewords in Reed-Muller codes over Galois rings, and the weight distribution governs the quality of the underlying code (see, e.g., [KLP12]).

1.2. Connections to Zeta Functions and Rational Points.

Efficiently counting roots in $\left(\mathbb{Z}/\langle p^k\rangle\right)^2$ of polynomials in $\mathbb{Z}[x_1, x_2]$ is a natural first step toward efficiently enumerating the roots in R^2 for polynomials in $R[x_1, x_2]$ for R a Galois ring. However, observe that the ring of p-adic integers $\mathbb{Z}_p$ is the inverse limit of $\mathbb{Z}/\langle p^k\rangle$ as $k \longrightarrow \infty$. It then turns out that the zero sets of polynomials over $\mathbb{Z}/\langle p^k\rangle$ inform the zero sets of polynomials over $\mathbb{Z}_p$ and beyond.

In particular, for any $f \in \mathbb{Z}[x_1, \ldots, x_n]$, one can form a fundamentally important generating function, and a related zeta function, as follows: Let $N_{p,k}(f)$ denote the number of roots in $\left(\mathbb{Z}/\langle p^k\rangle\right)^n$ of the mod p^k reduction of f and define the *Poincare series of f* to be $P_f(t) = \sum_{k=0}^{\infty} \frac{N_{p,k}(f)}{p^{kn}} t^k$. Also, letting $t := p^{-s}$, we define the *Igusa local zeta function of f* to be $Z_f(t) := \int_{\mathbb{Z}_p} |f(x_1, \ldots, x_n)|_p^s dx$, where $|\cdot|_p$ and dx respectively denote the standard p-adic absolute value on $\mathbb{Z}_p$ and Haar measure on $\mathbb{Z}_p$. (This function turns to be defined on the right open half-plane of $\mathbb{C}$, possibly with the exception of finitely many poles.) The precise definitions of $|\cdot|_p$ and dx won't matter for our algorithmic results, but what does matter is that Igusa discovered in the 1970s that $P(t) = \frac{1-tZ(t)}{1-t}$ and proved that Z (and thus P) is a rational function of t [Igu07].

[1] Also, major conferences such as ANTS consistently continue to feature papers on speeding up point-counting for various special families of curves and surfaces.

Igusa defined his zeta function Z with the goal of generalizing earlier work of Siegel (on counting representations of integers via quadratic forms) to high degree forms, e.g., how many ways can one write 239 as a sum of cubes? However, the algorithmic computation of these zeta functions has received little attention, aside from some very specific cases. Our results imply that one can compute Z for certain bivariate f in time polynomial in $dk\log p$. This extends earlier work on the univariate case [DMS20, Zhu20] to higher-dimensions and will be pursued in a sequel to this paper.

It should also be pointed out that recent algorithmic methods for finding rational points (over $\mathbb{Q}$) for curves of genus ≥ 2 proceed (among many other difficult steps) by finding the p-adic rational points on a related family of varieties (see, e.g., [BM20, Sec. 5.3]). So a long term goal of this work is to improve the complexity of finding the p-adic rational points on curves and surfaces, generalizing recent p-adic speed-ups in the univariate case [RZ20].

1.3. **From Finite Fields to Prime Power Rings.** Returning to point counting over prime power rings, the computation of $N_{p,k}(f)$ is subtle already for $n=1$: This special case has recently been addressed from different perspectives in [BLQ13, CGRW18, KRRZ19, DMS19], and was just recently proved to admit a *deterministic* algorithm of complexity $(dk\log p)^{O(1)}$, thanks to the last paper.

The special case $(n,k)=(2,1)$ of computing $N_{p,1}(f)$, just for f a cubic polynomial, is already of considerable interest in the design of cryptosystems based on the elliptic curve discrete logarithmic problem. In fact, even this very special case wasn't known to admit an algorithm polynomial in $\log p$ until Schoof's work in the 1980s [Sch85]. More recently, algorithms for computing $N_{p,1}(f)$ for arbitrary $f\in\mathbb{Z}[x_1,x_2]$ of degree d, with complexity $d^8(\log p)^{2+o(1)}\sqrt{p}$, have been derived by Harvey [Har15] (see also [Zhu20, Ch. 5]), and similar complexity bounds hold for arbitrary finite fields.

Our main result shows that counting points over $\mathbb{Z}/\langle p^k\rangle$ for arbitrary k is slower than the $k=1$ case only by a factor polynomial in k (neglecting the other parameters).

Theorem 1.1. *Suppose $f=g_1+g_2$ for some $g_i\in\mathbb{Z}[x_i]$, $\deg f=d\geq 1$, and all the coefficients of f are of absolute value at most p^k. Then there is a Las Vegas randomized algorithm that computes $N_{p,k}(f)$ in time $(dk\log p)^{O(1)}\sqrt{p}$. In particular, the number of random bits needed is $O(d^2k\log(dk)\log p)$, and the space needed is $O(d^4k\sqrt{p}\log p)$. Furthermore, if the zero set of f over the algebraic closure $\bar{\mathbb{F}}_p$ is smooth and irreducible, then $N_{p,k}(f)$ can be computed in quantum randomized time $(d(\log p))^{O(1)}k$.*

We prove Theorem 1.1 in Section 4.1. The central idea is to reduce to a moderate number of moderately sized instances of point counting over $\mathbb{F}_p$. Recall that *Las Vegas randomized time* simply means that our algorithm needs random bits and gives an answer that is correct with probability at least 1/2 and, in case of error, states that an error has occured. *Quantum* randomized time here will mean that we avail to a quantum computer, and instead obtain an algorithm that gives an answer that is correct with probability at least 2/3, but with no correctness guarantee.

In what follows, we call a polynomial of the form $f_\zeta(x_1,x_2):=\frac{1}{p^s}f(\zeta_1+px_1,\zeta_2+px_2)$, with $(\zeta_1,\zeta_2)\in\mathbb{F}_p^2$ a singular point of the zero set of f in $\mathbb{F}_p^2$ and s as large as possible with f_ζ still in $\mathbb{Z}[x_1,x_2]$, a *perturbation* of f. Our reduction to point counting over $\mathbb{F}_p$ will involve finding all *isolated* singular points of the zero set of f (as well as its perturbations) in $\mathbb{F}_p^2$, in order to categorize the base-p digits of the coordinates of the roots of f in $\left(\mathbb{Z}/\langle p^k\rangle\right)^2$. This yields a geometrically defined recurrence for $N_{p,k}(f)$ that is conveniently encoded by a tree. We detail this construction in Sections 2.2 and 4.1 below.

Remark 1.2. *A classical algebraic geometer may propose simply applying resolution of singularities, applying finite field point counting (with proper corrections at blown-up singular points), and then an application of Hensel's Lemma. We use a more direct approach that allows us to lift singular points individually and much more simply. In particular, it appears (from* [PR11]*) that resolution of singularities for a plane curve of degree d over $\mathbb{F}_p$ has complexity $O(d^5)$ (neglecting multiples depending on p), while our algorithm (if looked at more closely) has better dependence on d. More to the point, replacing an input bivariate polynomial by a higher degree complete intersection (the latter being the output after doing resolution of singularities) results in a more complicated input when one needs to avail to prime field point counting, thus compounding the complexity even further. Furthermore, in higher dimensions, resolution of singularities becomes completely impractical* [BGMWo11]. $\diamond$

Remark 1.3. *We can extend Theorem 1.1 to more general curves. The key obstruction is whether f, or one of its perturbations, fails to be square-free (see the final section of the Appendix). We hope to extend our methods to arbitrary curves in the near future. For now, we simply point out that many commonly used curves in practice are variable separated, e.g., many hyperelliptic curves used in current cryptography are zero sets of polynomials of the form $x_2^2 - g(x_1)$.* $\diamond$

2. Background

2.1. **Some Basics on Point Counting Over Finite Fields.**

One of the most fundamental results on point counting for curves over finite fields dates back to work of Hasse and Weil in the 1940s. In what follows, we use $|S|$ to denote the cardinality of a set S.

Theorem 2.1. [Wei49] *Let $\mathbb{F}_q$ be a finite field of order $q = p^m$, and let $\mathcal{C}$ be an absolutely irreducible smooth projective curve defined over $\mathbb{F}_q$. Let g denote the genus of $\mathcal{C}$ and $\mathcal{C}(\mathbb{F}_q)$ to be the set of $\mathbb{F}_q$-points of $\mathcal{C}$. Then $||\mathcal{C}(\mathbb{F}_q)| - q| \leq 2g\sqrt{q}$.*

The error bound above is optimal, and can be derived by proving a set of technical statements known as the *Weil Conjectures (for curves)*. The Weil Conjectures (along with corresponding point counts) were formulated for arbitrary varieties over finite fields and, in one of the crowning achievements of 20th century mathematics, were ultimately proved by Deligne in 1974 [Del74].

Efficient methods for computing $N_{p,1}(f)$ (and the number of points for a curve over any finite field) began to appear with the work of Schoof [Sch85], via so-called ℓ-adic methods. Let g denote the genus[2] of the curve $\mathcal{C}$. Via later work (e.g., [Pil90, AH01]) it was determined that $N_{p,1}(f)$ can be computed in time $(\log p)^{2^{g^{O(1)}}}$ for arbitrary curves. Kedlaya's algorithm [Ked01] then lowered this complexity bound to $(g^4 p)^{1+o(1)}$ for hyperelliptic curves, e.g., curves with defining polynomials of the form $x_2^2 - g(x_1)$. Kedlaya observed later that, on a quantum computer, one could compute (finite field) zeta functions for non-singular curves in time $(d \log p)^{O(1)}$ [Ked06]. (The precise definition of these zeta functions need not concern us

[2]The precise definition of the genus need not concern us, so we will simply recall that it is a birational invariant of $\mathcal{C}$ (i.e., it is invariant under rational maps with rational inverse) and is at most $(d-1)(d-2)/2$ for $\mathcal{C}$ the zero set of a degree d bivariate polynomial.

here: Suffice it to say that the computation of the zeta function of a curve over a finite field includes the computation of $N_{p,1}(f)$ as a special case.) More recently, Harvey [Har15] gave an efficient (classical) deterministic algorithm which, although asymptotically slower than Kedlaya's quantum algorithm, allows arbitrary input polynomials.

