Distribution Agreement

Signature:

______________________________ ______________

Noam Kantor Date

Rank-Favorable Bounds for Rational Points on Superelliptic Curves

By

Noam Kantor
Master of Science

Mathematics

Dr. David Zureick-Brown
Advisor

Dr. Suresh Venapally
Committee Member

Dr. Patrick Allitt
Committee Member

Accepted:

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

Date

Rank-Favorable Bounds for Rational Points on Superelliptic Curves

By

Noam Kantor

Advisor: Dr. David Zureick-Brown, PhD

An abstract of
A thesis submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Master of Science
in Mathematics
2017

Abstract

Rank-Favorable Bounds for Rational Points on Superelliptic Curves
By Noam Kantor

Let C be a curve of genus at least two, and let r be the rank of the rational points on its Jacobian. Under mild hypotheses on r, recent results by Katz, Rabinoff, Zureick-Brown, and Stoll bound the number of rational points on C by a constant that depends only on its genus. Yet one expects an even stronger bound that depends favorably on r: when r is small, there should be fewer points on C. In a 2013 paper, Stoll established such a bound for hyperelliptic curves using Chabauty's method. In the present work we extend Stoll's results to superelliptic curves. We also discuss a possible strategy for proving a rank-favorable bound for arbitrary curves based on the metrized complexes of Amini and Baker. Our results have stark implications for bounding the number of rational points on a curve, since r is expected to be small for "most" curves.

Rank-Favorable Bounds for Rational Points on Superelliptic Curves

By

Noam Kantor

Advisor: Dr. David Zureick-Brown, PhD

A thesis submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Master of Science
in Mathematics
2017

Acknowledgements

This thesis would never have been possible without the help of my advisor David Zureick-Brown and my graduate student advisor Jackson Morrow. They listened to my ideas – good and bad – and helped streamline the arguments that did work out. I am grateful also to Andrew Obus, who directed me to Raynaud's work on quotients of semistable curves. Last but not least, I am indebted to the entire Emory Mathematics Department for supporting me throughout my mathematical journey.

Contents

1 Background and Main Theorems **1**

1.1 Notation 5

2 Chabauty's Method and Stoll's Improvement **6**

2.1 The Technique 6

3 Automorphisms of the Annulus **9**

4 Annuli on Superelliptic Curves **10**

5 Bounding Zeros of Differentials **14**

6 The Final Count(down) **17**

6.1 p-Adic Rolle's Theorem 17

6.2 Uniform Bounds 18

7 Discussion **22**

7.1 Other Possible Attacks on Superelliptic Curves 22

8 Metrized Complexes **23**

9 A New Chabauty Divisor **27**

10 Bounding $\deg D_{\mathrm{chab}}$ **29**

10.1 Katz and Zureick-Brown's Analysis 29

10.2 $\deg D_{\mathrm{chab}}$ on the Metrized Complex 30

10.3 Applying the Lemma to Chabauty 33

11 Evidence for the Main Assumption **34**

1 Background and Main Theorems

Diophantine finiteness, in its most general sense, is the philosophy that the set of rational points on a variety of general type should be "small," for example not Zariski dense. There is a web of interconnected conjectures that make this philosophy precise, such as the Bombieri-Lang conjecture (see [7], p. 479.) In the dimension one case, Bombieri-Lang is just Faltings' Theorem, which says that the number of rational points on a curve of genus greater than one is finite.

But it is computationally and theoretically important to have control of rational points beyond the basic finiteness supplied by Faltings. In fact, one conjectures the following:

Conjecture 1 ([4]). *Let C be a curve of genus $g > 1$ defined over a number field K. Then there exists a constant $B(g, K)$ such that $\#C(K) < B(g, K)$.*

Significant progress has been made on this conjecture when the rank of the Jacobian of C is somewhat smaller than g using a technique called Chabauty's method. The classical application of the method constructs a p-adic analytic differential ω and an associated p-adic analytic function f_ω on $\mathrm{Jac}_\mathrm{C}(\mathbb{C}_p)$ such that the zero set of f_ω contains $C(\mathbb{Q})$. Chabauty's original paper reasons using basic facts about analytic functions to show that f_ω only has finitely many zeros. We now know that f_ω is really a p-adic integral of ω, and a robust theory of p-adic integration is the key to modern applications of Chabauty's ideas.

Indeed, a number of improvements have been made to Chabauty's method which make it both stronger and more widely applicable. First, Coleman [6] realized that one could apply more sophisticated p-adic analytic machinery such as Riemann-Roch and Newton polygons to explicitly estimate the number of zeros of ω, and thus the number of zeros of f_ω. Another limitation of the original method was that it required analysis in $\mathbb{C}_p$, for p a prime of good reduction for C. Since the first prime of good reduction might be arbitrarily large, even Coleman's Riemann-Roch improvement gave non-uniform bounds on the number of rational points on C. Michael Stoll, in the groundbreaking paper [13], described a theory of p-adic integration on annuli that allowed for uniform bounds using primes of bad reduction. (Lorenzini–Tucker [10] had actually already used primes of bad reduction, but they required a regular model of the curve in question, which led to non-uniform bounds.)

Prior to Stoll's work, it was known that the p-adic integral that Chabauty and Coleman used in the case of good reduction is actually two different integrals that happen to agree in this case. These two integrals, the Berkovich-Coleman and Abelian integrals, no longer agree in the bad reduction case because the "tube"

$$]y[= \{x \in C \mid \mathrm{red}(x) = y\}$$

is an annulus when y is a non-smooth point of the special fiber. (Here red refers to the reduction map from a curve to the special fiber.)

Stoll realized that one can use differentials for which the two integrals agree, and forcing the two integrals to be equal in this way consists of one linear condition on the space of differentials. (He imposes one more linear condition to choose a branch of the p-adic logarithm used in the Berkovich-Coleman integral.) When C is a hyperelliptic curve, Stoll explicitly describes the differentials on each tube. He then applies standard Newton polygon arguments and a little linear algebra on each tube, and then incorporates a sophisticated analysis of the combinatorics of the special fiber to prove the following:

Theorem 2 ([13], Theorem 1.4). *Suppose $C\colon y^2 = f(x)$ is a hyperelliptic curve defined over $\mathbb{Q}$. Assume further that $g \geq 3$ and $r := \operatorname{rank} \mathrm{Jac}_{\mathrm{C}}(\mathbb{Q}) \leq g - 3$. Then*

$$\#C(\mathbb{Q}) \leq 33(g-1) + 1 \textit{ if } r = 0$$

and

$$\#C(\mathbb{Q}) \leq 8rg + 33(g-1) - 1 \textit{ if } r \geq 1.$$

Note here that the given bound is linear in both g and r. Katz, Rabinoff, and Zureick-Brown [9] have shown that the number of rational points on any curve defined over $\mathbb{Q}$ with $r < g - 2$ is bounded quadratically by g, nearly completing the Chabauty program. Based on Stoll's results, however, one actually expects a bound that is linear in both r and g and reduces to their bound when r is close to g.

In this paper we extend Stoll's work to superelliptic curves, that is curves of the form $y^m = f(x)$ for $m \geq 2$ and f a rational polynomial of degree at least 4. Equivalently, superelliptic curves are curves with Galois m-to-one maps to $\mathbb{P}^1$. We obtain a linear bound in g and r for each fixed m, thus confirming the idea that small values of r should lead to better Chabauty bounds.

The following theorem is a consequence of our main theorem, though we have weakened the bounds considerably to make it easier on the reader's eyes.

Theorem 3. *Let C be the superelliptic curve defined above. Suppose $r < g - 2$, and let p be* *the smallest prime congruent to 1 modulo m. Then*

$$\#C(\mathbb{Q}) < 2(2g-2)(r+3) + \frac{2g-2}{m} + (5p+2)(g-1) + 4r.$$

Much of Stoll's work transfers from the hyperelliptic case to the superelliptic case. Perhaps surprisingly, the most technical input for our generalization is Raynaud's classification of order m automorphisms of the p-adic disc and annulus when $(p,m)=1$. (Incidentally, the $p \mid m$ case is much more involved and plays a large part in the solution of the local lifting conjecture in [11].)

