Point Counting on Curves over Prime Power Rings
Final Report for the Texas A&M Mathematics REU

Garrett Credi

1 Introduction

This project focused primarily on determining efficient methods to count the number of solutions a given polynomial $f \in \mathbb{Z}[x, y]$ possessed over prime power rings $\mathbb{Z}/p^k\mathbb{Z}$. By efficient methods, we seek to develop an algorithm that, given a polynomial $f \in \mathbb{Z}[x, y], k \in \mathbb{N}_+$ and p prime, will output the number of points on f over $\mathbb{Z}/p^k\mathbb{Z}$ in time polynomial in $\log(p), k$ and $deg(f)$. While there is already an existing explicit algorithm, it possess to major complexity-related drawbacks that this project sought to improve. While, unfortunately, no improvements were fully worked out, this document describes the work I was able to do, the paths that I explored, and the ideas that I had not fully worked out. Hopefully this will prove some use for future researchers examining this problem.

2 Background

The first part of this project was to understand the previous work done in understanding point counts. The first simplification to make is to work one prime p at a time, and to relate the different base rings $\mathbb{Z}/p^k\mathbb{Z}$ as k varies.

2.1 Truncation

The first connection to note is that the base rings themselves are related. Specifically, for $k' > k$, we have a truncation map $\pi_{p,k',k} : \mathbb{Z}/p^{k'}\mathbb{Z} \to \mathbb{Z}/p^k\mathbb{Z}$ that maps $[n] \mapsto [n \mod p^k]$.

Proposition 1. *The maps $\pi_{p,k',k}$ are well defined ring homomorphisms.*

Proof. Firstly, assume that $n, n' \in \mathbb{Z}$ both represent the same equivalence class in $\mathbb{Z}/p^{k'}\mathbb{Z}$, i.e. that $n = n' + rp^{k'}$. Then $\pi_{p,k',k}([n]) = [n \mod p^k] = [n' + rp^{k'} \mod p^k] = [n' \mod p^k]$ since $p^{k'} \equiv 0 \mod p^k$.

$\pi_{p,k',k}$ is also an obvious ring homomorphism since it is the projection map from $\mathbb{Z}/p^{k'}\mathbb{Z}$ to its quotient $\left(\mathbb{Z}/p^{k'}\mathbb{Z}\right)/(p^k)$. □

These $\pi_{p,k',k}$ do not just give maps on the base rings, but we will see that they also give rise to maps on the given solution sets we care about.

Definition 2. *Given $f \in \mathbb{Z}[x,y]$, p a prime, and $k \in \mathbb{N}_+$, let $\bar{f}$ be the image of f when projecting to $\mathbb{Z}/p^k\mathbb{Z}[x,y]$ and let $Z_{p,k}(f) = \left\{\zeta \in \left(\mathbb{Z}/p^k\mathbb{Z}\right)^2 \mid \bar{f}(\zeta) = [0]\right\}$. This is **the solution set of f over** $\mathbb{Z}/p^k\mathbb{Z}$.*

Thus the goal of this project can be reinterpreted as determining efficient algorithms to determine $|Z_{p,k}(f)|$.

Importantly, since each of the $\zeta \in Z_{p,k'}(f)$ live in $(\mathbb{Z}/p^{k'}\mathbb{Z})^2$, $\pi_{p,k',k}$ can give an element in $(\mathbb{Z}/p^k\mathbb{Z})^2$. Importantly, $\pi_{p,k',k}$ will also map solutions over $\mathbb{Z}/p^{k'}\mathbb{Z}$ to solutions over $\mathbb{Z}k$.

Theorem 3. *The maps $\pi_{p,k',k}$ give rise to natural maps $\pi_{p,k',k}(f) : Z_{p,k'}(f) \to Z_{p,k}(f)$.*

The first property to note is the compositionality of the $\pi_{p,k',k}$.

Proposition 4. *For $k'' > k' > k$ we have that $\pi_{p,k',k} \circ \pi_{p,k'',k'} = \pi_{p,k'',k}$.*

Proof. First note that $\pi_{p,k',k}([1]) = [1]$ since $1 + rp^{k'} = 1 + rp^{k'-k}p^k$. Then, since $[1]$ additively generates each of the $\mathbb{Z}/p^k\mathbb{Z}$'s, we have that

$$\begin{gathered}\pi_{p,k',k} \circ \pi_{p,k'',k'}([n]) = \pi_{p,k',k} \circ \pi_{p,k'',k'}([1] + [1] + ... + [1]) = \\ (\pi_{p,k',k} \circ \pi_{p,k'',k'})([1]) + (\pi_{p,k',k} \circ \pi_{p,k'',k'})([1]) + ... + (\pi_{p,k',k} \circ \pi_{p,k'',k'})([1]) = \\ [1] + [1] + ... + [1] = \pi_{p,k'',k}([1]) + \pi_{p,k'',k}([1]) + ... + \pi_{p,k'',k}([1]) = \\ \pi_{p,k'',k}([1] + [1] + ... + [1]) = \pi_{p,k'',k}([n])\end{gathered}$$

□

We can then restrict the generality needed to prove Theorem 3 by noticing that $\pi_{p,k',k} = \pi_{p,k+1,k} \circ ... \circ \pi_{p,k'-1,k}$.