2.2. **The Central Recurrence for Bivariate Point Counting.** In this section, we generalize the tools we used for root counting for univariate polynomials in [KRRZ19] to point counting for curves. It is not hard to see that these tools extend naturally to point counting for hypersurfaces of arbitrary dimension. The only subtlety is maintaining low computational complexity and keeping track of the underlying singular locus.

Let $\mathbf{x} := (x_1, x_2)$ denote the tuple of two variables, and let $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$ be a bivariate polynomial with integer coefficients of total degree $d \geq 1$. Then for any $\zeta := (\zeta_1, \zeta_2) \in \mathbb{Z}^2$, the Taylor expansion of f at ζ is $f(\mathbf{x}) = \sum_{j_1,j_2} \frac{D^{j_1,j_2} f(\zeta)}{j_1! j_2!}(x_1 - \zeta_1)^{j_1}(x_2 - \zeta_2)^{j_2}$, where j_1, j_2 are non-negative integers and $D^{j_1,j_2} f(\mathbf{x}) := \frac{\partial^{j_1+j_2}}{\partial x_1^{j_1} \partial x_2^{j_2}} f(\mathbf{x})$.

Let $\tilde{f}(\mathbf{x}) := (f(\mathbf{x}) \mod p)$ denotes the mod p reduction of f. Now let $\zeta = (0,0)$ and write $\tilde{f} = g_m + g_{m+1} + \cdots + g_n$ where g_i is a (homogeneous) form in $\mathbb{F}_p[\mathbf{x}]$ of degree i and $g_m \neq 0$. We then define m to be the *multiplicity of $\tilde{f}$ at $\zeta = (0,0)$*. Write $m = m_\zeta(\tilde{f})$. To extend this definition to a point $\zeta = (a,b) \neq (0,0)$, let T be the translation that takes $(0,0)$ to ζ, i.e. $T(x_1, x_2) = (x_1 + a, x_2 + b)$. Then $\tilde{f}^T := \tilde{f}(x_1 + a, x_2 + b)$ and we define $m_\zeta(\tilde{f}) := m_{(0,0)}(\tilde{f}^T)$. Then it is immediate from the definition that:

Lemma 2.2. *If $\tilde{f} = \prod \tilde{f}_r^{e_r} \in \mathbb{F}_p[\mathbf{x}]$ is a factorization of $\tilde{f}$ into irreducible polynomials over $\mathbb{F}_p$ then $m_\zeta(\tilde{f}) = \sum m_\zeta(\tilde{f}_r)$.*

We say ζ is a *smooth point* of $\tilde{f}$ if $m_\zeta(\tilde{f}) = 1$, and call it a *singular point* otherwise. In particular, by Lemma 2.2, a point ζ is a smooth point of $\tilde{f}$ if and only if ζ belongs to just one irreducible component $\tilde{f}_r$ of $\tilde{f}$, the corresponding exponent $e_r = 1$, and ζ is a smooth point of $\tilde{f}_r$.

Now we are ready to generalize the tools in [KRRZ19] for curves:

Definition 2.3. *Let $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$ and fix a prime p. Let $\mathrm{ord}_p : \mathbb{Z} \longrightarrow \mathbb{N} \cup \{0\}$ denote the usual p-adic valuation with $\mathrm{ord}_p(p) = 1$. We then define $s(f, \varepsilon) := \min_{j_1,j_2 \geq 0} \left\{ j_1 + j_2 + \mathrm{ord}_p \frac{D^{j_1,j_2} f(\varepsilon)}{j_1! j_2!} \right\}$ for any $\varepsilon \in \{0, \ldots, p-1\}^2$. Finally, fixing $k \in \mathbb{N}$, let us inductively define a set $T_{p,k}(f)$ of pairs $(f_{i,\zeta}, k_{i,\zeta}) \in \mathbb{Z}[\mathbf{x}] \times \mathbb{N}$ as follows: We set $(f_{0,0}, k_{0,0}) := (f, k)$. Then, for any $i \geq 1$ with $(f_{i-1,\mu}, k_{i-1,\mu}) \in T_{p,k}(f)$ and any* singular *point $\zeta_{i-1} \in (\mathbb{Z}/p\mathbb{Z})^2$ of $\tilde{f}_{i-1,\mu}$ with $s_{i-1} := s(f_{i-1,\mu}, \zeta_{i-1}) \in \{2, \ldots, k_{i-1,\mu} - 1\}$, we define $\zeta := \mu + p^{i-1}\zeta_{i-1}$, $k_{i,\zeta} := k_{i-1,\mu} - s_{i-1}$ and $f_{i,\zeta}(\mathbf{x}) := \left[\frac{1}{p^{s_{i-1}}} f_{i-1,\mu}(\zeta_{i-1} + p\mathbf{x}) \right] \mod p^{k_{i,\zeta}}$.*

Just as in the univariate case, the perturbations $f_{i,\zeta}$ of f will help us keep track of how the points of f in $(\mathbb{Z}/p^k\mathbb{Z})^2$ cluster, in a p-adic metric sene, about the points of $\tilde{f}$. It is clear that $\frac{D^{j_1,j_2} f(\varepsilon)}{j_1! j_2!}$ is always an integer as the coefficient of $x_1^{j_1} x_2^{j_2}$ in the Taylor expansion of $f(\mathbf{x} + \varepsilon)$ about $\mathbf{x} = (0,0)$. We will see in the next section how $T_{p,k}(f)$ is associated with a natural tree structure. Moreover, $T_{p,k}(f)$ is always a finite set by definition, as only $f_{i,\zeta}$ with $i \leq \lfloor (k-1)/2 \rfloor$ and $\zeta \in (\mathbb{Z}/p\mathbb{Z})^2$ are possible.

Lemma 2.4. *Following the notation above, let $n_p(f)$ denote the number of smooth points of $\tilde{f}$ in $(\mathbb{Z}/p\mathbb{Z})^2$. Then provided $k \geq 0$ and $\tilde{f}$ is not identically zero, we have*

$$N_{p,k}(f) \;=\; p^{k-1}n_p(f) + \left(\sum_{\substack{\zeta_0\in(\mathbb{Z}/p\mathbb{Z})^2\\ s(f,\zeta_0)\geq k}} p^{2(k-1)}\right) + \sum_{\substack{\zeta_0\in(\mathbb{Z}/p\mathbb{Z})^2\\ s(f,\zeta_0)\in\{2,\dots,k-1\}}} p^{2(s(f,\zeta_0)-1)}N_{p,k-s(f,\zeta_0)}(f_{1,\zeta_0}).$$

We will prove Lemma 2.4 in the next section, where it will be clear how Lemma 2.4 applies recursively. Then we show how Lemma 2.4 leads to our recursive algorithm for computing $N_{p,k}(f)$.

3. Generalized Hensel Lifting and the Proof of our Main Recurrence

Let us first prove the following alternative definition for multiplicity of a point on the curve. We will mainly use this definition for the rest of the discussion.

Lemma 3.1. *For any $\zeta \in \mathbb{F}_p^2$, $m := m_\zeta(\tilde{f})$ is the smallest nonnegative integer such that there exists $j_1, j_2 \geq 0$ with $j_1 + j_2 = m$, and $D^{j_1,j_2}f(\zeta) \neq 0 \mod p$.*

Proof. Fix $\zeta \in \mathbb{F}_p^2$, and let T be the translation that takes $(0,0)$ to ζ. Then for any $j_1, j_2 \geq 0$, $D^{j_1,j_2}\tilde{f}^T(0,0) = D^{j_1,j_2}\tilde{f}(\zeta)$. So it suffices to prove the statement for the case when $\zeta = (0,0)$.

Suppose $\tilde{f} = g_m + g_{m+1} + \cdots + g_n$, where g_i is a homogeneous form in $\mathbb{F}_p[\mathbf{x}]$ of degree i and $g_m \neq 0$. Then $\tilde{f}$ must have a nonzero monomial term $a_r x_1^r x_2^{m-r}$, for some integer $r \leq m$, and $a_r \in \mathbb{F}_p^\times$. Note that as $h_m \in \mathbb{F}_p[\mathbf{x}]$, we must have $r, m-r < p$ as well. Then for any $j_1, j_2 \geq 0$, we have $D^{j_1,j_2}\left(a_r x_1^r x_2^{m-r}\right) = a_r\binom{r}{r-j_1}\binom{m-r}{m-r-j_2}x_1^{r-j_1}x_2^{m-r-j_2}$. It is obvious that for any pair of nonnegative integers j_1, j_2 with $j_1 + j_2 < m$, either $r - j_1 > 0$ or $m - r - j_2 > 0$. Moreover, any other nonzero monomial term $a_t x_1^{t_1}x_2^{t_2}$ of $\tilde{f}$ must have $t_1 + t_2 \geq m$ and $t_1 \geq r$ or $t_2 \geq m - r$. Hence $t_1 - j_1 > 0$ or $t_2 - j_2 > 0$. So for such a pair of j_1, j_2, we must have $D^{j_1,j_2}\tilde{f}(0,0) = 0 \mod p$. Now take $j_1 = r$ and $j_2 = m - r$, then

$$D^{j_1,j_2}\tilde{f}(0,0) = a_r\binom{r}{r-j_1}\binom{m-r}{m-r-j_2} \neq 0 \mod p.$$

Conversely, if m is the smallest nonnegative integer such that there exists $j_1, j_2 \geq 0$ with $j_1 + j_2 = m$ and $D^{j_1,j_2}f(0,0) \neq 0 \mod p$, then there exists $a_j x_1^{j_1}x_2^{j_2}$ a nonzero monomial term of $\tilde{f}$ of smallest total degree. So $m = m_{(0,0)}(\tilde{f})$. $\square$

The classical Hensel's Lemma (see, e.g., [NZM91, Thm. 2.3, Pg. 87]) says that any *non-degenerate* root of a univariate polynomial in $\mathbb{Z}/p\mathbb{Z}$ lifts uniquely into any larger prime power ring $\mathbb{Z}/p^k\mathbb{Z}$. One expects similar nice behavior from a smooth point on a curve over $\mathbb{Z}/p\mathbb{Z}$. We prove the following analogue of Hensel's Lemma for curves in the Appendix:

Lemma 3.2. *Let $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$. If $f(\sigma) \equiv 0 \mod p^j$ for $j \geq 1$, and $\left(\zeta^{(0)} \equiv \sigma \mod p\right)$ is a smooth point on $\tilde{f}$, then there are exactly p many $\mathbf{t} \in (\mathbb{Z}/p\mathbb{Z})^2$ such that $f(\sigma + p^j\mathbf{t}) \equiv 0 \mod p^{j+1}$.*

For $k > j \geq 1$ and any $\sigma^{(j)} \in (\mathbb{Z}/p^j\mathbb{Z})^2$ such that $f(\sigma^{(j)}) \equiv 0 \mod p^j$, we call $\sigma^{(k)} \in (\mathbb{Z}/p^k\mathbb{Z})^2$ a *lift* of $\sigma^{(j)}$, if $f(\sigma^{(k)}) \equiv 0 \mod p^k$ and $\sigma^{(k)} \equiv \sigma^{(j)} \mod p^j$. Then by applying Lemma 3.2 inductively, we obtain:

Proposition 3.3. *Let $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$, and $k > j \geq 1$. If $f(\sigma^{(j)}) \equiv 0 \mod p^j$, and $(\sigma^{(j)} \mod p)$ is a smooth point of $\tilde{f}$, then $\sigma^{(j)}$ lifts to exactly p^{k-j} many roots of $(f \mod p^k)$.*

Lemma 3.4. *Following the notation above, suppose instead $\zeta^{(0)} \in (\mathbb{Z}/p\mathbb{Z})^2$ is a point on $\tilde{f}$ of (finite) multiplicity $m \geq 2$. Suppose also that $k \geq 2$ and that there is a $\sigma^{(k)} \in (\mathbb{Z}/p^k\mathbb{Z})^2$ with $\sigma^{(k)} \equiv \zeta^{(0)} \mod p$ and $f(\sigma^{(k)}) = 0 \mod p^k$. Then $s(f, \zeta^{(0)}) \in \{2, \ldots, m\}$.*

Proof. As $\zeta^{(0)}$ is a singular point on $\tilde{f}$, then $\frac{\partial f}{\partial x_i}(\zeta^{(0)}) = 0 \mod p$ for every $i = 1, \ldots, n$. Then for $\sigma^{(k)} = \zeta^{(0)} + p\tau \in (\mathbb{Z}/p^k\mathbb{Z})^2$ with $\tau := (\tau_1, \tau_2) \in (\mathbb{Z}/p^{k-1}\mathbb{Z})^2$,

$$f(\sigma^{(k)}) = f(\zeta^{(0)}) + p\left(\frac{\partial f}{\partial x_1}(\zeta^{(0)})\tau_1 + \frac{\partial f}{\partial x_2}(\zeta^{(0)})\tau_2\right) + \sum_{i_1+i_2 \geq 2} p^{i_1+i_2} D^{i_1+i_2} f(\zeta^{(0)})\tau_1^{i_1}\tau_2^{i_2} \tag{1}$$

to have solutions mod p^k, we need $f(\zeta^{(0)}) \equiv 0 \mod p^2$, as the second and the third summand in equation (1) has p-adic order at least 2.

As $\zeta^{(0)}$ is a singular point of multiplicity m on $\tilde{f}$, there exists an m-th Hasse derivative: $D^{j_1,j_2} f(\zeta^{(0)}) \neq 0 \mod p$ with $j_1 + j_2 = m$. So $s(f, \zeta^{(0)}) \leq \mathrm{ord}_p\left(p^{j_1+j_2} D^{j_1,j_2} f(\zeta^{(0)})\right) = m$. $\square$

We can now relate $N_{p,k}(f)$ to the recursive structure on $T_{p,k}(f)$.

Proof of Lemma 2.4: The lifting of smooth points of $\tilde{f}$ follows from Proposition 3.3.

Now assume that $\zeta_0 \in (\mathbb{Z}/p\mathbb{Z})^2$ is a singular point of $\tilde{f}$. Write $\zeta := \zeta_0 + p\sigma$ for $\sigma := \zeta_1 + p\zeta_2 + \cdots + p^{k-2}\zeta_{k-1} \in (\mathbb{Z}/p^k\mathbb{Z})^2$, and let $s := s(f, \zeta_0)$. Note that by Lemma 3.4, $s \geq 2$. Then by definition, $f(\zeta) = p^s f_{1,\zeta_0}(\sigma)$, for $f_{1,\zeta_0} \in \mathbb{Z}[\mathbf{x}]$ and f_{1,ζ_0} does not vanish identically mod p.

If $s \geq k$, then $f(\zeta) = 0 \mod p^k$ regardless of choice of σ. So there are exactly $p^{2(k-1)}$ values of $\zeta \in (\mathbb{Z}/p^k\mathbb{Z})^2$ such that $\zeta \equiv \zeta_0 \mod p$ and $f(\zeta) = 0 \mod p^k$.

If $s \leq k-1$ then ζ is a root of f if and only if $f_{1,\zeta_0}(\sigma) \equiv 0 \mod p^{k-s}$. But then $\sigma = \zeta_1 + p\zeta_2 + \ldots + p^{k-s-1}\zeta_{k-s} \mod p^{k-s}$, i.e., the rest of the base p digits $\zeta_{k-s+1}, \ldots, \zeta_{k-1}$ do not appear in the preceding mod p^{k-s} congruence. So the number of possible lifts ζ of ζ_0 is exactly $p^{2(s-1)}$ times the number of roots $(\zeta_1 + p\zeta_2 + \ldots + p^{k-s-1}\zeta_{k-s}) \in (\mathbb{Z}/p^{k-s}\mathbb{Z})^2$ of f_{1,ζ_0}. This accounts for the third summand in our formula. $\square$

Remark 3.5. *The algebraic preliminaries we concluded in this section and Definition 2.3 can be extended transparently for point counting for hypersurfaces of arbitrary dimensions.* $\diamond$

4. Bounding Sums of Multiplicities on Curves with at Worst Isolated Singularities

Suppose $F \in \mathbb{F}_p[\mathbf{x}]$ is a nonconstant polynomial of total degree D. Then F factors into a product of irreducible components $F = \prod_{i=1}^{l} F_i^{e_i} \in \mathbb{F}_p[\mathbf{x}]$ where each $F_i \in \mathbb{F}_p[\mathbf{x}]$ is irreducible, and $e_i \geq 1$. We say F is *squarefree* if $e_i = 1$ for every i. Suppose $G = \prod_{j=1}^{m} G_j^{c_i} \in \mathbb{F}_p[\mathbf{x}]$ with $G_i \in \mathbb{F}_p[\mathbf{x}]$ irreducible and $c_i \geq 1$. We say F and G *have no common component*, if $F_i \neq G_j$ for every pair of i, j.

Lemma 4.1. (Corollary of Bézout's Theorem) *Let $F, G \in \mathbb{F}_p[\mathbf{x}]$ be two curves with no common components, then $\sum_\zeta m_\zeta(F) m_\zeta(G) \leq \deg(F)\deg(G)$.*

Lemma 4.2. *Let $F \in \mathbb{F}_p[\mathbf{x}]$ be squarefree with degree d, then $\sum_\zeta m_\zeta(F)\left(m_\zeta(F) - 1\right) \leq d(d-1)$.*

Proof. As F is squarefree, then F and $D^{1,0}F(\mathbf{x})$ have no common component. It is also easy to deduct from Lemma 3.1 that for any $\zeta \in \mathbb{F}_p^2$, $m_\zeta(D^{1,0}F) \geq m_\zeta(F) - 1$. The conclusion thus follows by applying Lemma 4.1. □

Suppose $F = \prod_{i=1}^{l} F_i^{e_i} \in \mathbb{F}_p[\mathbf{x}]$ is a nonconstant polynomial. For each i, let $d_i := \deg(F_i)$ and let $d := \sum d_i^{e_i}$ be the total degree of F. Let $I \subseteq \{1, \ldots, l\}$ be an nonempty subset of indices, and let S_I denote the set of points in the intersection $\bigcap_{i\in I} F_i$, and let $T_I = \{\zeta \in S_I : \zeta \text{ is smooth on } F_i \text{ for all } i \in I\}$.

We then prove the following more generalized statement of Lemma 4.2 in the Appendix:

Lemma 4.3. *Using the notation above we have:*

$$\sum_{\substack{\zeta\in S_I\\ I\neq\emptyset}} m_\zeta(F)(m_\zeta(F) - \sum_{i\in I} e_i) + \sum_{\substack{\zeta\in T_I\\ |I|\geq 2}} m_\zeta(F) \leq d(d-1). \tag{2}$$

Observe that if $\zeta \in S_I$ and ζ is an isolated singular points on F, then either $\zeta \in T_I$ or $m_\zeta(F) > \sum_{i\in I} \mu_\zeta(F_i)$, and $m_\zeta(F) = \sum_{i\in I} \mu(F_i)$ if it is non-isolated. So only the part corresponding to the isolated singular points contribute to the sum on the left hand side of Equation 2. So we obtain the following:

Theorem 4.4. *Let $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$ be a nonconstant polynomial of degree d. Fix a prime p and suppose that $\tilde{f}$ does not vanish identically over $\mathbb{Z}/p\mathbb{Z}$. Then* $\sum_{\substack{\zeta \text{ isolated}\\ \text{singular on } \tilde{f}}} \deg \tilde{f}_{1,\zeta} \leq d(d-1)$.

Proof. This is immediate by observing that $\deg \tilde{f}_{1,\zeta} \leq s(f,\zeta) \leq m_\zeta(\tilde{f})$. □

However, bounding the degree of the perturbations $\tilde{f}_{1,\zeta}$ corresponding to *non-isolated singular points* of $\tilde{f}$ can be hard. This is evident in the discussion in the final section of the Appendix: lifting non-isolated singular points for certain families of curves requires extra care.

4.1. **Algorithms and Complexity Analysis: Proof of Theorem 1.1.** For this section, let us consider bivariate polynomials $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$ of the form $f(\mathbf{x}) = g(x_1) + h(x_2)$. One broad family of examples of such bivariate polynomials is the family of superelliptic curves: $f(\mathbf{x}) = x_2^d - g(x_1)$.