Example. *For concreteness, consider the case $m=3$ and $r=0$. The curve*

$$2y^3 = z^4 - 10z^3 + 35z^2 - 50z + 24$$

is one such curve with genus 3. Then Theorem 3 gives a coarse bound of $\frac{149}{3}(g-1)$. In contrast, the techniques of [9] give a bound of $76g^2 - 82g + 22$.

In the second half of the paper we extend ideas of Amini–Baker and Katz–Zureick-Brown to show how one might prove a rank-favorable bound for arbitrary curves satisfying the Chabauty rank hypothesis. In particular, we exhibit a bound similar to theirs that does not require a regularity hypothesis. While we are not yet able to prove uniform statements, we do establish a combinatorial criterion – "the Main Assumption" – from which one can bound $\#C(\mathbb{Q})$. If one could prove that this criterion holds in a uniform way for all curves then uniform rank-favorable bounds would follow.

1.1 Notation

We now fix some notation for the rest of our discussion. From here onward, p will always be a prime that does not divide m. We will let K be a number field, and k the completion of K at some prime $\mathfrak{p}$ lying over $p \in \mathbb{Z}$. We denote by ζ_m a primitive mth root of unity in k or K depending on the context. For any extension of p-adic fields, e will denote the ramification index of the extension, and f will denote the residue degree.

Let C be a superelliptic curve over a number field K defined by an affine equation $y^m = f(x)$ for some $m \geq 2$ and $f \in K[X]$, and we will assume C defines a curve of genus at least two. In any case, C comes with an automorphism τ of order m given by $y \mapsto \zeta_m y$, and the quotient of C by the subgroup generated by τ is $\mathbb{P}^1$.

Recall that the p-adic unit open disc of radius r is just the set

$$D_{0,k} = \{y \in \mathbb{C}_p \mid |x - y| < r\}.$$

Similarly, we define the p-adic annulus with outer radius r and inner radius α by

$$A_{\alpha,k} = \{y \in \mathbb{C}_p \mid \alpha < |x - y| < r\}.$$

Finally, C has a g-dimensional vector space of p-adic analytic differentials denoted $H^0(C, \Omega^1_{\mathbb{Q}_p})$.

2 Chabauty's Method and Stoll's Improvement

2.1 The Technique

We now describe in more detail the technology behind Chabauty's method. For this section C is any curve of genus at least two. We may replace C by a semistable model so that its special fiber has no cusps. (See [8], Prop. 4.1 for a precise argument.) If x is any point of the special fiber of C, then by [3] $]x[$ is either a tube or an annulus depending on whether x is smooth or singular, respectively. From here on we suppose that there is some K-rational $P \in]x[$; otherwise our bounds are trivially satisfied

The key lemma of Chabauty's method revolves around a p-adic integral which vanishes on the p-adic closure of $\mathrm{Jac}_{\mathrm{C}}(\mathbb{Q})$.

Definition 4. *Let A be an abelian variety over $\mathbb{C}_p$. The abelian logarithm on A is the unique homomorphism of $\mathbb{C}_p$-Lie groups $\log\colon A(\mathbb{C}_p) \to \mathrm{Lie}(A)$ such that*

$$d\log\colon \mathrm{Lie}(A) \to \mathrm{Lie}(\mathrm{Lie}(A)) = \mathrm{Lie}(A)$$

is the identity map. Now by definition, $\mathrm{Lie}(A)$ is the dual of $\Omega^1_{A/\mathbb{C}_p}(A)$, and we denote the evaluation pairing by $\langle \cdot, \cdot \rangle$.

Finally, for $x, y \in A(\mathbb{C}_p)$ and $\omega \in \Omega^1_{A/\mathbb{C}_p}(A)$ we set

$$^{Ab}\int_O^x := \langle \log(x), \omega \rangle$$

and we call $^{Ab}\int$ the abelian integral on A. We also define $^{Ab}\int_x^y \omega = {}^{Ab}\int_O^x - {}^{Ab}\int_O^y$. (We have *let O denote the identity of A.) When C is a curve and ι is an Abel-Jacobi mapping into its Jacobian, we define the integral from $x \in C$ to $y \in C$ by integrating from $\iota(x)$ to $\iota(y)$ in* $\mathrm{Jac_C}$.

Fixing a residue tube $]x[$, Chabauty's method centers around the subspace V_{chab} of differentials such that

$$f_\omega(z) := {}^{Ab}\!\!\int_P^z \omega = 0$$

for all $z \in \mathrm{Jac_C}(\mathbb{Q}) \cap]x[$. Since $C(\mathbb{Q}) \subseteq \mathrm{Jac_C}(\mathbb{Q})$, every rational point of C is a zero of f_ω for all $\omega \in V_{\mathrm{chab}}$. The following lemma powers Chabauty's method.

Lemma 5. *Let* $r := \mathrm{rank}(\mathrm{Jac_C}(\mathbb{Q}))$*. Then* $\dim(V_{chab}) \geq g - r$*.*

The second integral that one can define on the curve C is called the Berkovich-Coleman integral – it is essentially defined by formal anti-differentiation and is therefore (1) dependent just on the endpoints of the integral on annuli, discs, and more general "wide opens" and (2) it is amenable to the tools of Newton polygons.

Definition 6. *Suppose ω is an analytic differential on the annulus $A_{\alpha,k}$ with local coordinate T. Consider the local power series expansion of ω:*

$$\omega = g(T)\frac{dT}{T} = \sum_{n=-\infty}^{\infty} a_n T^n \frac{dT}{T}$$

where $g(T)$ converges on $A_{\alpha,k}$. Denote by f the (incomplete) formal antiderivative of g given by

$$\sum_{n \neq 0} \frac{a_{n-1}}{n} T^n.$$

Finally, let

$$\mathrm{Log}(T) = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{(T-1)^n}{n}$$

be the branch of the logarithm where $\mathrm{Log}(p) = 0$. *Then we define the Berkovich-Coleman integral on* C *via the formula*

$$^{BC}\!\!\int_x^y \omega = (f(y) + a_0 \,\mathrm{Log}(y)) - (f(x) - a_0 \,\mathrm{Log}(x)).$$

Stoll proves in [13] the following integral comparison theorem, which [9] generalizes using analytic uniformization of p-adic abelian varieties.

Theorem 7. *Suppose* $x, y \in C$ *lie in the same residue annulus. Then there is a codimension two subspace of* $\Omega^1_{C/\mathbb{C}_p}(C)$ *such that*

$$^{Ab}\!\!\int_x^y \omega = {}^{BC}\!\!\int_x^y \omega$$

Now Lemma 5 and Theorem 7 help produce a differential which satisfies the equality of the Berkovich-Coleman and Abelian integrals and also lies in V_{chab}: if $r < g-2$ we conclude that there is a nonzero subspace W_{chab} of differentials of dimension at least $g - r - 2$ for which the Abelian and Coleman integrals are equal, and which vanishes on the closure of the p-adic points of the Jacobian. The main work of this paper is to bound the number of zeros of such a differential on a residue tube.

To summarize this section in one sentence: there is a "relatively large" subspace W_{chab} such that: (1) the integrals of elements of W_{chab} vanish on $C(\mathbb{Q})$ (owing to the Abelian integral), and (2) the integral can be computed using formal antiderivatives (owing to the

Berkovich–Coleman integral).

3 Automorphisms of the Annulus

Emulating [13], we proceed with the following process: A tube on the curve is a disc or an annulus, so it is sent to another disc or annulus by the superelliptic automorphism (it may be sent to itself). Since we know the possible automorphisms of the disc and annulus, we know the behavior of the discs and annulus under the quotient map. Finally, once we know the ramification behavior of the quotient map we use it to write explicit equations for the annulus on the curve. Having this equation in hand, we can explicitly describe a basis for the differentials on the annulus or disc.