Proposition 5. *If $\zeta \in (\mathbb{Z}/p^k\mathbb{Z})^2$ has $\bar{f}(\zeta) \neq [0]$ then for $\zeta' \in (\mathbb{Z}/p^{k+1}\mathbb{Z})^2$ such that $\pi_{p,k+1,k}(\zeta') = \zeta$, $\bar{f}(\zeta') \neq [0]$*

Proof. To begin, if $\pi_{p,k+1,k}(\zeta') = \zeta$, we must have that $\zeta' = \zeta + p^k\xi$ for some $\xi \in \mathbb{Z}/p^{k+1}\mathbb{Z}$.

Since $f(\zeta + p^k\xi) \equiv f(\zeta) + f_x(\zeta)p^k\xi_x + f_y(\zeta)p^k\xi_y \mod p^{k+1}$ (easily proven by manipulation) if the right hand side were to be zero, by choosing integer representatives for each term, we would have to have that $\nu_p(f(\zeta) + f_x(\zeta)p^k\xi_x + f_y(\zeta)p^k\xi_y) \geq k+1$ where ν_p denotes the p-adic valuation. However, since $f(\zeta) \neq 0$, any representative of ζ would have to have $\nu_p(f(\zeta)) < k$. Then, by the ultrametric inequality

$\nu_p(f(\zeta + p^k\xi)) \geq \min\{\nu_p(f(\zeta)), k + \nu_p(f_x(\zeta)\xi_x), k + \nu_p(f_y(\zeta)\xi_y)\}$ with equality if the min is attained uniquely. In this case, since $\nu_p(f(\zeta)) < k$ we have unique attainment, so $\nu_p(f(\zeta + p^k\xi)) = \nu_p(f(\zeta)) < k < k+1$.

Therefore $f(\zeta') \neq 0$. □

Proof of Theorem 3. Applying the contrapositive of Prop 5, we see that if $\zeta \in Z_{p,k+1}(f)$ then $\pi_{p,k+1,k}(\zeta) \in Z_{p,k}(f)$. Therefore for $k' > k$, $\pi_{p,k',k}$ restricts to a well defined function on $Z_{p,k'}(f)$ which we denote by $\pi_{p,k',k}(f)$. □

Remark 6. *This idea of truncation comes from the behavior of $\pi_{p,k',k}$ when writing a given $n \in \mathbb{Z}$ in base p. For example, if $n = d_0 + d_1 p + d_2 p^2 + ... + d_{k'} p^{k'}$, then $\pi_{p,k',k}([n]) = [d_0] + [d_1 p] + [d_2 p^2] + ... + [d_{k-1} p^{k-1}] + [d_k \cdot 0] + ... + [0] = [d_0] + [d_1 p] + [d_2 p^2] + ... + [d_{k-1} p^{k-1}]$. Thus $\pi_{p,k',k}$ really truncates a class down to its first k p-adic digits.*

2.2 Lifting

The importance of constructing these $\pi_{p,k',k}(f)$ is to relate $|Z_{p,k'}(f)|$ to $|Z_{p,k}(f)|$ in order to set up an inductive algorithm to point count.

The 'base case' for our purposes will be determining $|Z_{p,1}(f)|$ i.e. the number of points on f over $\mathbb{F}_p$. This is because the power of algebraic geometry can be readily applied in the case where the base ring is a field, and we will see the application of one such method later in the paper. But, for now, assume we have an efficient algorithm to determine $|Z_{p,1}(f)|$.

However, not all $\mathbb{F}_p$ points ζ are made equal.