Lemma 4.5. *Let $F(x_1, x_2) = g(x_1) + h(x_2) \in \mathbb{F}_p[\mathbf{x}]$ such that g, h are nonconstant polynomials. Then F is squarefree.*

Proof. Suppose F is not squarefree and let $F = \prod_{i=1}^{l} F_i^{e_i} \in \mathbb{F}_p[\mathbf{x}]$ be the irreducible factorization of F, and $e_i \geq 1$. Without loss of generality assume $e_1 > 1$ and $g'(x_1) = D^{1,0}F \neq 0$. Let $G = F/F_1^{e_1} = \prod_{i=2}^{l} F_i^{e_i}$. Differentiating F with respect to x_1, we have

$$\begin{aligned} g'(x_1) &= e_1 F_1^{e_1-1} D^{1,0}F \cdot G + F_1^{e_1} \cdot D^{1,0}G \\ &= F_1^{e_1-1}\left(e_1 D^{1,0}F \cdot G + F_1 \cdot D^{1,0}G\right). \end{aligned}$$

So $F_1(x_1, x_2)$ must divide $g'(x_1)$, implying that $h(x_2)$ is a constant, a contradiction. □

We now have enough ingredients to state our main algorithm:

Algorithm 4.6 ($\texttt{PrimePowerPointCounting}(f, p, k)$).
Input. $(f, p, k) \in \mathbb{Z}[\mathbf{x}] \times \mathbb{N} \times \mathbb{N}$ with p prime and $f(\mathbf{x}) = g(x_1) + h(x_2)$.
Output. An integer $M \leq N_{p,k}(f)$ that, with probability at least $\frac{2}{3}$, is exactly $N_{p,k}(f)$.
Description.

Let $v := s(f)$ and $f_{0,0} := f$.
If $v \geq k$
    Let $M := p^{2k}$. **Return**.
If $v \in \{1, \dots, k-1\}$
    Let $M := p^{2v}\texttt{PrimePowerPointCounting}\left(\frac{f_{0,0}(\mathbf{x})}{p^v}, p, k-v\right)$. **Return**.
End(If).
If $s(g) = s(h) = 0$
    Let $M := p^{k-1} n_p(f)$.
    For $\zeta^{(0)} \in (\mathbb{Z}/p\mathbb{Z})^2$ a singular point of $\tilde{f}_{0,0}$ **do**[2]
        Let $s := s(f_{0,0}, \zeta^{(0)})$.
        If $s \geq k$
            Let $M := M + p^{2(k-1)}$.
        Elseif $s \in \{2, \dots, k-1\}$
            Let $M := M + p^{2(s-1)}\texttt{PrimePowerPointCounting}\left(f_{1,\zeta^{(0)}}, p, k-s\right)$.
        End(If).
    End(For).
Elseif $s(g) \geq 0$ or $s(h) \geq 0$ accordingly
    Let $M := p^k n_p(g)$ or $p^k n_p(h)$.
    For $\zeta^{(0)} \subseteq (\mathbb{Z}/p\mathbb{Z})^2$ a set of singular points of $\tilde{f}_{0,0}$ from a degenerate root of $\tilde{g}$ or $\tilde{h}$ **do**
        Let $s := s(f_{0,0}, \zeta^{(0)})$.
        If $s \geq k$
            Let $M := M + p^{2k-1}$.
        Elseif $s \in \{2, \dots, k-1\}$
            Let $M := M + p^{2s-1}\texttt{PrimePowerPointCounting}\left(f_{1,\zeta^{(0)}}, p, k-s\right)$.
        End(If).
    End(For).
End(If).
Return.

There are some remaining details to clarify about our algorithm. First, let $s(f)$ denote the largest power of p that divides all the coefficients of f. Then by Definition 2.3, we see that any polynomial in $T_{p,k}(f)$ should also be of the form $g(x_1) + h(x_2)$ with $s(g) = 0$ or $s(h) = 0$. By Lemma 4.5, we see that when $s(g) = s(h) = 0$, then $\tilde{f} \mod p$ is squarefree. Now without loss of generality, suppose $0 = s(g) < s(h) = c$, then $\tilde{f}(\mathbf{x}) = \tilde{g}(x_1) \mod p$. Then any singular point on $\tilde{f}$ should be of the form $(\zeta_1^{(0)}, y)$ for any *degenerate root* $\zeta_1^{(0)}$ of the univariate polynomial $\tilde{g}(x_1) \in \mathbb{F}_p[x_1]$ and any choice of $y \in \{0, 1, \dots, p-1\}$. So it makes sense to consider the perturbation of f in the direction of x_1 only.

Let $\zeta_1^{(0)}$ be any degenerate root of $\tilde{g}$. Abusing notation, let $\zeta^{(0)} := \{\zeta_1^{(0)}\} \times \mathbb{F}_p = \{(\zeta_1^{(0)}, y) : y \in \{0, 1, \dots, p-1\}\}$, the set of singular points of $\tilde{f}$ with the first coordinate being $\zeta_1^{(0)}$. Consider $f(\zeta_1^{(0)} + px_1, x_2) = g(\zeta_1^{(0)} + px_1) + h(x_2)$. Let $s(f, \zeta^{(0)}) := s(f(\zeta_1^{(0)} + px_1, x_2)) = \min(s(g, \zeta_1^{(0)}), c)$, the largest p's power dividing all the coefficients of the perturbation, and let $f_{1,\zeta^{(0)}} = \frac{1}{p^{s(f,\zeta^{(0)})}} f(\zeta_1^{(0)} + px_1, x_2)$.

We prove the following more specific version of Lemma 2.4 in the Appendix:

Lemma 4.7. *Let $f(\mathbf{x}) = g(x_1) + h(x_2)$ with $0 = s(g) < s(h) = c$. Let $n_p(g)$ denote the number of non-degenerate root of $\tilde{g}$ in $\mathbb{F}_p$, and following the notation above:*

$$N_{p,k} = p^k n_p(g) + \left(\sum_{\substack{\zeta^{(0)} \subseteq (\mathbb{Z}/p\mathbb{Z})^2 \\ s(f,\zeta^{(0)}) \geq k}} p^{2k-1} \right) + \sum_{\substack{\zeta^{(0)} \subseteq (\mathbb{Z}/p\mathbb{Z})^2 \\ s(f,\zeta^{(0)}) \leq k-1}} p^{2s(f,\zeta^{(0)})-1} N_{p,k-s(f,\zeta^{(0)})}(f_{1,\zeta^{(0)}})$$

By symmetry, a variant of our preceding lemma also holds when $0 = s(h) < s(g) = c$. Similarly, for any degenerate root $\zeta_2^{(0)}$ of the univariate polynomial $\tilde{h}(x_2) \in \mathbb{F}_p$, we denote $\zeta^{(0)} := \mathbb{F}_p \times \{\zeta_2^{(0)}\}$ to be the set of singular points of $\tilde{f}$ with the second coordinate being $\zeta_2^{(0)}$.

Notation 4.8. *Suppose $\zeta^{(i-1)} = \{\zeta_1^{(i-1)}\} \times \mathbb{F}_p$ is the set of singular points on $\tilde{f}_{i-1,\zeta}$ for some polynomial in $T_{p,k}(f)$ and $\zeta = (\zeta_1, \zeta_2)$, we write*

$$\zeta + p^{i-1}\zeta^{(i-1)} = \{(x_1, x_2) : x_1 = \zeta_1 + p^{i-1}\zeta_1^{(i-1)}, x_2 \in \{\zeta_2 + p^{i-1} \cdot 0, \ldots \zeta_2 + p^{i-1} \cdot p - 1\}\}$$

as element-wise operations for set. We also use this notation similarly when $\zeta^{(i-1)} = \mathbb{F}_p \times \{\zeta_2^{(i-1)}\}$.

We are now ready to prove the correctness of our main algorithm.

Proof of Correctness of Algorithm 4.6: Assume temporarily that Algorithm 4.6 is correct when $s(f) = 0$, i.e. when $f_{0,0}$ is *not* identically 0 mod p. Since for any integers a with $a \leq k$, and any elements $\mathbf{x}, \mathbf{y} \in (\mathbb{Z}/p^k\mathbb{Z})^2$, $p^a\mathbf{x} = p^a\mathbf{y} \bmod p^k \iff \mathbf{x} = \mathbf{y} \bmod p^{k-a}$, Steps 1–6 of our algorithm then dispose of the case where f is identically 0 in $(\mathbb{Z}/p\mathbb{Z})[\mathbf{x}]$. So let us now prove correctness when f is *not* identically 0 in $(\mathbb{Z}/p\mathbb{Z})[\mathbf{x}]$.

Recall from the discussion at the very beginning of this section, we see that any polynomial in $T_{p,k}(f)$ should be of the form $f_{i,\zeta^{(i-1)}}(\mathbf{x}) := g_i(x_1) + h_i(x_2)$ with $s(g_i) = 0$ or $s(h_i) = 0$. Applying Lemma 2.4 and Lemma 4.7 accordingly, we then see that it is enough to prove that the value of M is the value of our formula for $N_{p,k}(f)$ when the two **For** loops of Algorithm 4.6 runs correctly.

When $s(g) = s(h) = 0$, Steps 7–16 (once the **For** loop is completed) then simply add the second and third summands of our formula in Lemma 2.4 to M thus ensuring that $M = N_{p,k}(f)$. On the other hand, when $s(g) > 0$ or $s(h) > 0$, Steps 17–26 (once the **For** loop is completed) handles add the second and third summands of our formula in Lemma 4.7 to M thus ensuring that $M = N_{p,k}(f)$. So we are done. □

In [KRRZ19], we defined a recursive tree structure for root counting for univariate polynomial in $\mathbb{Z}/p^k\mathbb{Z}$. We define similarly a recursive tree for $f(\mathbf{x}) = g(x_1) + h(x_2)$ that will enable our complexity analysis.