In [13], Michael Stoll classifies the involutions of the p-adic disc and annulus "by hand." We here describe more general results of Raynaud that classify finite order automorphisms of the p-adic annulus and disc using algebraic methods.

Theorem 8 ([12], Props. 2.3.1, 2.3.2)**.**

1. *Let $\tau\colon D_{0,k} \to D_{0,k}$ be an analytic map of order m, with m coprime to p. Then after an analytic change of coordinates, τ is just multiplication by an m^{th} root of unity. In particular, τ has only one fixed point and $D_{0,k}/\langle\tau\rangle$ is a disc.*

2. *Let $\tau\colon A_{\alpha,k} \to A_{\alpha,k}$ be an analytic map of order m such that $|\tau(z)| = |z|$ for all $z \in A_\alpha(\mathbb{C}_p)$. Then after an analytic change of coordinates, τ is just multiplication by an m^{th} root of unity. In particular, τ has no fixed points and $A_{\alpha,k}/\langle\tau\rangle$ is an annulus.*

3. *Now suppose m is even. Let $\tau\colon A_{\alpha,k} \to A_{\alpha,k}$ be an analytic map of order m such that*

$|\tau(z)| = \alpha/|z|$ for some α. Then, after a change of coordinates, τ is given by either $z \mapsto \zeta/z$, where ζ_m is a primitive m^{th} root of unity, or $z \mapsto \zeta' a/z$ for some a with $|a| = \alpha$ and ζ' is a primiitve $m/2^{th}$ root of unity. In particular, τ has two fixed points and $A_{\alpha,k}/\langle\tau\rangle$ is a disc.

4 Annuli on Superelliptic Curves

Based on the preceding lemmas, we can write explicit equations for the annuli on the curve. We again follow Stoll's notation. Denote by Θ the set of branch points of the m-fold cover of $\mathbb{P}^1$ associated to C, alias the zeros of f. By a change of coordinates we can always assume that infinity is not a branch point, so we can write

$$y^m = f(x) = c \prod_{\theta \in \Theta} (x - \theta).$$

If $\theta \neq 0$, we have the functions

$$f_\theta^+(x) = \left(1 - \frac{\theta}{x}\right)^{1/m}, f_\theta^-(x) = \left(1 - \frac{x}{\theta}\right)^{1/m}.$$

These converge when $|x| < |\theta|$ and $|x| > |\theta|$, respectively. They satisfy the equations

$$x - \theta = x f_\theta^+(x)^m \text{ and } x - \theta = -\theta f_\theta^-(x)^m.$$

The following lemmas are generalizations of those in [13]. The first gives equations for both discs and annuli whose quotient is a disc. In this case, such a quotient is either completely ramified or completely split.

Lemma 9. *Let $\phi\colon D_{0,k} \to D \subseteq \mathbb{P}^1_k$ be a parametrization of an open disc by the unit open disc.*

1. *Suppose $D(\mathbb{C}_p) \cap \Theta = \emptyset$. If, for any $c \in D(k)$, $f(c)$ is not an m^{th} power in k, then $\pi^{-1}(D) \cap C(k)$ is empty. If $f(c)$ is an m^{th} power, $\pi^{-1}(D)$ is the disjoint union of m disjoint open discs in C, each isomorphic to D via π.*

2. *Suppose $D(\mathbb{C}_p) \cap \Theta = \{\theta_1\}$, and that D has radius r. Then $\theta_1 \in k$; further assume that there is some $c \in k$ such that $r|f'(\theta_1)| = |c|^m$. Then $\pi^{-1}(D)$ is a disc on C, and up to an analytic change of coordinates the map π is just the m-th power map (i.e. the superelliptic automorphism acts by rotation.)*

3. *Suppose $D(\mathbb{C}_p) \cap \Theta = \{\theta_1, \theta_2\}$. Then $(x-\theta_1)(x-\theta_2)$ has coefficients in k. Furthermore, the set $\pi^{-1}(D)$ is either contained in the preimage of the smallest closed disc containing θ_1 and θ_2, or $\pi^{-1}(D)$ is an annulus A in C such that, after an analytic change of coordinates, $\pi(z) = z^{m/2} + \beta/z^{m/2}$ for some $\beta \in k^\times$. In this case τ acts as $z \mapsto \zeta_{m/2}\beta/z$.*

Proof. The proofs follow Stoll's work *mutatis mutandis*, so we simply record here the relevant parameterizations of the discs and annuli on C.

1. When $\pi^{-1}(D) \cap C(\mathbb{C}_p) \neq \emptyset$, there exists $\gamma \in k^\times$ such that $f(0) = \gamma^m$. Then we have the parameterization

$$D^{(i)} = \{(z, \zeta_m^i \gamma h(z) \mid z \in D)\}, \ i = 1, \ldots, m.$$

2. By assumption there exist $\gamma \in k^\times$ and $u \in \mathcal{O}_k^\times$ such that $\gamma^m = uf'(\theta_1)$. We have the parameterization
$$D = \{uz^m, \gamma z h(uz^m) \mid z \in D_{0,k}\}.$$

3. Changing coordinates so that $\theta_1 + \theta_2 = 0$, the equation for the curve is given by $y^m = c'(x^2 - a)h(x)^m$. The convergence properties of the m-th root when $(p, m) = 1$ is independent of m, and so if $|x| > |a|$ and c' is not an m-th power in k, x cannot be the coordinate of a k-point. In this case the preimage of D in C is contained in the preimage of $\{|x| < \theta_1\}$.

   If $c' = \gamma^m$ with $\gamma \in k$ then a parameterization $\phi : A_{\alpha^{1/m}} \to \pi^{-1}(A)$ is given by
$$z \mapsto \left(z^{m/2} + \frac{a}{4}z^{-m/2}, \gamma(z^{m/2} - \frac{a}{4}z^{-m/2})^{2/m} h(z^{m/2} + \frac{a}{4}z^{-m/2})\right).$$

$\square$

The second lemma describes equations for annuli whose quotient is an annulus, in the process generalizing Stoll's lemma to account for the many possible behaviors of an automorphism of an m-to-one mapping. Indeed, while the case $m = 2$ separates cleanly into odd and even annuli, we have many more cases.

Lemma 10. *Let $\phi\colon A_{\alpha,k} \to A \in \mathbb{P}^1_k$ be an open annulus such that $A \cap \Theta = \emptyset$ and $A(k) \neq \emptyset$. The complement of A in $\mathbb{P}^1_k$ is the disjoint union of two closed discs, which partition Θ into Θ_0 and Θ_1. This partition induces a factorization $f(x) = cf_0(x)f_\infty(x)$ with f_0 and f_∞ monic such that the roots of f_0 are the elements of Θ_0 and the roots of f_∞ are the elements of Θ_0.*

Write $l = \gcd(\#\Theta_0, m)$. If c' is not an l^{th} power in k, then $\pi^{-1}(A) \cap C(k) = \emptyset$. Assume otherwise, and further suppose that there exists $\gamma \in k$ with $u^l c' = \gamma^m$. Then $\pi^{-1}(A)$ is a union of l annuli, and after an analytic change of coordinates the superelliptic automorphism acts by both interchanging these annuli and rotating them by $\zeta_m^{m/l}$.

Proof. By construction $\#\Theta_0$ is invertible modulo m/l; let $\#\Theta_0^{-1}$ denote the smallest integer representative for its inverse in this ring. Then $\#\Theta_0 \#\Theta_0^{-1} = 1 + n(m/l)$ for some integer n.

The parameterization is given by

$$A^{(j)} = \{(u^{l\#\Theta_0^{-1}} z^{m/l}, \zeta_m^{jm/l} \gamma u^n z^{\#\Theta_0/l} h(u^{l\#\Theta_0^{-1}} z^{m/l})) \mid z \in A_{0,k}\},\ j = 1, \ldots, l.$$

□

Combining Raynaud's classification of automorphisms of the annulus with the preceding two lemmata, we conclude:

Lemma 11. *The preceding two lemmas provide an exhaustive list of the parameterizations of maximal annuli on the curve C.*

Proof. Suppose $\phi \colon A \to C$ is a maximal annulus such that $A(k) \neq 0$. Then A can be parametrized using Lemmas 9 and 10. We have two cases.