Definition 7. *A point $\zeta \in Z_{p,1}(f)$ is singular if $\dfrac{\partial f}{\partial x}(\zeta) \equiv \dfrac{\partial f}{\partial y}(\zeta) \equiv 0 \mod p$. A point is non-singular if it is not singular.*

The importance of singularity can be seen as an extension of the reasoning present in the proof of Proposition 5. Since the ability for a point $\zeta \in Z_{p,k}(f)$ to have a lift to a point $\zeta' \in Z_{p,k+1}(f)$ depended on whether or not

$$f(\zeta) + f_x(\zeta) p^k \xi_x + f_y(\zeta) p^k \xi_y \equiv 0 \mod p^{k+1} \tag{1}$$

While this equation may seem far removed from the finite field case, we can perform a few reductions to see that this equation is truly over $\mathbb{F}_p$. First, note that only the behavior of $\xi \mod p$ affects Eqn. 1 since any term of order p or higher in ξ (i.e. any digits d_i for $i \geq 1$) would give a term of order p^{k+1} or higher, which reduces to $0 \mod p^{k+1}$. Thus we can instead force "$\xi \in \mathbb{F}_p^2$". This is a slight abuse of notation, but it is a convenient one for what is to come. Secondly, if we assume that $\zeta \in Z_{p,k}(f)$ then $f(\zeta) = rp^k$. Thus, Eqn. 1 simplifies to

$$p^k (r + f_x(\zeta)\xi_x + f_y(\zeta)\xi_y) \equiv 0 \mod p^{k+1} \tag{2}$$

$$r + f_x(\zeta)\xi_x + f_y(\zeta)\xi_y \equiv 0 \mod p \tag{3}$$

Which clearly reveals the dependency on whether or not ζ is a singular or non-singular point on f over $\mathbb{F}_p$. Importantly,

Theorem 8. *If ζ is a nonsingular $\mathbb{F}_p$ point of f, then there are exactly p^{k-1} points $\zeta' \in Z_{p,k}(f)$ such that $\pi_{p,k,1}(\zeta') = \zeta$.*

Proof. As was done with the proof for 3, we will restrict to demonstrating that a point $\zeta \in Z_{p,k}(f)$ which is non-singular over $\mathbb{F}_p$ lifts to exactly p points in $Z_{p,k}(f)$.

As was noted above, to determine whether or not $\zeta + p^k \xi$ is a solution $\mod p^{k+1}$ depends on whether or not Eqn. 3 is satisfied. Since ζ is assumed to be a nonsingular point over $\mathbb{F}_p$ we can

assume that either $f_x(\zeta) \neq 0 \mod p$ or $f_y(\zeta) neq 0 \mod p$. Without loss of generality, assume that $f_x(\zeta) \neq 0$, as what follows can just be substituted in the case that $f_x(\zeta) = 0$ and $f_y(\zeta) \neq 0$. Then, since $\mathbb{F}_p$ is a field, $f_x(\zeta)$ possesses a multiplicative inverse, so we can re-arrange Eqn. 3 to

$$\xi_x = f_x(\zeta)^{-1}(f_y(\zeta)\xi_y - r) \mod p \tag{4}$$

Thus any choice of ξ_y determines a ξ_x. Since there are p choices for ξ_y, we get p unique lifts.

Now, since singularity is not affected by lifting (if $\zeta \equiv \zeta' \mod p^k$ then their mod p reductions are also equal $\pi_{p,k',1}(\zeta') = \pi_{p,k,1} \circ \pi_{p,k',k}(\zeta') = \pi_{p,k,1}(\zeta)$), this allows us to inductively lift from $Z_{p,1}(f)$ to $Z_{p,2}(f)$ to $Z_{p,3}(f)$ and so on, up to $Z_{p,k}(f)$. Since $k-1$ lifts occurred, and each lifting stage introduced p points for every point in the base, we get p^{k-1} lifts in total. □

The case for singular $\mathbb{F}_p$ points is more subtle, but is still closely related to the case for non-singular points. The important observation is that the expansion needed for the non-singular case (Eqn. 1) is the first order Taylor-expansion of f at ζ. In order to work in the singular case, we will need to use a higher order expansion.

Proposition 9. *For $\zeta \in Z_{p,1}(f)$ and $\xi \in Z_{p,k-1}(f)$,*

$$f(\zeta + p\xi) \equiv \sum_{n=0}^{\infty} p^n \sum_{\substack{i+j=n \\ i,j\geq 0}} \frac{1}{i!j!}\frac{\partial^n f}{\partial x^i \partial y^j}(\zeta)\xi_x^i \xi_y^j \mod p^k$$

Proof. As a function $f : \mathbb{R} \to \mathbb{R}$ this is an equality on the nose via Taylor's Theorem, so it remains to show that each of the coefficients are well defined over $\mathbb{Z}/p^k\mathbb{Z}$. The only possible offenders are the coefficients of the form $\frac{p^n}{i!j!}$, and we will have to demonstrate that $n \geq \nu_p(i!) + \nu_p(j!)$.