Definition 4.9. *Let us identify the elements of $T_{p,k}(f)$ with nodes of a lablled rooted directed tree $\mathcal{T}_{p,k}(f)$.*

1. *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 $\mathcal{T}_{p,k}(f)$.*
2. *There is an edge from node $(f_{i',\zeta'}, k_{i',\zeta'})$ to node $(f_{i,\zeta}, k_{i,\zeta})$ if and only if $i' = i - 1$ and there is a (set of) singular points $\zeta^{(i-1)}$ in $(\mathbb{Z}/p\mathbb{Z})^2$ of $\tilde{f}_{i',\zeta'}$ with $s(f_{i',\zeta'}, \zeta^{(i-1)}) \leq k_{i',\zeta'} - 1$ and $\zeta = \zeta' + p^{i-1}\zeta^{(i-1)}$ in $(\mathbb{Z}/p^i\mathbb{Z})^2$.*

(3) *Suppose* $f_{i',\zeta'} = g_{i'}(x_1) + h_{i'}(x_2)$. *The label of a directed edge from node* $(f_{i',\zeta'}, k_{i',\zeta'})$ *to node* $(f_{i,\zeta}, k_{i,\zeta})$ *is* $p^{2\left(s\left(f_{i',\zeta'},(\zeta-\zeta')/p^{i'}\right)-1\right)}$ *or* $p^{2s\left(f_{i',\zeta'},(\zeta-\zeta')/p^{i'}\right)-1}$ *respectively when* $s(g_{i'}) = s(h_{i'}) = 0$ *or otherwise.*

In particular, the labels of the nodes lie in $\mathbb{Z}[\mathbf{x}] \times \mathbb{N}$.

Remark 4.10.
1. Just as the tree structure for the univariate polynomial in [KRRZ19], *our trees* $\mathcal{T}_{p,k}(\cdot)$ *encode algebraic expressions for our desired root counts* $N_{p,k}(\cdot)$. *In particular, the children of a node labelled* (f_i, k_i) *yield terms that one sums to get the root count* $N_{p,k_i}(f_i)$, *and the edge labels yield weights multiplying the corresponding terms.*
2. One main difference is that the correspondence between polynomials in $T_{p,k}(f)$ *with the label in the tree* $\mathcal{T}_{p,k}(f)$ *is no longer one-to-one. In particular, in the case when* $f_{i,\zeta}(\mathbf{x}) = g_i(x_1) + h_i(x_2)$ *with* $s(g_i) > 0$, *its child node polynomial* $f_{i+1,\zeta'}$ *for* $\zeta' - \zeta = \{\zeta_1^{(i)}\} \times \mathbb{F}_p$, *correspond to* a set of *singular points of* $\tilde{f}_{i,\zeta}$ *with the first coordinate equaling to a degenerate root* $\zeta_1^{(i)}$ *of* $\tilde{g}_i$. $\diamond$

The following lemma, proved in the Appendix, will be central in our complexity analysis.

Lemma 4.11. *Let* $f(\mathbf{x}) = g(x_1) + h(x_2) \in \mathbb{Z}[\mathbf{x}]$ *be a nonconstant polynomial of degree* d. *Following the notation of Definition 4.9, we have that:*

(1) *The depth of* $\mathcal{T}_{p,k}(f)$ *is at most* k.
(2) *The degree of the root node of* $\mathcal{T}_{p,k}(f)$ *is at most* $\binom{d}{2}$.
(3) *The degree of any* non-*root node of* $\mathcal{T}_{p,k}(f)$ *labeled* $(f_{i,\zeta}, k_{i,\zeta})$, *with parent* $(f_{i-1,\mu}, k_{i-1,\mu})$ *and* $\zeta^{(i-1)} := (\zeta - \mu)/p^{i-1}$, *is at most* $s(f_{i-1,\mu}, \zeta^{(i-1)})$. *In particular,* $\deg \tilde{f}_{i,\zeta} \leq s(f_{i-1,\mu}, \zeta^{(i-1)}) \leq k_{i-1,\mu} - 1 \leq k - 1$ *and*

$$\sum_{\substack{(f_{i,\zeta},k_{i,\zeta}) \text{ a child} \\ \text{of } (f_{i-1,\mu},k_{i-1,\mu})}} \deg \tilde{f}_{i,\zeta}\left(\deg \tilde{f}_{i,\zeta} - 1\right) \leq \deg \tilde{f}_{i-1,\mu}\left(\deg \tilde{f}_{i-1,\mu} - 1\right)$$

(4) $\mathcal{T}_{p,k}(f)$ *has at most* $\binom{d}{2}$ *nodes at depth* $i \geq 1$, *and thus a total of no more than* $1 + (k-1)\binom{d}{2}$ *nodes.*

Proof of Theorem 1.1: Since we already proved that Algorithm 4.6, it suffices to prove the stated complexity bound for Algorithm 4.6. The proof consists of three parts: (a) the point counting algorithm over $\mathbb{F}_p$ from [Har15], (b) the univariate reduction and the factorization algorithm, and (c) applying Lemma 4.11 to show that the number of necessary factorization and point counting, and p-adic valuation calculations is well-bounded.

More specifically the **For** loops and the recursive calls of Algorithm 4.6 can be seen as the process of building the tree $\mathcal{T}_{p,k}(f)$. We begin at the root node by applying the algorithm in [Har15] to find the number of roots of $\tilde{f}$ in $\mathbb{F}_p$. This computation takes time $O(d^8 p^{1/2} \log^{2+\varepsilon} p)$ and space $O(d^4 p^{1/2} \log p)$ by [Har15]. (Specifically, one avails to Theorem 3.1, Lemmata 3.2 and 3.4, and Proposition 4.4 from Harvey's paper.)

To find singular points of $\tilde{f}$, it suffices to find the roots of the 2×2 polynomial system $F := (\tilde{f}(\mathbf{x}), D^{1,0}\tilde{f}(\mathbf{x}))$ over $\mathbb{F}_p$. This is done by first transforming the problem to factorization of a univariate polynomial U_F via univariate reduction over the finite field (see, e.g. [Roj99]). In particular $\deg U_F \leq d^2$ and roots of U_F will encode information on tuple (x_1, x_2) as solutions to the polynomial system F. Computing U_F can be done in time polynomial in the *mixed area* of the Newton polygons of F, and thus takes time $d^{O(1)}$. Then we use the fast randomized Kedlaya-Umans factoring algorithm in [KU08] to find solutions to U_F in $\mathbb{F}_p$, and thereby the singular points of $\tilde{f}$. This takes time $(d^3 \log p)^{1+o(1)} + (d^2 \log^2 p)^{1+o(1)}$ and requires $O(d^2 \log p)$ random bits.

In order to continue the recursion, we need to compute p-adic valuations of polynomial coefficients to determine $s(f_{0,0}, \zeta^{(0)})$ and the edges emanating from the root node. Expanding $f(\zeta^{(0)} + p\mathbf{x}) \mod p^k$ takes time no worse than $d^2(k \log p)^{1+o(1)}$ via Horner's method and fast finite ring arithmetic (see, e.g., [BS96, vzGG13]). Computing $s(f_{0,0}, \zeta^{(0)})$ thus takes time $d(k \log p)^{1+o(1)}$ by evaluating p-adic valuations using standard tools such as binary methods. By Assertion (2) of Lemma 4.11, there are no more than $\binom{d}{2}$ many such $\zeta^{(0)}$. So the total work so far is $(dk)^{O(1)} p^{1/2+\varepsilon}$. Note that computing $N_{p,1}(f)$ via algorithm in [Har15] dominates the computation.

The remaining work can also be well-bounded similarly by Lemma 4.11. In particular, the sum of the degress if $\tilde{f}_{i,\zeta}$ at level i of the tree $\mathcal{T}_{p,k}(f)$ is no greater than $\binom{d}{2}$.

Now observe that for $i \geq 2$, the amount of work needed to determine the polynomials at level i via computing $s(f_{i-1,\mu}, \zeta^{(i-1)})$ is no greater than $\binom{d}{2} d(k \log p)^{1+o(1)}$. By the basic inequality that $r_1^t + \cdots + r_\ell^t \leq (r_1 + \cdots + r_\ell)^t$ (valid for any $r, t \geq 1$), the total amount of work for point counting over $\mathbb{F}_p$, univariate reduction and factorization for each subsequent level of $\mathcal{T}_{p,k}(f)$ will be $(dk)^{O(1)} p^{1/2+\varepsilon}$ with $O(d^2 \log p)$ random bits needed. The expansion of the $f_{i,\zeta}$ at level i will take time no greater than $d^3(k \log p)^{1+o(1)}$ to compute. So the total work at each subsequent level is $(dk)^{O(1)} p^{1/2+\varepsilon}$.

Therefore the total amount of work for our tree will be $(dk)^{O(1)} p^{1/2+\varepsilon}$, and the number of random bits needed is $O(d^2 k \log p)$.

The argument proving the Las Vegas properties of our algorithm can be done similarly as in [KRRZ19]. In particular, we run factorization algorithm for sufficiently many times to reduce the *overall* error probability to less than $2/3$. Thanks to Lemma 4.11, it is enough to enforce a success probability of $O(\frac{1}{d^2 k})$ for each application of factorization, and to run the algorithm from [KU08] for $O(\log(dk))$ times for each time we need univariate factorization. So a total of $O(d^2 k \log(dk) \log p)$ many random bits is needed.

Our algorithm proceeds with building the tree structure $\mathcal{T}_{p,k}(f)$, so we only need to keep track of collections of $f_{i,\zeta}$. A bivariate polynomial of degree d with integer coefficients all of absolute value less than p^k requires $O(dk \log p)$ bits to store, and there are no more than $\binom{d}{2} k$ many polynomials in $\mathcal{T}_{p,k}(f)$. Combining with the space needed from algorithm in [Har15], we only need $O(d^4 k p^{1/2} \log p)$ space.