1. $\tau(A) \cap A = \emptyset$.

   In this case π is an isomorphism from A onto its image, and we conclude that $\pi(A)$ is an annulus containing no ramification points. Thus, by Lemma 10, each component of $\pi^{-1}(\pi(A))$ can be parametrized via the map we described there. In particular, the lemma provides a parameterization of A itself. Of course, this parametrization depends on the number of branch points within $\pi(A)$, or rather the greatest common divisor of that number with m.

2. $\tau(A) \cap A \neq \emptyset$, and τ preserves the orientation of the chain corresponding to A.

   In this case we have $\tau(A) = A$, and a standard argument from non-archimedean geometry shows that the orientation condition implies that $|(\phi^*\tau)(z)| = |z|$. Raynaud's classification (Theorem 8) implies that (up to an analytic change of coordinates) $\phi^*\tau$ is a rotation of order dividing m on $A_{\alpha,k}$. Thus $\pi(A)$ is an annulus containing no branch points.

   We conclude that A is parametrized via Lemma 10. Furthermore, in this case $(m, \Theta_0) = 1$ since $\pi^{-1}(\pi(A)) = A$.

3. $\tau(A) \cap A \neq \emptyset$, and τ reverses the orientation of the chain corresponding to A.

   As above, $\tau(A) = A$ and the orientation condition implies that $|\phi^*\tau(z)| = \alpha/|z|$ for $z \in A_{\alpha,k}$. Raynaud's classification tells us that $\phi^*\tau$ is an inversion composed with a rotation of order dividing m, so $\pi(A)$ is a disc containing two branch points. It is thus parametrized via the proof of Lemma 9.

□

5 Bounding Zeros of Differentials

Our goal in this section is to obtain bounds on the number of zeros of differentials on the annuli which cover our superelliptic curve C. In general, the Weierstrass Preparation

Theorem ([5], Theorem 2.4.3) says that an analytic function on an annulus A can be written in the form $f = gu$, where g is a Laurent polynomial with finitely many exponents and u has no zeros on A. One of Stoll's key insights was that there is a basis for Ω^1_C in which every basis element has the same u in the Weierstrass decomposition. This uniform description allows us to induce cancellation in the g-component.

Differentials on a smooth plane curve with equation $f(x,y) = 0$ have an explicit basis given by

$$\omega^{(j)} = x^i \frac{dx}{\frac{\partial f}{\partial y}},$$

for $0 \le i \le g-1$. For C, after clearing an m from the denominator these terms are of the form

$$x^i \frac{dx}{y^{m-1}}.$$

Based on our classification of annuli and discs on C, we thus have the following two local descriptions of differentials:

1. A arises as in Lemma 9.

   A similar computation shows

$$\phi^*\omega^{(j)} = \frac{\left(z^{m/2} + \frac{a}{4}z^{-m/2}\right)^j \left(\frac{m}{2}z^{m/2} - \frac{ma}{8}z^{-m/2}\right)}{\left(z^{m/2} - \frac{a}{4}z^{-m/2}\right)^{2/m}} \eta(z) \frac{dz}{z},$$

   where again $\eta(z)$ is an analytic function with no zeros on A.

2. A arises as in Lemma 10.

Recalling the notation $l = (m, \#\Theta_0)$, we have

$$\phi^*\omega^{(j)} = z^{mj/l+m/l-\#\Theta_0(m-1)/l}\eta(z)\frac{dz}{z}.$$

Here $\eta(z)$ is a constant multiple of a power of the function $h(z)$ and hence is nonzero on A.

Theorem 12. *Let $V \neq 0$ be a subspace of codimension at least one of Ω^1_C. Then there exists a nonzero differential $\omega \in V$ such that $\phi^*\omega = g(z)u(z)dz/z$, where g is a finite Laurent series with its highest and lowest exponents differing by at most $m(r+3)+1$ (in the case of annuli arising as in Lemma 9) or $m(r+2)/l+1$ (in the case of annuli arising as in Lemma 10) and u is an analytic function that is nonzero on A.*

Proof. Let W be the subspace of Ω^1_C spanned by $\{\omega^{(j)} \mid 0 \leq j \leq r+2\}$. The $\omega^{(j)}$ form a basis of Ω^1_C, and so W has dimension $r+3$. By assumption V has dimension at least $g-r-2$, so $W \cap V \neq \{0\}$. Let ω be a nonzero element in this intersection.

1. Suppose A arises as in Lemma 10. Then, since $\omega \in W$, we have the description

$$\phi^*\omega = \sum_{j=0}^{r+2} z^{mj/l+m/l-\#\Theta_0(m-1)/l}\eta(z)\frac{dz}{z}.$$

Factoring out $\eta(z)\frac{dz}{z}$, the remaining sum has exponents that range from $m/l-\#\Theta_0(m-1)/l$ to $m(r+2)/l+m/l-\#\Theta_0(m-1)/l$. Thus the highest and lowest exponents differ by at most $m(r+2)/l+1$, inclusive. Here $u = \eta$.

2. Suppose A arises as in Lemma 9. Then any nonzero $\omega \in V \cap W$ is of the form

$$\phi^*\omega = \sum_{j=0}^{r+2} \frac{(z^{m/2} + \frac{a}{4}z^{-m/2})^j(\frac{m}{2}z^{m/2} - \frac{ma}{8}z^{-m/2})}{(z^{m/2} - \frac{a}{4}z^{-m/2})^{2/m}}\eta(z)\frac{dz}{z}.$$

Factoring out the denominator and $\eta(z)\frac{dz}{z}$, the remaining polynomial has an exponent spread of at most $m(r+3)+1$.

The denominator, $(z^{m/2} - \frac{a}{4}z^{-m/2})^{2/m}$, is analytic on the annulus A, and so the end term we factored out is nonzero on A. We set $u = \frac{\eta}{(z^{m/2} - \frac{a}{4}z^{-m/2})^{2/m}}$. Our conclusion follows.

□

6 The Final Count(down)

6.1 p-Adic Rolle's Theorem

The strategy of effective Chabauty hinges on being able to control the zeros of $\int_P^z \omega$ based on the zeros of ω using a p-adic Rolle's theorem. The theory of Newton polygons (and in more recent applications, sophisticated tropical geometry) provides the tools for this analysis.

If $p > e+1$, we define

$$\mu := 1 + \frac{e}{p-e-1}.$$

Proposition 13 ([13], Prop. 7.7)**.** *Suppose a p-adic analytic differential ω has a Newton* *polygon of length d in a p-adic annulus A, and that $p > 2g$. Then*

$$f_\omega(z) := \int_P^z \omega$$

has at most μd zeros on A.

6.2 Uniform Bounds

Remark 14. *Throughout this section we will assume that $K = \mathbb{Q}$ and $k = \mathbb{Q}_p$ for a prime p that we will pick in the next section. It seems possible to obtain a bound when K is an arbitrary number field in the same way Stoll does, but we feel as though a diversion in this direction would detract from the explicit nature of the bounds over $\mathbb{Q}$.*

It is possible – though certainly difficult – to bound the number of rational points that reduce to smooth points of the special fiber of C. This bound has been established in two ways in the literature: Via the alternative rank functions of Katz and Zureick-Brown [9] and through the theory of metrized complexes defined by Amini and Baker [2]. We restate their theorem here.

For the purposes of this section, it is helpful to denote by C_D the portion of $C(\mathbb{Q}_p)$ covered by discs and by C_A the portion covered by annuli.

Proposition 15 ([13], Lemma 7.1, using [8], Thm. 4.4)**.** *Let $V \neq 0$ be a linear subspace of codimension r of the space of regular differentials on C and let N_D denote the number of discs whose union is $C_D(k)$. Suppose further that $p > e + 1$. Then the integrals f_ω for $\omega \in V$ have at most*

$$N_D + 2\mu r \leq (5q + 2)(g - 1) - 3q(t - 1) + 2\mu r.$$

common zeros in $C_D(k)$.