We now try to approximate $\nu_p(i!) = \sum_{m=1}^{i} \nu_p(m)$. Note that $\nu_p(i!) = \sum_{m=1}^{\infty} m\cdot$(number of multiples of p^m less than i) $\leq \sum_{m=1}^{\infty} \frac{mi}{p^m} = \frac{pi}{(p-1)^2}$ where the first equality is just reintepreting $\nu_p(i!) = \sum_{m=1}^{i} \nu_p(m)$ and the last equality comes from using the generating function $f(x) = \sum_{m=1}^{\infty} p^{xm}$. Therefore $n - \nu_p(i!) - \nu_p(j!) = i - \nu_p(i!) + j - \nu_p(j!) \leq i - \frac{p}{(p-1)^2}i + j - \frac{p}{(p-1)^2}j = (1 - \frac{p}{(p-1)^2})(i+j)$ which is greater than or equal to zero since $0 \leq \frac{p}{(p-1)^2} \leq 1$. Therefore our expansion is valid, as all coefficients are well-defined over $\mathbb{Z}/p^k\mathbb{Z}$! □

Now, just as we used the fact that $\zeta \in Z_{p,k}(f)$ to get the factor of p^k out of Eqn. 1, we will try and factor out a high power of p as well.

Definition 10. *Let $s(f,\zeta) = \min_{i,j\geq 0}\left\{i + j + \frac{1}{i!j!}\frac{\partial^n f}{\partial x^i \partial y^j}(\zeta)\right\}$, which is just the minimal power of p that divides all the coefficients of $f(\zeta + p\xi)$. Furthermore, let $f_{\zeta,k}(\xi) = p^{-s(f,\zeta)}f(\zeta + p\xi)$.*

With that notation in hand, we arrive at the following fact.

Proposition 11. *A point $\zeta' = \zeta + p\xi \in (\mathbb{Z}/p^k\mathbb{Z})^2$ is a lift of $\zeta \in Z_{p,1}(f) \mod p^k$ if and only if $f_{\zeta,k}(\xi) \equiv 0 \mod p^{k-s(f,\zeta)}$*

Proof.

$$f(\zeta + p\xi) = p^{s(f,\zeta)} f_{\zeta,k}(\xi) \mod p^k$$

Which is zero if and only if $f_{\zeta,k}(\xi)$ is zero $\mod p^{k-s(f,\zeta)}$. □

Remark 12. *Eqn. 1 is extremely similar to this setup, where $s(f,\zeta) = k$ and $f_{\zeta,k}(\xi) = r + f_x(\zeta)\xi_x + f_y(\zeta)\xi_y$.*

Remark 13. *Additionally, if $s(f,\zeta) \geq k$, then $f_{\zeta,k}(\xi) = 0$. This means that we would have $p^{2(k-1)}$ points in $\mathbb{Z}/p^k\mathbb{Z}$ that lift from $\zeta \in Z_{p,1}(f)$ since the first x and y coordinates of the lift are fixed, while the remaining $k-1$ coordinates on both axes are free. Thus for such a ζ, we would get $p^{2(k-1)}$ unique points in $Z_{p,k}(f)$.*

If $s(f,\zeta) < k$ then every root of $f_{\zeta,k}(\xi)$ would lift to $p^{2(s(f,\zeta)-1)}$ points in $Z_{p,k}(f)$ since ζ determines the first x and y coordinate, ξ determines the next $k - s(f,\zeta)$ roots and so there are $2k - 2(k - s(f,\zeta) + 1) = 2(s(f,\zeta) - 1)$ coordinates that are free.

If we have a singular root ζ of f over $\mathbb{F}_p$, we can use the above proposition to reduce our exponent, and move to a different equation in order to simplify counting. We then arrive at a recursive formula to determine the number of points on a curve over $\mathbb{Z}/p^k\mathbb{Z}$.

Theorem 14. *Let $S_p(f)$ be the set of singular points of f over $\mathbb{F}_p$, and $n_p(f)$ be the number of non-singular points of f over $\mathbb{F}_p$. Then*

$$|Z_{p,k}(f)| = n_p(f)p^{k-1} + \sum_{\substack{\zeta \in S_p(f) \\ s(f,\zeta) \geq k}} p^{2(k-1)} + \sum_{\substack{\zeta \in S_p(f) \\ s(f,\zeta) < k}} p^{2(s(f,\zeta)-1)} |Z_{p,k-s(f,\zeta)}(f_{\zeta,k})|$$

Note that this formula provides a recursive algorithm to compute the point count for any f, p, k. *However*, the computational complexity of the corresponding algorithm is not on the order that we would like it to perform on, and the goals of this project were focused on improving such bounds. Unfortunately, we were unable to achieve such improvements, but we will describe avenues explored during this REU, their pitfalls, and their viability for future exploration now.

3 Current Work

3.1 Determining $|Z_{p,1}(f)|$

The first quantity necessary in the computation of 14 is the number of non-singular points on f over $\mathbb{F}_p$. However, a more tractable quantity to compute is the number of points on f, singular or non-singular, over $\mathbb{F}_p$. Knowing this quantity, along with $S_p(f)$ allows us to compute $n_p(f)$ as $|Z_{p,1}(f)| - |S_p(f)|$.