If $\tilde{f}$ defines a smooth and irreducible curve over the algebraic closure $\bar{\mathbb{F}}_p$ of $\mathbb{F}_p$ then the second part of the theorem follows immediately by combining our bivariate version of Hensel's Lemma (Lemma 3.2) with Kedlaya's quantum point counting algorithm from [Ked06]. $\square$

5. Appendix: Remaining Proofs and Finessing Exceptional Curves

5.1. **Proof of Lemma 3.2 (Higher-Dimensional Hensel's Lemma).** Consider the Taylor expansion of f at σ by $p^j\mathbf{x}$,

$$\begin{aligned} f(\sigma + p^j\mathbf{x}) &= f(\sigma) + p^j\left(\frac{\partial f}{\partial x_1}(\sigma)x_1 + \frac{\partial f}{\partial x_2}(\sigma)x_2\right) + \sum_{i_1+i_2\geq 2} p^{j(i_1+i_2)}D^{i_1,i_2}f(\sigma)x_1^{i_1}x_2^{i_2} \\ &= f(\sigma) + p^j\left(\frac{\partial f}{\partial x_1}(\sigma)x_1 + \frac{\partial f}{\partial x_2}(\sigma)x_2\right) \mod p^{j+1}, \end{aligned}$$

as $j(i_1+i_2) \geq j+1$ for all $i_1+i_2 \geq 2$. Then $\mathbf{t} := (t_1, t_2)$ is such that $(\sigma + \mathbf{t}p^j)$ is a solution to $f \equiv 0 \mod p^{j+1}$ if and only if

$$\frac{\partial f}{\partial x_1}(\sigma)t_1 + \frac{\partial f}{\partial x_2}(\sigma)t_n = -\frac{f(\sigma)}{p^j} \mod p. \tag{3}$$

As $(\zeta^{(0)} = \sigma \mod p)$ is a smooth point on $\tilde{f}$, then there exists an i such that $\frac{\partial f}{\partial x_i}(\sigma) = \frac{\partial f}{\partial x_i}(\zeta^{(0)}) \neq 0 \mod p$. Then left hand side of (3) does not vanish identically, and thus define a nontrivial linear relation in $(\mathbb{Z}/p\mathbb{Z})^2$. So fixing ζ, there are exactly p many $\mathbf{t} \in (\mathbb{Z}/p\mathbb{Z})^2$ satisfying (3). $\square$

5.2. **The Proof of Lemma 4.3.** We prove by induction on the number of irreducible components of F.

When $l = 1$, $F = F_1^{e_1}$. By Lemma 2.2, $m_\zeta(F) = e_1 m_\zeta(F_1)$ for every $\zeta \in \mathbb{F}_p^2$. Then by Lemma 4.2 and expanding

$$\sum_{\zeta \text{ on } F_1} \frac{m_\zeta(F)}{e_1}\left(\frac{m_\zeta(F)}{e_1} - 1\right) \leq d_1(d_1-1),$$

the conclusion holds.

Now suppose the inequality holds for $l-1 > 1$, and let $F' = \prod_{i=1}^{l-1} F_i^{e_i}$ and d' be its degree, and F_l is irreducible and has no common component with F'. Then $\sum_{\zeta \text{ on } F_l} m_\zeta(F_l^{e_l})\left(m_\zeta(F_l^{e_l}) - e_l\right) \leq e_l d_l(e_l d_1 - e_l)$, and

$$\sum_{J\subseteq\{1,\ldots,l-1\}}\left[\sum_{\zeta\in S_J} m_\zeta(F')\left(m_\zeta(F') - \sum_{j\in J} e_j\right) + \sum_{\substack{\zeta\in T_J \\ |J|\geq 2}} m_\zeta(F')\right] \leq d'(d'-1)$$

By Lemma 4.1, we must have $\sum_\zeta m_\zeta(F')m_\zeta(F_l^{e_l}) \leq d' d_l e_l$. Summing over all $J \subseteq \{1,\ldots,l-1\}$, we have

$$\begin{aligned} &\sum_J\left[\sum_{\zeta\in S_J} m_\zeta(F')\left(m_\zeta(F') - \sum_{j\in J} e_j\right) + \sum_{\substack{\zeta\in T_J \\ |J|\geq 2}} m_\zeta(F')\right] + 2\sum_\zeta m_\zeta(F')m_\zeta(F_l^{e_l}) + \sum_{\zeta \text{ on } F_l} m_\zeta(F_l^{e_l})\left(m_\zeta(F_l^{e_l}) - e_l\right) \\ &\leq d'(d'-1) + 2d'd_l e_l + (d_l e_l)^2 - e_l^2 d_l \leq (d' + d_l e_l)^2 - d' - e_l^2 d_l \leq d(d-1). \end{aligned}$$

Note that for each $J \subseteq \{1, \ldots, l-1\}$ and each $\zeta \in S_J$ such that ζ is not a point of F_l, $m_\zeta(F') = m_\zeta(F)$. If $\zeta \in S_{J\cup\{l\}} \backslash T_{J\cup\{l\}}$, then $m_\zeta(F_l^{e_l}) + m_\zeta(F') > e_l + \sum_{j\in J} e_j$, and

$$\begin{aligned}
&m_\zeta(F')(m_\zeta(F') - \sum_{i\in J} e_i) + 2m_\zeta(F_l^{e_l})m_p(F') + m_\zeta(F_l^{e_l})(m_\zeta(F_l^{e_l}) - e_l) \\
&= (m_\zeta(F') + m_\zeta(F_l^{e_l}))^2 - \sum_{i\in J} e_i m_\zeta(F') - e_l m_\zeta(F_l^{e_l}) \\
&\geq m_\zeta(F)(m_\zeta(F) - \sum_{i\in J\cup\{l\}} e_i)
\end{aligned}$$

So we can rewrite

$$\begin{aligned}
A :=& \sum_{J\subseteq\{1,\ldots,l-1\}} \left[\sum_{\zeta\in S_J} m_\zeta(F')(m_\zeta(F') - \sum_{j\in J} e_j) + 2\sum_{\zeta\notin T_{J\cup\{l\}}} m_\zeta(F')m_\zeta(F_l^{e_l}) \right] + \sum_{\zeta\in S_{\{l\}}} m_\zeta(F_l^{e_l})\left(m_\zeta(F_l^{e_l}) - e_l\right) \\
\geq& \sum_{J\subseteq\{1,\ldots,l-1\}} \left[\sum_{\zeta\in S_J} m_\zeta(F)(m_\zeta(F) - \sum_{j\in J} e_j) + \sum_{\zeta\in S_{J\cup\{l\}}} m_\zeta(F)(m_\zeta(F) - \sum_{j\in J\cup\{l\}} e_i) \right] + \sum_{\zeta\in S_{\{l\}}} m_\zeta(F)(m_\zeta(F) - e_l) \\
=& \sum_{I\in\{1,\ldots,l\}} \sum_{\zeta\in I} m_\zeta(F)(m_\zeta(F) - \sum_{i\in I} e_i).
\end{aligned}$$

On the other hand, if $\zeta \in T_{J\cup\{l\}}$, we must have $m_\zeta(F_l^{e_l}) + m_\zeta(F') = e_l + \sum_{j\in J} e_j$. Then summing over all $J \subseteq \{1, \ldots, l-1\}$, and

$$\begin{aligned}
B :=& \sum_J \sum_{\substack{\zeta\in T_J \\ |J|\geq 2}} m_\zeta(F') + 2\sum_{\zeta\in T_{J\cup\{l\}}} m_\zeta(F')m_\zeta(F_l^{e_l}) \\
=& \sum_J \left[\sum_{\substack{\zeta\in T_J \\ |J|\geq 2}} m_\zeta(F) + 2\sum_{\substack{\zeta\in T_{J\cup\{l\}} \\ |J|\geq 2}} m_\zeta(F')m_\zeta(F_l^{e_l}) \right] + \sum_{i=1}^{l-1} \sum_{\zeta\in T_{\{i,l\}}} m_\zeta(F')m_\zeta(F_l^{e_l}) \\
\geq& \sum_J \left[\sum_{\substack{\zeta\in T_J \\ |J|\geq 2}} m_\zeta(F) + \sum_{\substack{\zeta\in T_{J\cup\{l\}} \\ |J|\geq 2}} m_\zeta(F) \right] + \sum_{i=1}^{l-1} \sum_{\zeta\in T_{\{i,l\}}} m_\zeta(F') = \sum_I \sum_{\substack{\zeta\in T_I \\ |I|\geq 2}} m_\zeta(F).
\end{aligned}$$

The last inequality holds because for $a, b \geq 1$, we must have $2ab \leq a + b$.

Combining all of above computations, we have

$$\sum_I \left[\sum_{\zeta\in S_I} m_\zeta(F)(m_\zeta(F) - \sum_{i\in I} e_i) + \sum_{\substack{\zeta\in T_I \\ |I|\geq 2}} m_\zeta(F) \right] \leq A + B \leq d(d-1).$$

The conclusion thus follows. $\square$

5.3. **The Proof of Lemma 4.11.**

Assertion (1): By Definitions 2.3 and 4.9, each $(f_{i,\zeta}, k_{i,\zeta})$ whose parent node is $(f_{i-1,\mu}, k_{i-1,\mu})$, must satisfies $1 \leq k_{i-1,\mu} - k_{i,\zeta} \leq k_{i-1,\mu} - 1$, and $1 \leq k_{i,\zeta} \leq k - 1$ for all $i \geq 1$. So considering any root to leaf path in $\mathcal{T}_{p,k}(f)$, it is clear that the depth of $\mathcal{T}_{p,k}(f)$ can be no greater than $1+(k-1)=k$.

Assertion (2): If $s(g)=s(h)=0$, then by Lemma 4.5, $\tilde{f}(\mathbf{x})\in\mathbb{F}_p[\mathbf{x}]$ is square-free. As the multiplicity of any singular point is at least 2, by Lemma 4.2, $\tilde{f}$ has at most $\binom{d}{2}$ many singular points. In this case, each edge emanating from the root of $\mathcal{T}_{p,k}(f)$ corresponds to a unique singular point of $\tilde{f}_{0,0}$.

Suppose otherwise, and without loss of generality $0=s(g)<s(h)=c$, then each edge emanating from the root node correspond to the set $\{\zeta_1^{(0)}\}\times\mathbb{F}_p$ for a unique degenerate root $\zeta_1^{(0)}$ of the univariate polynomial $\tilde{g}(x_1)$. As $\tilde{g}$ has at most $\left\lfloor\frac{\deg\tilde{g}}{2}\right\rfloor\leq\left\lfloor\frac{d}{2}\right\rfloor\leq\binom{d}{2}$ degenerate roots, we are done.