To calculate the number of rational points lying in $C_A(\mathbb{Q}_p)$, we conclude the discussions of the previous sections.

Proposition 16. *Suppose $p > e+1$. Then the number of common zeros in $C_A(\mathbb{Q}_p)$ of all f_ω for $\omega \in V_{chab}$ is bounded by*

$$((2g-2)/m-2))\,\mu\,(m(r+3)+1)\,.$$

Proof. Our residue tube analysis shows that, in fact, each orbit of annuli that arise in the case of a non-inverting action of τ contains at most $m(r+2)+l \leq m(r+3)$ shared zeros. When τ inverts the annulus, the number of shared zeros on the annulus is at most $m(r+3)+1$. We therefore take $m(r+3)+1$ as a uniform bound for the number common zeros of the differentials in V on any orbit of annuli. Applying Stoll's Newton polygon calculation for the optimal differential in the residue orbit, this leaves at most $\mu(m(r+3)+1)$ common zeros of the integrals f_ω.

How many such orbits of annuli can there be? Each orbit corresponds to an edge in the image of a minimal skeleton of C via the analytification of the map π. This image is obtained by starting with the convex hull of the ramification points in $\mathbb{P}^1$ (a tree with d leaves) and removing the leaves. This leaves a tree with at most $d-2$ nodes, hence at most $d-3$ edges.

Now a simple computation with Riemann-Hurwitz shows that

$$2g-2 = md - m - d - (m,d).$$

Certainly, then,

$$d \leq (2g-2)/m + 1,$$

so there are at most $(2g-2)/m - 2$ orbits of annuli.

Combining the bounds on C_D and C_A, the total number of common zeros of all f_ω for $\omega \in V_{\text{chab}}$ is bounded by

$$((2g-2)/m - 2))\,\mu\,(m(r+3)+1) + (5q+2)(g-1) - 3q(t-1) + 2\mu r.$$

This completes the proof of Theorem 3, barring the choice of a prime p.

$\square$

Now we optimize the choice of a prime that satisfies $p > e+1$ and $(p, m) = 1$. A naive approach might proceed as follows: We have that

$$[\mathbb{Q}_p(\zeta_m) : \mathbb{Q}_p] \leq \phi(m).$$

A standard theorem of ramification theory then says that

$$ef = [\mathbb{Q}_p(\zeta_m) : \mathbb{Q}_p],$$

where f is the residue degree of the extension. Thus e is certainly bounded above by $\phi(m)$, and we conclude that we must choose $p > \phi(m) + 1 \leq m$. By Bertrand's Postulate there exists such a p that is less than $2n$. Using the trivial bound $f \leq \phi(m)$ from above, we see that

$$q \leq (2m-1)^{m-1}.$$

We can immediately improve this astronomical bound as follows: The $(p-1)^{\text{th}}$ roots of unity are the only roots of unity in $\mathbb{Q}_p$ for p odd, and so $\zeta_m \in \mathbb{Q}_p$ if and only if $p = 1 \pmod{m}$. For such a prime we have $e = f = 1$.

The problem of finding the smallest prime in an arithmetic progression is answered by Linnik's Theorem. The theorem says that there exists an L and m_0 such that for all $m > m_0$, the smallest prime congruent to one modulo m is less than a constant times m^L. Recent work has shown that L can be taken a little under 5, but at the cost of astronomical bounds for m_0. Under the GRH, the smallest prime congruent to one modulo m is less than $m(\log m)^2$. We will content ourselves with an easily digested exponential bound, and trust the reader to search for the smallest prime in an arithmetic progression in any one specific case or use a polynomial bound in general if they want to:

Theorem 17 ([14])**.** *The smallest prime congruent to one modulo m is at most $2^{\phi(m)} - 1$.*

Putting all of this together, we conclude that for each m, $\#C(\mathbb{Q})$ is bounded by a bilinear polynomial in r and g. Furthermore, the dependence on m is at worst polynomial in nature, and can be bounded easily as a function of $2^{\phi(m)} - 1$.

This concludes the proof of the main theorem.

7 Discussion

One might wonder whether the bound in the main theorem can be taken completely independent of m. Based on our methods, even the most fanciful conjectures on the least prime in an arithmetic progression still inject some dependence on m into our final bounds. In any case, under RH the dependence is relatively small.

The main theorem thus provides infinitely many classes of curves C (as we vary m in the main theorem) for which the number of rational points on C is be bounded linearly in r and g, and raises the question of whether such a bound might hold for all curves.

7.1 Other Possible Attacks on Superelliptic Curves

In this section we note how one might use a tower of superelliptic curves to glean arithmetic information about the individual curves.

Given the curve $y^m = f(x)$ as above, and the curve $y^s = f(x)$ for any divisor s of m, we always have a cover

$$\rho : C \to C' : y^s = f(x).$$

This cover is given by the map $(x, y) \mapsto (x, y^{m/s})$.

Proposition 18. *Suppose that $\#C'(K) \leq B$. Then $\#C(K) \leq R(K)B$, where $R(K)$ denotes the number of m/s-th roots of unity in K.*

Proof. Suppose P is a K-rational point of C. Then $\rho(P)$ is K-rational, too. Thus given $Q = (x_0, y_0) \in C'(K)$, it lifts to a K-rational point on C if and only if y_0 is an m/s-th power in K, say $y_0 = s_0^{m/s}$. In this case any other lift of Q is given by $(x_0, \zeta_{m/s}^i s_0)$. □

We would like to then replace m by its smallest prime divisor and analyze the resulting curve C', but it is not always true that C' satisfies the rank hypothesis of Chabauty's method. It is true, however, that one can relate the genus of the two curves using Riemann-Hurwitz, so if Chabauty's method does work on C' then we may obtain a genus-dependent (but rank-ignorant) bound on the rational points of C'.

8 Metrized Complexes

In this section we use the theory of metrized complexes to understand how one might prove a rank-favorable bound for arbitrary curves satisfying the Chabauty rank hypothesis. Omid Amini and Matthew Baker invented metrized complexes to interpolate between two existing theories of the degeneration of linear series. On one hand, the Eisenbud-Harris theory is best applied to curves whose dual graph is a tree. On the other hand, previous work of Baker established a theory which worked mostly with curves whose dual graph has the same Betti number as the curve itself. We give here a bare-bones account of the theory; see [2] for more details and proofs.

Definition 19. *A metrized complex over a field κ consists of:*

1. *A metric graph Γ;*
2. *a collection of irreducible and nonsingular curves $\{C_v\}$, one for each vertex $v \in \Gamma$; and*
3. *for each edge $e \in \Gamma$ and endpoint v of e, an identification of v with a κ-point of C_v. We let A_v denote set of the points of C_v which have been identified with vertices of Γ.*

Now let X be a proper curve over a discrete valuation ring K with quotient field κ. Associated to X is a metrized complex: The curves C_v are just the components of the special fiber, the edges are identified with the non-smooth points of the special fiber, and the endpoints of each edge are identified with the evident non-smooth points on the curves C_v. To define the length of an edge e: Recall that if X is semistable over K, it is locally of the form

$$\frac{K[[x,y]]}{xy = p^l}$$

around a point P which reduces to a non-smooth point of the special fiber, where p is a uniformizer of K. Then we let l be the length of the edge associated to P. Notice that Γ is a metric version of the dual graph of the special fiber of X.

However, the most conceptually satisfying description of l is as the logarithm of the inner radius of the residue annulus $]\bar{P}[$. We describe the Berkovich picture more clearly now, because it underlies much of the theory. The Berkovich analytification X^{an} contains $X(K)$. As we would expect of an analytic object, however, this set plays a minimal role in the analytification. The most important part of X^{an} for us will be its skeleton – a canonically embedded metric graph which is isomorphic to the metric graph Γ defined in the previous paragraph. Berkovich proved that there is a canonical retraction map from $X(K)$ to the skeleton; see [9] for a complete account of this story. For a point P for which $\bar{P}$ is singular, it is defined as follows: Once we have identified the residue tube $]\bar{P}[$ with an annulus of outer radius one, P lies at some radius r_0 on this annulus. We then retract P to the point on e of distance $\log_p r_0$ from from the vertex v.