The main resource for computing this quantity came from (Harvey) with the following two results.

Theorem 15 ((Trace formula) (Harvey) Theorem 3.1)**.** *Let $\bar{F} \in \mathbb{F}_q[x]_d$ and let X be the hypersurface in $\mathbb{T}^n_{\mathbb{F}_q}$ cut out by $\bar{F}$. Let r, λ and τ be positive integers satisfying*

$$\tau \geq \frac{\lambda}{(p-1)ar}. \tag{5}$$

Let $F \in \mathbb{Z}_q[x]_d$ be any lift of $\bar{F}$. Then

$$|X(\mathbb{F}_{q^r})| = (q^r - 1)^n \sum_{s=0}^{\lambda+\tau-1} \alpha_s \operatorname{tr}(A^{ar}_{F^s}) \pmod{p^\lambda},$$

where

$$\alpha_s = (-1)^s \sum_{t=0}^{\tau-1} \binom{-\lambda}{t} \binom{\lambda}{s-t} \in \mathbb{Z},$$

and where A_{F^s} is regarded as a linear operator on $\mathbb{Z}_q[x]_{ds}$.

and

Lemma 16 ((Harvey)Lemma 3.2)**.** *Let $F \in \mathbb{Z}_q[x]_d$. The matrix of $A^a_{F^s}$ on $\mathbb{Z}_q[x]_{ds}$, with respect to the basis B_{ds}, is given by*

$$\phi^{a-1}(M_s) \cdots \phi(M_s) M_s,$$

where M_s is the square matrix defined by

$$(M_s)_{v,u} = (F^{(p-1)s})_{pv-u}$$

for $u, v \in B_{ds}$, and where ϕ acts componentwise on matrices.

In the case of point counting on a curve X over $\mathbb{F}_p$, we note that, since Harvey's theorem works over $\mathbb{F}_q$ where $q = p^a$ and point counts over $\mathbb{F}_{q^r}$ we can set $a = r = 1$. Similarly, as the maximum number of points on X over $\mathbb{F}_p$ is p^2, we can compute $|X(\mathbb{F}_p)| \mod p^2$, and thus we can work with $\lambda = 2$. Thus we need to choose a τ such that $\tau \geq \frac{\lambda}{(p-1)ar} = \frac{2}{p-1}$. To keep τ as small as possible, we then set $\tau = \begin{cases} 2, & p = 2 \\ 1, & p \neq 2 \end{cases}$. Finally, with Harvey's terminology, working with curves amounts to working with homogenous polynomials over $\mathbb{Z}[x_0, x_1, x_2]$ so $n = 2$.

Thus 15 simplifies to computing $|X(\mathbb{F}_p)| = (p-1)^2 \sum_{s=0}^{\tau+1} \alpha_s \operatorname{tr}(A_{F^s}) \mod p^2$. As this now depends on computing $\operatorname{tr}(A_{F^s})$, simplifying 16 will allow us to efficiently compute the point count. Moreover, as s and τ can only take a finite set of values, we can compute α_s in constant time via a simple lookup table.

$$\alpha_s = \begin{cases} 1, & s = 0 \wedge \tau = 1 \\ -2, & s = 1 \wedge \tau = 1 \\ 1, & s = 2 \wedge \tau = 1 \\ 1, & s = 0 \wedge \tau = 2 \\ 0, & s = 1 \wedge \tau = 2 \\ -3, & s = 2 \wedge \tau = 2 \\ 2, & s = 3 \wedge \tau = 2 \end{cases}$$

Since $a = 1$, the product representation $A_{F^s} = \phi^{a-1}(M_s)...\phi(M_s)M_s$ collapses to just $A_{F^s} = M_s$. Importantly, since we don't have to worry the matrix multiplications causing non-diagonal entries of M_s to interact with the overall trace, we can restrict ourselves to computing *only* the diagonal entries of M_s.

Since, by definition, $(M_s)_{v,u} = (F^{(p-1)s})_{pv-u}$ (where the parenthetical notation on the LHS denotes matrix entries and the notation on the RHS denoting the monomial coefficient of $\mathbf{x}^{pv-u}$ in $F^{(p-1)s}$, the diagonal entries of M_s are $(M_s)_{u,u} = (F^{(p-1)s})_{pu-u} = (F^{(p-1)s})_{(p-1)u}$. Translating that from the parenthetical notation into the monomial notation: the diagonal entries of M_s are exactly the coefficients of monomials in $F^{s(p-1)}$ that are of the form $\mathbf{x}^{(p-1)}$ with x a degree $s \cdot deg(F)$ homogenous monomial.