Assertion (3): Suppose $f_{i-1,\mu}=g_{i-1}(x_1)+h_{i-1}(x_2)\in\mathbb{Z}[\mathbf{x}]$ with $s(g_{i-1})=s(h_{i-1})=0$. Then $\zeta^{(i-1)}$ is a singular point of $\tilde{f}_{i-1,\mu}$, and let

$$s:=s(f_{i-1,\mu},\zeta^{(i-1)})=\min_{0\leq i_1+i_2\leq k_{i,\zeta-1}}\left\{(i_1+i_2)+\mathrm{ord}_p\left(D^{i_1,i_2}f_{i-1,\mu}(\zeta^{(i-1)})\right)\right\}$$

So then for each pair of (ℓ_1,ℓ_2) with $\ell_1+\ell_2\geq s+1$, the coefficient of $x_1^{\ell_1}x_2^{\ell_2}$ in the perturbation $f_{i-1,\mu}(\zeta^{(i-1)}+p\mathbf{x})$ must be divisible by p^{s+1}. In other words, the coefficient of $x_1^{\ell_1}x_2^{\ell_2}$ in $f_{i,\zeta}(\mathbf{x})$ must be divisible by p. So $\deg\tilde{f}_{i,\zeta}\leq s$.

Now by Lemma 3.4, we know that the multiplicity of $\zeta^{(i-1)}$ on $\tilde{f}_{i-1,\mu}$: $m_{\zeta^{(i-1)}}(\tilde{f}_{i-1,\mu})\geq s(f_{i-1,\mu},\zeta^{(i-1)})$. Combining with 4.2, we have

$$\begin{aligned}\sum_{\substack{(f_{i,\zeta},k_{i,\zeta})\text{ a child}\\ \text{of }(f_{i-1,\mu},k_{i-1,\mu})}}\deg\tilde{f}_{i,\zeta}\left(\deg\tilde{f}_{i,\zeta}-1\right)&\leq\sum_{\substack{\zeta^{(i-1)}\text{ sing.}\\ \text{point on }\tilde{f}_{i-1,\mu}}}m_{\zeta^{(i-1)}}(\tilde{f}_{i-1,\mu})\left(m_{\zeta^{(i-1)}}(\tilde{f}_{i-1,\mu})-1\right)\\&\leq\deg\tilde{f}_{i-1,\mu}\left(\deg\tilde{f}_{i-1,\mu}-1\right).\end{aligned}$$

Suppose without loss of generality, $0=s(g_{i-1})<s(h_{i-1})=c$. Then by a similar argument $\deg\tilde{f}_{i,\zeta}\leq s(f_{i-1,\mu},\zeta^{(i-1)})=\min(s(\tilde{g},\zeta_1^{(i-1)}),c)\leq s(\tilde{g},\zeta_1^{(i-1)})$. By Lemma 4.11 we have that $\sum_{\substack{\zeta_1^{(i-1)}\text{ a deg.}\\ \text{root of }\tilde{g}_{i-1}}}s(\tilde{g}_{i-1},\zeta_1^{(i-1)})\leq\deg\tilde{g}_{i-1}$, so then $\sum_{\substack{(f_{i,\zeta},k_{i,\zeta})\text{ a child}\\ \text{of }(f_{i-1,\mu},k_{i-1,\mu})}}\deg\tilde{f}_{i,\zeta}\leq\deg\tilde{g}_{i-1}$. We are done, simply by observing that for $\deg\tilde{f}_{i,\zeta}\geq 2$ and any collections of $a_i>2$, we must have $\sum a_i(a_i-1)\leq\left(\sum a_i\right)\left(\sum a_i-1\right)$.

Assertion (4): This is immediate from Assertions (1) and (3). □

5.4. **The Proof of Lemma 4.7.** Any points over $\mathbb{F}_p$ on $\tilde{f}(\mathbf{x})$ is nonsingular if and only if $D^{1,0}(\tilde{f})=\tilde{g}'(x_1)\neq 0\mod p$, as $h(x_2)$ is identically 0 mod p. In other words, any nonsingular point on $\tilde{f}$ should be of the form $(\zeta_1^{(0)},y)$ where $\zeta_1^{(0)}$ is a non-degenerate root of $\tilde{g}$, and any choice of $y\in\{0,1,\ldots,p-1\}$. So the number of non-singular point on $\tilde{f}$ is: $n_p(f)=p\cdot n_p(g)$. Then the first summand in the equation is obvious by plugging into the first summand in Lemma 2.4.

Now suppose $\zeta_0:=\zeta_1^{(0)}$ is a degenerate root of the univariate polynomial $\tilde{g}$, and $\zeta^{(0)}=\{\zeta_0\}\times\mathbb{F}_p$. Write $\sigma=\zeta_0+p\tau$, where $\tau:=\zeta_1+\ldots+p^{k-2}\zeta_{k-1}\in\mathbb{Z}/p^{k-1}\mathbb{Z}$ via base-p expansion. Then by definition $f(\zeta_0+px_1,x_2)=p^{s(f,\zeta^{(0)})}f_{1,\zeta^{(0)}}(x_1,x_2)$, where $f_{1,\zeta^{(0)}}\in\mathbb{Z}[x_1,x_2]$ does not vanish identically mod p.

If $k \geq s(f, \zeta^{(0)})$, then $f(\sigma, y) = 0 \mod p^k$ regardless of choice of $\tau \in \mathbb{Z}/p^{k-1}\mathbb{Z}$ and $y \in \mathbb{Z}/p^k\mathbb{Z}$. So there are exactly $p^{k-1} \cdot p^k = p^{2k-1}$ many pairs of $(\sigma, y) \in (\mathbb{Z}/p^k\mathbb{Z})^2$ such that $\sigma = \zeta_0 \mod p$ and $f(\sigma, y) = 0 \mod p^k$.

If $s(f, \zeta^{(0)}) \leq k-1$, then $f(\sigma, y) = 0 \mod p^k$ if and only if

$$f_{1,\zeta^{(0)}}(\tau, y) = 0 \mod p^{k-s(f,\zeta^{(0)})}. \tag{4}$$

Let $s := s(f, \zeta^{(0)})$, then $\tau = \zeta_1 + p\zeta_2 + \ldots + p^{k-s-1}\zeta_{k-s} \mod p^{k-s}$ and $y := \sum_{i=0}^{k-1} p^i y_i = y_0 + \ldots + p^{k-s-1}y_{k-s-1} \mod p^{k-s}$. So the rest of the base-p digits, $\zeta_{k-s+1}, \ldots, \zeta_{k-1}$ and $y_{k-s}, \ldots, y_{k-1}$ respectively does not appear in Equality (4). The possible lifts ζ where the first coordinate mod p is ζ_0 is thus exactly $p^{s-1} \cdot p^s$ times the number of roots $(\tau, y) \in (\mathbb{Z}/p^{k-s}\mathbb{Z})^2$ of $f_{1,\zeta^{(0)}}$. □

5.5. **Exceptional Curves.** Let $f(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$ be a nonconstant polynomial, and let $s(f)$ denote the largest p-th power dividing all the coefficients of f.

Consider $f(\mathbf{x}) = g^d(\mathbf{x}) + p^{cd}h(\mathbf{x}) \in \mathbb{Z}[\mathbf{x}]$, with $d \geq 2$ and $c \geq 1$. Moreover, $f(\mathbf{x}) \equiv g^d(\mathbf{x}) \mod p$ and f is irreducible mod p.

For $k \leq cd$, $f(\mathbf{x}) = g^d(\mathbf{x}) \mod p^k$. Now suppose $\zeta^{(0)}$ is a smooth point on $(g \mod p)$. Then by Hensel's Lemma (Lemma 3.2), $\zeta^{(0)}$ lifts to $p^{\lceil \frac{k}{d} \rceil - 1}$ many roots of g mod $p^{\lceil \frac{k}{d} \rceil}$. Suppose σ is one of the lift, then $\sigma + p\tau$ for any $\tau \in (\mathbb{Z}/p^{k-\lceil \frac{k}{d} \rceil}\mathbb{Z})^2$ is a root of $(g^d \mod p^k)$. So each $\zeta^{(0)}$ lifts to $p^{\lceil \frac{k}{d} \rceil - 1} \cdot p^{2(k-\lceil \frac{k}{d} \rceil)} = p^{2k-\lceil \frac{k}{d} \rceil - 1}$ many roots of f mod p^k.

Now suppose $k > cd$, and let ζ be a root of $f \mod p^{cd}$ such that $\zeta^{(0)} \equiv \zeta \mod p$ is a smooth point on g. Consider the Taylor expansion of f at ζ:

$$\begin{aligned} f(\zeta + p^{cd}\mathbf{x}) &= [g(\zeta) + T(\mathbf{x})]^d + p^{cd}h(\zeta + p^{cd}\mathbf{x}) \\ &= \left[g(\zeta)^d + p^{cd}h(\zeta)\right] + \sum_{l=1}^{d} g(\zeta)^{d-l}T(\mathbf{x})^l + \sum_{i_1+i_2 \geq 1} D^{i_1,i_2}h(\zeta)p^{cd(i_1+i_2+1)}x_1^{i_1}x_2^{i_2} \end{aligned} \tag{5}$$

where $T(\mathbf{x}) := g(\zeta + p^{cd}\mathbf{x}) - g(\zeta) = \sum_{i_1+i_2 \geq 1} D^{i_1,i_2}g(\zeta)p^{cd(i_1+i_2)}x_1^{i_1}x_2^{i_2}$. As $\zeta^{(0)}$ is a smooth point on g, either $D^{1,0}g(\zeta)$ or $D^{0,1}g(\zeta)$ is not zero mod p. Then $s(T) = cd$, and each term in the second summand of Equality (5) has valuation $(d-l)\operatorname{ord}_p g(\zeta) + lcd$.

If $\zeta^{(0)}$ is also a point on $h \mod p$, then ζ continues to lift, and by Lemma 4.1, there are at most d^2 many such $\zeta^{(0)}$. However, there are cases when $h(\zeta) \neq 0 \mod p$, yet ζ continues to lift to p^k for $k > cd$.