We can use Berkovich's retraction to define our own retraction τ_* from K-points of X

to points on $\mathfrak{CX}$. When $P \in X(K)$ reduces to a smooth point $\bar{P}$ of the special fiber on the component C_v, we define $\tau_*(P)$ to be $\bar{P} \in C_v$ (such a reduction is defined because our model is proper.) When $\bar{P}$ is a singular point of the special fiber, $\tau_*(P)$ is a point on the edge corresponding to the reduction of P. We retract it using the retraction map to the skeleton of the Berkovich analytification X^{an}, so that its exact location is determined by where it lies on the annulus $]\bar{P}[$.

Extending τ_* by linearity, there is a map from divisors on X to divisors on $\mathfrak{CX}$. A divisor on $\mathfrak{CX}$ is what one would expect: a finite combination of points that are either on the curves C_v or on the open edges $e \in \Gamma$. There is a parallel theory of divisors and linear equivalence on $\mathfrak{CX}$. We state these precisely – they will be important to us.

Definition 20. *A rational function $\mathfrak{f}$ on $\mathfrak{CX}$ consists of a rational function f_v on C_v in the usual sense, along with a piecewise-linear function f_Γ on the metric graph Γ.*

If $\mathfrak{CX}$ is the metrized complex associated to a curve X, there is a tropicalization map trop that takes a rational function f on X and produces a rational function $\mathfrak{f} = \mathrm{trop}(f)$ on $\mathfrak{CX}$. It is important to note here that there may be more rational functions on $\mathfrak{CX}$ than those that come from tropicalizations, since in the above definition we do not require any compatibility between f_v and f_Γ. One can use the new notion of rational functions to define linear equivalence on $\mathfrak{CX}$. However, it will be sufficient for us to know the following two facts:

Proposition 21. *If f is a rational function on X, then $\tau_*(\mathrm{div}(f)) = \mathrm{div}(\mathrm{trop}(f))$.*

In other words, whatever the correct notion of the divisor of a piecewise-linear function is, taking the divisor of a tropicalized function is the same as first taking the divisor on X

and then retracting. The second fact that we'll need is that linear equivalence on $\mathfrak{CX}$ is built up from simpler chip-firing moves.

Proposition 22. *The notion of equivalence for* $\mathrm{Div}(\mathfrak{CX})$ *is generated by the following primitive equivalences.*

1. *(Equivalence on C_v.) A divisor supported only on a curve C_v is equivalent to a divisor on C_v which is equivalent in the usual sense.*

2. *(Equivalence on an edge e.) Suppose $q \in e$, and let ϵ be a number which is less than or equal to the smallest distance from q to a vertex. Let $q_\epsilon^\pm$ be the two points at distance ϵ from q on the edge e. Then $2q$ is equivalent to $p_\epsilon^+ + p_\epsilon^-$.*

3. *(Firing from a vertex to surrounding edges.) Suppose v is a vertex of Γ with outgoing edges $e_1, \ldots, e_k$, and ϵ is a number smaller than the length of each e_k. Let $q_\epsilon^1, \ldots, q_\epsilon^k$ be points at distance ϵ from v on the respective edges. Finally, suppose $s_1, \ldots, s_k$ are the points on C_v which have been identified with the endpoints of $e_1, \ldots, e_k$. Then the divisor $\sum_i s_i$ is equivalent to the divisor $\sum_i q_\epsilon^i$.*

It turns out that the definition of divisor rank on curves – based on the dimension of $L(D)$ – is extremely ill-behaved for tropical divisors. Instead, we have the following definition.

Definition 23. *Let D be a divisor on $\mathfrak{CX}$. The rank $r(D)$ is the largest integer such that $D - E$ is equivalent to an effective divisor for any effective divisor E. If D is not equivalent to an effective divisor, then we set $r(D) = -1$.*

Now consider a set $\mathcal{H} = \{H_v\}$, where H_v is a set of rational functions on the curve C_v. Then the restricted rank of the pair $(D, \mathcal{H})$ is the largest integer such that $D - E$ is equivalent *to an effective divisor for any effective divisor E, and where equivalences of type 2 must be contained in the set H_v. Note that the rank agrees with the restricted rank when $H_v = \kappa(C_v)$.*

If D has degree d and $(D, \mathcal{H})$ has rank r then we say $(D, \mathcal{H})$ is a $\mathfrak{g}^r_d$.

The notions of divisors and equivalence on $\mathfrak{C}\mathfrak{X}$ are compatible with the parallel notions on X. For example, τ_* takes principal divisors to principal divisors, effective divisors to effective divisors, and equivalent divisors to equivalent divisors.

Definition 24. *Let K_v denote a canonical divisor on C_v. A canonical divisor $\mathcal{K}$ on $\mathfrak{C}\mathfrak{X}$ is any divisor equivalent to*

$$\sum_{v \in V} K_v + A_v.$$

Amini and Baker use their Riemann-Roch for metrized complexes to show that the retraction of a canonical divisor on X is a canonical divisor on $\mathfrak{C}\mathfrak{X}$. They also prove a version of Clifford's theorem which we state here. Recall that a special divisor is a subdivisor of a canonical divisor.

Theorem 25 ([2], Theorem 3.4)**.** *Suppose $D \in \mathrm{Div}(\mathfrak{C}\mathfrak{X})$ is special. Then*

$$r(D) \leq \frac{\deg(D)}{2}.$$

9 A New Chabauty Divisor

We define a Chabauty divisor on $\mathfrak{C}\mathfrak{X}$ as follows. Say Q is a smooth point of the special fiber. Then, mimicking KZB and Stoll, we choose a differential $\omega \in V$ that vanishes to the least degree on the point Q after normalizing ω so that it doesn't have any poles or zeros along the whole component C_v on which Q lives. In the notation of KZB, choosing an optimal differential for each smooth point of the special fiber results in a divisor

$$D^g := \sum_{Q \in X(\mathbb{F}_p)^{sm}} n_Q Q.$$

If Q is a singular point of the special fiber then it has an associated edge e in the metric graph portion of $\mathfrak{C}\mathfrak{X}$. For each differential $\omega \in V$, we may retract its divisor to $\mathfrak{C}\mathfrak{X}$ using the map τ_*. Consider the divisor

$$D'_e(\omega) := \tau_*(\operatorname{div} \omega) \restriction_e$$

which just takes the divisor associated to a differential, retracts it to $\mathfrak{C}\mathfrak{X}$, and forgets all points that don't lie on the edge e. Then some differential ω_e will minimize the degree of $D_e(\omega)$; we finally set

$$D_e = D'_e(\omega_e).$$

Packaging these edge-divisors together, we define the "graphical component" of the Chabauty divisor by

$$D^\Gamma := \sum_{e \in \Gamma} D_e.$$

Last but not least, the Chabauty divisor on our metrized complex is

$$D_{\text{chab}} = D^\Gamma + D^g.$$

10 Bounding $\deg D_{\text{chab}}$

10.1 Katz and Zureick-Brown's Analysis

At this point in the analysis, Katz and Zureick-Brown require two main ingredients to their proof.