Tying it all together, $\operatorname{tr}(A_{F^s})$ is then the sum of the coefficients of $(p-1)$ powers of $s \cdot deg(F)$-degree monomials present in the expansion of $F^{(p-1)s}$. Finding a way to efficiently compute this will provide an efficient method to compute $|X(\mathbb{F}_p)|$ (hopefully faster than $\sqrt{p}\log^{1+\epsilon}(p)$-time that Harvey provides).

(Note: for some reason Harvey's definition of $|X(\mathbb{F}_p)|$ is the number of points of X over $\mathbb{F}_p$ *none of whose coordinates are zero*. To determine the true amount, we would first compute $|X(\mathbb{F}_p)|$ by (the modified version of) Harvey's algorithm and then determine the number of roots of $F(x, 0, 1)$ (denoted N_x) and $F(0, y, 1)$ (denoted N_y). Then the true number of roots would be $|X(\mathbb{F}_p)| + N_x + N_y - \begin{cases} 1, & F(0,0,1) = 0 \\ 0, & \text{otherwise} \end{cases}$ where the subtraction is to prevent double counting of the root $(0, 0)$ if it is present.)

3.1.1 Example

For example, let's take $f(x, y) = x^2 + y^2$, so $F(x, y, z) = x^2 + y^2$ and let's work over general p. Then $F^{(p-1)s} = \sum_{i=0}^{(p-1)s} \binom{(p-1)s}{i} x^{2i} y^{2((p-1)s-i)}$. If $p = 2$ then $F^{(p-1)s} = F^s$ so the sum of $(p-1)$ power monomial coefficients is just the sum of the coefficients thus $\operatorname{tr}(A_{F^s}) = F(1, 1, 1)^s = 2^s$. So $|X(\mathbb{F}_2)| = \sum_{s=0}^{2+1} \alpha_s \operatorname{tr}(A_{F^s}) = 1 + 0 \cdot 2^1 - 3 \cdot 2^2 + 2 \cdot 2^3 = 1 \mod 2^2$. This gives the total number of nonzero roots ((1,1) being the only one), and finding the other roots will give us 1 additional root (0,0) thus giving us 2 roots over $\mathbb{F}_2$.

For higher p, we need only to be concerned with $s = 0, 1, 2$. As above, if $s = 0$, $\operatorname{tr}(A_{F^s}) = 1$. For $s = 1$, $F^{(p-1)} = \sum_{i=0}^{p-1} \binom{p-1}{i} x^{2i} y^{2(p-1-i)}$. For this to be a $p-1$ power, we need that both

$$\begin{cases} 2i \equiv 0 \mod p-1 \\ 2(p-1-i) \equiv -2i \equiv 0 \mod p-1 \end{cases}$$

which are the same equation, so we need to characterize the i such that $2i \equiv 0 \mod p-1$. It is easy to see that the only i that satisfy this equation (that are in the range $0 \leq i \leq p-1$) are $i = 0, \frac{p-1}{2}, p-1$. Therefore $\operatorname{tr}(A_{F^1}) = \binom{p-1}{0} + \binom{p-1}{\frac{p-1}{2}} + \binom{p-1}{p-1} = 2 + \binom{p-1}{\frac{p-1}{2}}$.

For $s = 2$, $F^{(p-1)\cdot 2} = \sum_{i=0}^{2(p-1)} \binom{2(p-1)}{i} x^{2i} y^{2((2(p-1)-i))}$, which, by a similar argument as above,

implies that the trace is $\text{tr}(A_{F^2}) = \sum_{i=0}^{4} \binom{2(p-1)}{\frac{i(p-1)}{2}}$ therefore

$$|X(\mathbb{F}_p)| = (p-1)^2(1 \cdot 1 - 2 \cdot (2 + \binom{p-1}{\frac{p-1}{2}}) + 1 \cdot \sum_{i=0}^{4} \binom{2(p-1)}{\frac{i(p-1)}{2}}) + 1 \mod p^2$$

3.1.2 Approaches to Computation

While the above computation may suggest that computing the sum is an easy calculation, in general it is much more subtle. For example, take $F(x,y,z) = zy^2 - x^3 - zx^2$ and $p = 3$. While the monomials present in F are zy^2, x^3, and zx^2, $F^{p-1} = x^6 + 2x^5z + x^4z^2 - 2x^3y^2z - 2x^2y^2z^2 + y^4z^2 = ((x^3)^2 + (zx^2)^2 - (xyz)^2 + (zy^2)^2) + 2x^5z - 2x^3y^2z$ which has $p-1$ powers of zy^2, x^3, and zx^2 *as well as* xyz. Occurrences like these require more subtle methods to determine these coefficient sums.