This could only happen when $g(\zeta)^d + p^{cd}h(\zeta) \equiv 0 \mod p^{cd+1}$, and in which case $\operatorname{ord}_p g(\zeta) = c$. Now the second summand in Equality (5) must have order $(d-1)c + cd$, whereas the third summand has order $\geq 2cd$. So now $s(f, \zeta) = \min\left\{\operatorname{ord}_p\left(g(\zeta)^d + p^{cd}h(\zeta)\right), (d-1)c + cd\right\}$. If $s(f, \zeta) < (d-1)c + cd$ then $\tilde{f}_{cd,\zeta} = \frac{f(\zeta + p^{cd}\mathbf{x})}{p^{s(f,\zeta)}} \mod p$ is a nonzero constant, and thus ζ does not lift. Suppose otherwise. Then

$$\tilde{f}_{cd,\zeta} = \frac{g(\zeta)^d + p^{cd}h(\zeta)}{p^{s(f,\zeta)}} + \frac{dg(\zeta)^{d-1}}{p^{(d-1)c}}\left(D^{1,0}g(\zeta)x_1 + D^{0,1}g(\zeta)x_2\right) \mod p,$$

which defines a line! By Hensel's Lemma, we are done!

So the problem boils down to determining a criterion for when $\operatorname{ord}_p\left(f(\zeta)^d + p^{cd}h(\zeta)\right) \geq (d-1)c + cd$ and $h(\zeta) \neq 0 \mod p$ happens. Also, we need to compute $\operatorname{ord}_p\left(f(\zeta)^d + p^{cd}h(\zeta)\right)$ for every lift $\zeta \mod p^{cd}$ for each non-isolated singular points $\zeta^{(0)}$, and there are exactly p^{cd-1} many such ζ.

In summary, computing perturbations for each and every singular point of $\tilde{f}$ can be very expensive going into higher dimensions: the underlying singular locus might not be zero-dimensional, and thus imply the calclulation of a number of perturbations super-linear in p.

It turns out for *some* families of curves, non-isolated singular points partitioned into groups that each lift uniformly. We will pursue this improvement in future work.

Acknowledgements

We are grateful to Daqing Wan for helpful comments on curves and error correcting codes.

References

[AH01] Leonard M. Adleman and Ming-Deh Huang. Counting points on curves and abelian varieties over finite fields. *Journal of Symbolic Computation*, 32(3):171 – 189, 2001.

[BGMWo11] Edward Bierstone, Dima Grigoriev, Pierre Milman, and Jarosł aw Wł odarczyk. Effective Hironaka resolution and its complexity. *Asian J. Math.*, 15(2):193–228, 2011.

[BLQ13] Jèrèmy Berthomieu, Grègoire Lecerf, and Guillaume Quintin. Polynomial root finding over local rings and application to error correcting codes. *Appl. Algebra Eng. Commun. Comput.*, 24:413–443, 2013.

[BM20] Jennifer S. Balakrishnan and J.S̃teffen Müller. Computational tools for quadratic chabauty. preprint, Boston University, 2020. draft of lecture notes for 2020 Arizona Winter School on Nonabelian Chabauty.

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

[BW10] Maheshanand Bhaintwal and Siri Krishan Wasan. Generalized Reed-Muller codes over $\mathbb{Z}_{||}$. *Des. Codes Cryptogr.*, 54(2):149–166, 2010.

[CDV06] Wouter Castryck, Jan Denef, and Frederik Vercauteren. Computing zeta functions of nondegenerate curves. Technical report, International Mathematics Research Papers, vol. 2006, article ID 72017, 2006.

[CGRW18] Qi Cheng, Shuhong Gao, J. Maurice Rojas, and Daqing Wan. Counting roots for polynomials modulo prime powers. In *Proceedings of ANTS XIII (Algorithmic Number Theory Symposium, July 16–20, 2018, University of Wisconsin, Madison)*. Mathematical Sciences Publishers (Berkeley, California), 2018.

[CH15] Henry Cohn and Nadia Heninger. Ideal forms of Coppersmith's theorem and Guruswami-Sudan list decoding. *Advances in Mathematics of Communications*, 9(3):311–339, 2015.

[CL08] Antoine Chambert-Loir. Computer (rapidement) le nombre de solutions d'équations dans les corps finis. *Séminaire Bourbaki*, 2006/2007:39–90, 2008.

[Del74] Pierre Deligne. La conjecture de weil. i. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 43(1):273–307, Dec 1974.

[DMS19] Ashish Dwivedi, Rajat Mittal, and Nitin Saxena. Counting basic-irreducible factors mod p^k in deterministic poly-time and p-adic applications. *arXiv e-prints*, page arXiv:1902.07785, Feb 2019.

[DMS20] Ashish Dwivedi, Rajat Mittal, and Nitin Saxena. Computing Igusa's Local Zeta Function of Univariates in Determinstic Polynomial-Time. In S. K. Galbraith, editor, *Proceedings of ANTS 2020 (Algorithmic Number Theory Symposium)*. Mathematical Sciences Publishers (Berkeley, California), 2020.

[GCM91] Javier Gomez-Calderon and Gary L. Mullen. Galois rings and algebraic cryptography. *Acta Arith.*, 59(4):317–328, 1991.

[GG16] Steven D. Galbraith and Pierrick Gaudry. Recent progress on the elliptic curve discrete logarithm problem. *Des. Codes Cryptogr.*, 78(1):51–72, 2016.

[GS99] V. Guruswami and M. Sudan. Improved decoding of reed-solomon and algebraic-geometry codes. *IEEE Transactions on Information Theory*, 45(6):1757–1767, Sep. 1999.

[GSS00] Venkatesan Guruswami, Amit Sahai, and Madhu Sudan. "Soft-decision" decoding of Chinese remainder codes. In *41st Annual Symposium on Foundations of Computer Science (Redondo Beach, CA, 2000)*, pages 159–168. IEEE Comput. Soc. Press, Los Alamitos, CA, 2000.

[Har15] David Harvey. Computing zeta functions of arithmetic schemes. *Proceedings of the London Mathematical Society*, 111(6):1379–1401, 11 2015.

[HKC$^+$94] A. Roger Hammons, Jr., P. Vijay Kumar, A. R. Calderbank, N. J. A. Sloane, and Patrick Solé. The $\mathbf{Z}_4$-linearity of Kerdock, Preparata, Goethals, and related codes. *IEEE Trans. Inform. Theory*, 40(2):301–319, 1994.

[Igu07] Jun-Ichi Igusa. *An Introduction to the Theory of Local Zeta Functions*. AMS/IP Studies in Pure Maths Rep Series. American Mathematical Society, 2007.

[Ked01] Kiran S. Kedlaya. Counting points on hyperelliptic curves using Monsky-Washnitzer cohomology. *J. Ramanujan Math. Soc.*, 16(4):323–338, 2001.

[Ked06] Kiran S. Kedlaya. Quantum computation of zeta functions of curves. *Comput. Complexity*, 15(1):1–19, 2006.

[KLP12] Tali Kaufman, Shachar Lovett, and Ely Porat. Weight distribution and list-decoding size of Reed-Muller codes. *IEEE Trans. Inform. Theory*, 58(5):2689–2696, 2012.

[Kob87] Neal Koblitz. Elliptic curve cryptosystems. *Math. Comp.*, 48(177):203–209, 1987.

[KRRZ19] Leann Kopp, Natalie Randall, J. Maurice Rojas, and Yuyu Zhu. Randomized Polynomial-Time Root Counting in Prime Power Rings. *Mathematics of Computation*, in production, 2019.

[KU08] Kiran Kedlaya and Christopher Umans. Fast polynomial factorization and modular composition. In P. Bro Miltersen, R. Reischuk, G. Schnitger, and D. van Melkebeek, editors, *Computational Complexity of Discrete Problems*, number 08381 in Dagstuhl Seminar Proceedings, Dagstuhl, Germany, 2008. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany.

[LW08] Alan G. B. Lauder and Daqing Wan. Counting points on varieties over finite fields of small characteristic. In *Algorithmic number theory: lattices, number fields, curves and cryptography*, pages 579—612, Cambridge, 2008. Math. Sci. Res. Inst. Publ., 44, Univ. Press.

[Mil86] Victor S. Miller. Use of elliptic curves in cryptography. In *Advances in cryptology—CRYPTO '85 (Santa Barbara, Calif., 1985)*, volume 218 of *Lecture Notes in Comput. Sci.*, pages 417–426. Springer, Berlin, 1986.

[NZM91] I. Niven, H.S. Zuckerman, and H.L. Montgomery. *An Introduction to the Theory of Numbers*. Wiley, 1991.

[Pil90] J. Pila. Frobenius maps of abelian varieties and finding roots of unity in finite fields. *Mathematics of Computation*, 55(192):745–763, 1990.

[PR11] Adrien Poteaux and Marc Rybowicz. Complexity bounds for the rational Newton-Puiseux algorithm over finite fields. *Appl. Algebra Engrg. Comm. Comput.*, 22(3):187–217, 2011.

[Roj99] J. Maurice Rojas. Solving degenerate sparse polynomial systems faster. *Journal of Symbolic Computation*, 28(1):155 – 186, 1999.

[RZ20] J. Maurice Rojas and Yuyu Zhu. A complexity chasm for solving sparse polynomial equations over p-adic fields. *arXiv e-prints*, page arXiv:2003.00314, 2020.

[Sch85] René Schoof. Elliptic curves over finite fields and the computation of square roots mod p. *Mathematics of Computation*, 44(170):483–494, 1985.

[Sud97] Madhu Sudan. Decoding of Reed Solomon codes beyond the error-correction bound. *J. Complexity*, 13(1):180–193, 1997.

[vdG01] Gerard van der Geer. Curves over finite fields and codes. In *European Congress of Mathematics, Vol. II (Barcelona, 2000)*, volume 202 of *Progr. Math.*, pages 225–238. Birkhäuser, Basel, 2001.

[vzGG13] Joachim von zur Gathen and Jürgen Gerhard. *Modern Computer Algebra*. Cambridge University Press, 3rd edition, 2013.

[Wan08] Daqing Wan. Algorithmic theory of zeta functions over finite fields. In *Algorithmic number theory: lattices, number fields, curves and cryptography*, pages 551–578. Math. Sci. Res. Inst. Publ., 44, Univ. Press, Cambridge, 2008.

[Wei49] André Weil. Numbers of solutions of equations in finite fields. *Bull. Amer. Math. Soc.*, 55(5):497–508, May 1949.

[Zhu20] Yuyu Zhu. Trees, point counting beyond fields, and root separation. Ph.d. thesis, Texas A& University, 2020.