D_{chab} is Special

The Chabauty divisor of Stoll–Katz–Zureick-Brown is special because every differential vanishes to order at least n_Q at the point Q: if we pick an arbitrary differential ω, $\operatorname{div}\omega$ is a canonical divisor that vanishes to at least the same order as D_{chab} at every point, and thus contains D_{chab} as a subdivisor. They then use a modified version of Clifford's theorem on the special fiber which shows that, since $K - D_{\text{chab}}$ is also special,

$$r(K - D_{\text{chab}}) \leq \deg(K - D_{chab})/2.$$

D_{chab} has Large Rank

Again, because every differential in V_{chab} satisfies D_{chab}, we have that $h^0(X_{\mathbb{F}_p}, \Omega^1_{\mathbb{Q}_p}(-D_{\text{chab}})) \geq g - r$, and so $r(K - D_{\text{chab}}) \geq g - r - 1$. Putting these two inequalities together,

$$g - r - 1 \leq (2g - 2 - \deg D_{\text{chab}})/2,$$

and solving for the degree gives the desired inequality $\deg D_{\text{chab}} \leq 2r$.

10.2 $\deg D_{\text{chab}}$ on the Metrized Complex

One runs into a number of immediate issues when generalizing the argument of [8] to the case of a non-regular model of the curve, and from now on D_{chab} will only refer to the metrized complex Chabauty divisor defined in this paper.

D_{chab} is Special

D_{chab} is no longer special, since not only the order of vanishing of a differential matters, but also the location of its zeros along an edge. Thus not every differential satisfies D_{chab} on the metrized complex. It is however true that $D^g + D_e$ is special for each edge $e \in \Gamma$, since it is a subdivisor of the canonical divisor $\operatorname{div} \omega_e$ corresponding to the optimal differential for that edge. So not every differential satisfies the conditions imposed by D_{chab}, but one differential satisfies part of it, and that is all we will need. Recall that the degree of a canonical divisor on the metrized complex $\mathfrak{C}$ is $2g-2$. Applying Amini–Baker's Clifford's theorem gives us the inequality

$$r_{\mathfrak{C}}(\mathcal{K} - D^g - D_e) \leq \frac{\deg(\mathcal{K} - D_{\text{chab}})}{2} = g - 1 - \frac{\deg(D^g) + \deg(D_e)}{2}.$$

$K - D_{\text{chab}}$ has Large Rank

Here again Katz and Zureick-Brown used the fact that every differential satisfied the conditions imposed by D_{chab}; since this is no longer the case for our Chabauty divisor, we need a further assumption (the "Main Assumption" below.) The following theorem is key to the Chabauty investigations of Amini and Baker:

Proposition 26. *Let X be a smooth proper curve over K, $\mathfrak{X}$ a strongly semistable model for X, and $\mathfrak{C}\mathfrak{X}$ the metrized complex associated to $\mathfrak{X}$. Let D be a divisor on X and let $L_\eta = (\mathcal{L}(D), H_\eta)$ for $H_\eta \subseteq H^0(X, \mathcal{L}(D)) \subseteq K(X)$, be a $\mathfrak{g}^r_d$ on X. For any vertex $v \in V$, define H_v as the κ-vector space defined by the reduction to $\kappa(C_v)$ of all the rational functions in H_η, and let $\mathcal{H} = \{H_v\}_{v \in V}$. Furthermore set $\mathcal{D} = \tau_*(D)$*

Let E° be an effective divisor of degree e supported on the smooth locus of the special fiber of $\mathfrak{X}$, and define $\mathcal{E}^\circ = \sum_{v \in V} E^\circ_v \in \mathrm{Div}(\mathfrak{C}\mathfrak{X})$, where E°_v is the restriction of E° to C_v. Suppose that $H_v \subseteq L(D_v - E^\circ_v)$ for all v. Then the pair $(\mathcal{D} - \mathcal{E}^\circ, \mathcal{H})$ is a limit $\mathfrak{g}^r_{d-e}$ on $\mathfrak{C}\mathfrak{X}$.

Remark 27. *Two crucial notes apply to this statement. The fact that the rank of $(\mathcal{D}-\mathcal{E}^\circ, \mathcal{H})$ has rank at least r is exactly saying that, for any degree r effective divisor $\mathcal{F} \in \mathrm{Div}\, \mathfrak{C}\mathfrak{X}$ there exists a linear equivalence $\mathfrak{f}_1$ such that*

$$\mathcal{D} - \mathcal{E}^\circ - \mathcal{F} + \mathrm{div}\, \mathfrak{f}_1 \geq 0.$$

First, $\mathcal{F}$ can be taken to have support only on the smooth portion of the special fiber. Second, $\mathfrak{f}_1$ is constructed as the tropicalization of an element $f_1 \in H_\eta$.

We upgrade their lemma, under an important assumption, to include divisors supported on the metric graph portion of the metrized complex $\mathfrak{C}\mathfrak{X}$.

Lemma 28. *Carrying over the setup of the previous proposition, let $\mathcal{E}^\times$ be an effective divisor of degree h supported only on the edges of $\mathfrak{C}\mathfrak{X}$. For any rational function, denote its tropicalization using a fraktur font.*

Assumption 29 (Main Assumption)**.** *Suppose furthermore that for every $f_1 \in H_\eta$, we have that $\mathcal{E}^\times$ is equivalent to a subdivisor of* $\operatorname{div}(\mathfrak{f}_1)$ *via an equivalence $\mathfrak{f}_2$ that is supported only on edges of $\mathfrak{CX}$.*

Then the pair $(\mathcal{D} - \mathcal{E}^\circ - \mathcal{E}^\times, \mathcal{H})$ is a limit $\mathfrak{g}^r_{d-e}$ on $\mathfrak{CX}$.

Remark 30. *Although Lemma 28 is quite general, we will in restrict ourselves to the case in which $\mathfrak{f}_2$ makes $\mathcal{E}^\times$ equivalent to a subdivisor of* $\operatorname{div}(\mathfrak{f}_1)$ *restricted to one particular edge $e \in \Gamma$.*

Proof. By the proposition from Amini–Baker, for any degree r effective divisor $\mathcal{F} \in \operatorname{Div} \mathfrak{CX}$ there exists a linear equivalence $\mathfrak{f}_1$ such that

$$\mathcal{D} - \mathcal{E}^\circ - \mathcal{F} + \operatorname{div} \mathfrak{f}_1 \geq 0.$$

Let's examine the divisor

$$\mathcal{D} - \mathcal{E}^\circ - \mathcal{E}^\times - \mathcal{F} + \operatorname{div} \mathfrak{f}_1 + \operatorname{div} \mathfrak{f}_2.$$

On any smooth point of the special fiber this divisor is effective because both $\operatorname{div}(\mathfrak{f}_2)$ and $\mathcal{E}^\times$ are supported on the edges of the metrized complex.

On the other hand, the main assumption says that $\operatorname{div}(\mathfrak{f}_1) + \operatorname{div}(\mathfrak{f}_2) - \mathcal{E}^\times$ is effective. Since $\mathcal{D}$ is effective and $\mathcal{F}$ is supported on the smooth part of the special fiber, we see that

$$\mathcal{D} - \mathcal{E}^\circ - \mathcal{E}^\times - \mathcal{F} + \operatorname{div} \mathfrak{f}_1 + \operatorname{div} \mathfrak{f}_2 \geq 0,$$

and we are done.

$\square$

10.3 Applying the Lemma to Chabauty

Amini and Baker use the above lemma with $\mathcal{D} = \mathcal{K}$, a canonical divisor, and $\mathcal{E}^\circ = D_{\text{chab}}$. We instead take $\mathcal{E}^\circ = D^g$, while $\mathcal{E}^\times$ will be a subdivisor of D_e, with $\deg(\mathcal{E}^\times) \sim \deg(D_e)$. For convenience $\mathcal{K}$ will be the canonical divisor associated to ω_e, the optimal differential for the edge e. They furthermore let H_η be the subspace of $L(K)$ corresponding to V_{chab}. Following through the same steps as Amini–Baker, we get a bound on the degree of $D^g + \mathcal{E}^\times$ and thus a bound on the degree of $D_g + D_e$. Finally, we sum over all edges e and obtain the following theorem:

Theorem 31. *By Lemma 26, for every edge $e \in \Gamma$ there is a rational function $\mathfrak{f}_{1,e}$ that displays the required equivalence. Let δ_e be the differential associated to $\mathfrak{f}_{1,e}$ in the bijection between $h^0(\Omega^1)$ and $L(-K)$. Assume that for every edge $e \in \Gamma$ there exists an effective divisor $\mathcal{E}_e^\times \leq D_e$ such that $\mathcal{E}_e^\times$ is equivalent to a subdivisor of $D_e'(\delta_e)$. (In other words, the Main Assumption is satisfied for D_e.) Suppose furthermore that there exists a constant $\alpha_e \in \mathbb{N}$ such that* $\deg(D_e) \leq \alpha_e \deg(\mathcal{E}^\times)$. *Then*

$$\#E \deg(D^g) + \sum_{e \in E} \frac{1}{\alpha_e} \deg(D_e) \leq 2r\#E.$$

Proof. Applying Lemma 28 and Clifford's theorem, we see that

$$\deg(D_g) + \deg(\mathcal{E}_e^\times) \leq 2r,$$

and hence

$$\deg(D_g) + \frac{1}{\alpha_e}\deg(D_e) \leq 2r.$$

Summing over E completes the proof.