The most fruitful method we have found involves a reinterpretation as a linear system of equations. Note that if $F(x,y,z) = \sum_{i=1}^{n} c_i \mathbf{x}^{e_i}$ then

$$F^{p-1} = \sum_{\substack{m_1+m_2+\ldots+m_n=p-1 \\ m_i \geq 0}} c_i' \mathbf{x}^{m_1e_1+m_2e_2+\ldots+m_ne_n}$$

Thus, if we want to determine the coefficient of $\mathbf{m}^(p-1)$ for some $deg(F)$ monomial $\mathbf{m}$, we could first determine which m_i's give rise to it. More formally, if $e_i = \langle \alpha_i, \beta_i, \gamma_i \rangle$ (so that $\mathbf{x}^{e_i} = x^{\alpha_i}y^{\beta_i}z^{\gamma_i}$) and we're attempting to determine the coefficient of $(x^ay^bz^c)^{p-1}$ then we would like to determine the m_i's so that

$$\sum_i m_i \begin{bmatrix} \alpha_i \\ \beta_i \\ \gamma_i \end{bmatrix} = (p-1) \begin{bmatrix} a \\ b \\ c \end{bmatrix} \tag{6}$$

$$\sum_i m_i = (p-1) \tag{7}$$

While it would seem natural to want to solve for the m_i's over $\mathbb{N}$, solving such linear system seems to be NP-complete as it is an instance of the Knapsack Problem. Instead, we could either solve the system over $\mathbb{Z}$, and carefully add m_i's that live in the corresponding nullspace to get a solution vector $\mathbf{m} = \langle m_i \rangle \in \mathbb{N}^n$. For possible information about this, check this MathOverflow question.

Additionally, we could also note that since the m_i's are positive and sum to $p-1$, $0 \leq m_i \leq p-1$ so we could solve the system over $\mathbb{F}_p$ as well. However in this case, we would want to be able to solve Equation 6 'over $\mathbb{F}_p$' while solving Equation 7 'over $\mathbb{Z}$'.

3.2 Dealing with non-squarefree-curves

Another important element of 14 is being able to iterate over $S_p(f)$ for the recursive steps. This would then lead to a complexity factor of $|S_p(f)|$ for the corresponding algorithm, and, in bad cases, would lead to a factor of p in the complexity. Since we desire to have $\log(p)$ complexity, this case is particularly troubling.

Before we discuss the bad cases, let us first examine the good cases.

Definition 17. *We say that a polynomial $f \in R[x,y]$ over a U.F.D. R is square-free if it factorizes as $f = \prod_i f_i^{e_i}$ where each $e_i = 1$ (with f_i irreducible and pairwise relatively prime).*

For such an f, it is easy to demonstrate that $|S_p(f)|$ does not add any factors of p to the complexity. This is due to the following result:

Proposition 18. *Let $f \in R[x,y]$, and f_x, f_y the formal derivatives with respect to x and y. Then f is not square-free if and only if $\gcd(f, f_x, f_y) \notin R$.*

Proof. Firstly, assume that f is not square free so $f = \prod_i f_i^{e_i}$. Assume without loss of generality that $e_0 > 1$. Then $f_x = (\frac{d}{dx} f_0)(e_0 f_0^{e_0-1})(\prod_{i \neq 0} f_i^{e_i}) + f_0^{e_0} \cdot \frac{d}{dx}(\prod_{i \neq 0} f_i^{e_i})$. Since $e_0 > 1, e_0 - 1 > 0$ so $f_0 | f_x$ and $f_0 | f$ so $f_0 | \gcd(f, f_x)$. A similar line of reasoning demonstrates that $f_0 | \gcd(f, f_y)$. Therefore $f_0 | \gcd(f, f_x, f_y)$.

Now assume that $\gcd(f, f_x, f_y) = g$ where $g \notin R$. Without loss of generality, assume that g is irreducible (otherwise we can choose an irreducible factor of g and proceed). Then, since $g | f$, we must have $g = f_i$ for some i. Again, after possible relabeling, assume that $g = f_0$. Then, by using the same expansion of f_x as above, we see that $e_0 - 1 > 0$ so $e_0 > 1$ and therefore f is not square-free. □

Theorem 19. *For f square-free, $|S_{p,1}(f)| = O(deg(f)^2)$.*

Proof. By Bezout's Theorem (see (Cox et al.)Thm, 7 p. 456, for C and D projective curves without common components, $|C \cap D| \leq deg(C)deg(D)$. By 18, for f square free at least one of $\gcd(f, f_x), \gcd(f, f_y)$ is a scalar, assume without loss of generality that $\gcd(f, f_x) \in R$, so the curves they represent have no common components. Since $|S_{p,1}(f)| = |\{\zeta \in \mathbb{F}_p^2 | f(\zeta) = f_x(\zeta) = f_y(\zeta) = 0\}|$, $|S_{p,1}(f) \leq |\{\zeta \in \mathbb{F}_p^2 | f(\zeta) = f_x(\zeta) = 0\}| \leq deg(f)deg(f_x) \leq deg(f)^2$. □