$\square$

Remark 32. *Given this lemma, we have that*

$$\sum \deg(D_e) \leq 2r\#E \max_{e\in\Gamma}\{\alpha_e\}.$$

Since $\#E \leq 3g-3$ for a semistable curve ([9], Lemma 4.14), an absolute bound on the constants α_e would lead to a rank-favorable bound on $\deg(D_{\mathrm{chab}})$.

We argue in the next section that the p-adic Rolle's theorem provides evidence for the main assumption. We have to modify $\mathcal{E}^\times$ even further there: $\mathcal{E}^\times$ will be the subdivisor of zeros of ω_e which lie in between *common* zeros of the integrals of $\omega \in V_{\mathrm{chab}}$

11 Evidence for the Main Assumption

What does the Main Assumption amount to, given the discussion of applying our lemma to Chabauty's method? It requires that D_e is equivalent to a subdivisor of

$$\tau_*(\operatorname{div}\delta_e)$$

for the differential $\delta_e \in V_{\mathrm{chab}}$ via an equivalence supported only on the edge e. This section is meant to discuss the implications of a strict version of Rolle's Theorem for the differentials

δ_e and ω_e. A strict Rolle's theorem does not hold for all power series (there are easy counterexamples for polynomials even), but the point is that we might reasonably expect that the Main Assumption holds for just the pertinent differentials δ_e and ω_e. In fact, a weak version of the p-adic Rolle's theorem we discuss below does hold.

Here's the idea. We'd like to bound the number of common zeros of the integrals. Over the real numbers, Rolle's Theorem tells us that the zeros of the differentials interlace with the zeros of the integrals. If two functions have two common zeros, then the derivatives of the functions each have a zero in the open interval dictated by the common zeros. We would like to use a similar constraint in the p-adic case to prove that the Main Assumption holds if a strict version of Rolle's Theorem holds.

Remark 33. *We say that a strict version of Rolle's theorem holds for a p-adic analytic function f if, for any distinct zeros $a \neq b$ of f with $|a| \geq |b|$, there is a zero c of f' with $|b| \leq |c| \leq |a|$.*

A weaker version of Rolle's theorem is known for discs, with a "fudge factor" in the interval:

Theorem 34 ([1], p. 316)**.** *Let $f \in \mathbb{C}_p[[X]]$ have convergence radius $r_f > 1$. If f has two distinct zeros $a \neq b$ in the closed unit ball satisfying $|a-b| \leq |p|^{\frac{1}{p-1}}$ then f' has a zero in the closed unit ball.*

As far as the author knows, a generalization of this statement for p-adic annuli is missing from the literature.

Assume that a strict version of Rolle's theorem holds for both ω_e and δ_e on the edge e. Suppose further that the integrals of δ_e and ω_e have no zeros on e with the same valuation.

We will construct $\mathcal{E}^{\times}$ in two steps. First, let $\mathcal{A}^{\times}$ be the subdivisor of D_e constructed as follows: For every two neighboring common zeros of the functions $\int \delta_e$ and $\int \omega_e$, keep the zero of ω_e that lies in between them. In the spirit of chip firing, we can think of our task as a game. We have the divisors $D_e'(\delta_e)$ and $\mathcal{A}^{\times}$. Consider z_1, z_2, z_3 common zeros of the integrals $\int \delta_e$ and $\int \omega_e$. Then by the strict version of Rolle's theorem, both δ_e and ω_e have zeros in the intervals $[|z_1|, |z_2|]$ and $[|z_2|, |z_3|]$. Consider these two zeros of ω_e as a subdivisor of $\mathcal{R}^{\times}$. It's not hard to see that an elementary equivalence can take make two zeros of δ_e cover at least one of the zeros of ω_e. Add the zero that can be covered to $\mathcal{E}^{\times}$.

Continuing this process, we get a divisor $\mathcal{E}_e^{\times}$ on the edge e. By construction, the degree of $\mathcal{E}_e^{\times}$ is approximately half of the number of common zeros of $\int \omega_e$ and $\int \delta_e$ on the edge e. This observation replaces the use of p-adic Rolle's in [8]: whereas they bound $\deg(D_{\mathrm{chab}})$ and then pass to zeros of the integrals, the present strategy aims to use a Rolle's theorem first and then conclude a bound on the zeros of the integrals by construction.

One upshot of this strategy is that it uses rational functions on $\mathfrak{C}\mathfrak{X}$ that don't come from rational functions on X. It is thus largely combinatorial, so for example it is not be hard to see that the numbers α_e could be small if the edge-divisors D_e and $D_e'(\delta_e)$ are relatively equidistributed on the edge e. Of course, any hope of a uniform bound with either strategy depends on combinatorial properties of $\tau_*(\mathrm{div}(\omega_e))$ on the edge e that must wait until future work.

References

[1] Robert M. Alain. *A course in p-adic Analysis*, volume 198 of *Graduate Texts in Mathematics*. Springer-Verlag, 2000.

[2] Omid Amini and Matthew Baker. Linear series on metrized complexes of algebraic curves. *M. Math. Ann.*, 362:55–106, 2015.

[3] Siegfried Bosch, Ulrich Güntzer, and Reinhold Remmert. *Non-Archimedean Analysis*. Springer-Verlag, 1984.

[4] Lucia Caporaso, Joe Harris, and Barry Mazur. Uniformity of rational points. *Journal of the American Mathematical Society*, 10(1):1–35, January 1997.

[5] William Cherry. Lectures on non-archimedean function theory. In *Advanced School on p-Adic Analysis and Applications*, Trieste, Italy, 2009. The Abdus Salam International Centre for Theoretical Physics.

[6] Robert F. Coleman. Effective chabauty. *Duke Math J.*, 52(3):765–770, 1985.

[7] Marc Hindry and Joseph H. Silverman. *Diophantine Geometry: An Introduction*. Springer, 2000.

[8] Eric Katz and David Zureick Brown. The chabauty-coleman bound at a prime of bad reduction and clifford bounds for geometric rank functions. *Compositio Mathematica*, 149(11), 2013.

[9] Eric Katz, Joseph Rabinoff, and David Zureick-Brown. Uniform bounds for the number of rational points on curves of small Mordell–Weil rank. *Duke Mathematical Journal*, 165(16), 2016.

[10] Dino Lorenzini and Thomas J. Tucker. Thue equations and the method of chabauty-coleman. *Inventiones mathematicae*, 48:47–77, 2002.

[11] Florian Pop. The oort conjecture on lifting covers of curves. *Annals of Mathematics*, 180(1):285–322, 2014.

[12] Michel Raynaud. Spécialisation des revêtements en caractéristique $p > 0$. *Ann. scient. Éc. Norm. Sup.*, 32(1):87–126, 1999.

[13] Michael Stoll. Uniform bounds for the number of rational points on hyperelliptic curves of small Mordell-Weil rank. *Duke Mathematical Journal*, April 2015.

[14] R. Thangadurai and A. Vatwani. The least prime congruent to one modulo n. *The American Mathematical Monthly*, 118(8):737–742, 2011.