Importantly, for f not square-free, the above result fails *terribly*. For example, if g defines a genus 0 curve, then there are immediately $p-1$ points on g over $\mathbb{F}_p$ by a simple application of the Hasse-Weil bound. Then, for $f = g^2$, every point on g is a singular point on f and so $|S_{p,1}(f)| = O(p)$ which adds terrible complexity onto the runtime of the point counting algorithm.

3.2.1 Current Work

Initially, one may think that points on g may lead to similar behavior, so iterating over the $O(p)$ many points on it could be replaced by a single 'generic reduction'. I.e. given a curve $f \equiv \prod_i f_i^{e_i} \mod p$ with, for example, $e_0 > 1$, and two points ζ, ζ' on f_0 one might expect that $s(f, \zeta) = s(f, \zeta')$ and that $|Z_{p,k-s(f,\zeta)}(f_{\zeta,k})| = |Z_{p,k-s(f,\zeta')}(f_{\zeta',k})$. However, one can construct cases where two points on f_0 with differing p-adic valuations $\nu_p(f_0(\zeta)), \nu_p(f_0(\zeta'))$ lead to different s values and point counts. Even worse, we can find two points such that their behavior on f_0 is exactly the same, but that give different s values and point counts upon reduction.

This lead to the notion of a (g,p,k) valuative decomposition

Definition 20. *For $f \in \mathbb{Z}[x,y]$ such that $g^2 | f \mod p$, a (g,p,k) valuative decomposition is an identity of the form $f \equiv \sum_{i=0}^{k-1} p^i g^{e_i} h_i \mod p^k$.*

This decomposition should allow for a continuation of the ideas above, i.e. partitioning $S_{p,1}(f)$ into groups that 'share the same behavior' with respect to this decomposition, and continuing from there. However, such a process may be quite complicated, but we nonetheless wish to offer it as a method in the case that something fruitful may come of it.

We can also use such a decomposition to examine $|Z_{p,k}(f) \cap Z_{p,k}(g)|$ which, by inclusion-exclusion methods may allow us to compute $|Z_{p,k}(f)|$. In the cases where $e_i > 0$, it is quite obvious that $|Z_{p,k}(f) \cap Z_{p,k}(g)| = |Z_{p,k}(g)|$ but something more interesting happens in the case where there is at least one $e_i = 0$.

To begin, note that $e_0 > 0$ since we stipulate that $g^2 | f \mod p$, so $e_0 \geq 2$. Then we rearrange the decomposition in the following manner $f \equiv \sum_{i,e_i=0} p^i h_i + \sum_{i,e_i>0} p^i g^{e_i} h_i$. If $\zeta \in Z_{p,k}(g)$ then $f(\zeta) \equiv \sum_{i,e_i=0} p^i h_i \mod p^k$. However, since $\min\{i | e_i = 0\} = \iota > 0$, $\sum_{i,e_i=0} p^i h_i = p^\iota \sum_{i,e_i=0} p^{i-\iota} h_i$. Therefore, forcing $f(\zeta) \equiv p^\iota \sum_{i,e_i=0} p^{i-\iota} h_i \equiv 0 \mod p^k$ is equivalent to forcing $\sum_{i,e_i=0} p^{i-\iota} h_i \equiv 0 \mod p^{k-\iota}$.

This allows us to reduce the system

$$\begin{cases} f \equiv 0 \mod p^k \\ g \equiv 0 \mod p^k \end{cases}$$

to the system

$$\begin{cases} g \equiv 0 \mod p^k \\ \sum_{i,e_i=0} p^{i-\iota} h_i \equiv 0 \mod p^{k-\iota} \end{cases}$$

which may offer more tractability in solving it.

Possible methods towards approaching such a system might be to solve both equations $\mod p^{k-\iota}$ using a modified version of Hensel lifting, keeping track of how many solutions are singular and nonsingular points of g, and then lifting the nonsigular $\mod p^{k-\iota}$ points to unique points $\mod p^k$ while doing *something* for the singular points. Hopefully this leads to something fruitful.

4 References

Works Cited

Cox, David A., et al. *Ideals, Varieties, and Algorithms*. Springer International Publishing, 2015, doi:10.1007/978-3-319-16721-3.

Harvey, David. Computing zeta functions of arithmetic schemes. 2014. doi:10.48550/ARXIV.1402.3439.