USING WREATH PRODUCTS TO FIND PERIODIC POINTS

LUCY EPSTEIN

Abstract. This paper will cover the results of Wreath Products and Proportions of Periodic Points, a 2016 paper by Juul, Kurlberg, Madhu, and Tucker [JKMT16]. The goal of this paper is to give an approximate sketch of the proofs in [JKMT16] while still being comprehensible to undergraduates with an intermediate understanding of arithmetic dynamics. As a motivating example, we go over the rho method of factorization, a significantly faster probabilistic method of finding a prime factor by Pollard [Pol75], and its connections to [JKMT16].

Contents

1. Introduction

Arithmetic dynamics is a subject focused on the images of points under repeated applications of functions. This makes periodicity one of the most important properties a point can have as it specifies the entire orbit of the point. Preperiodic points, or points which become periodic after a finite number of applications of a function, are also studied in detail. This paper aims to provide a summary of the results of [JKMT16], a set of results classifying the proportion of periodic points of multiple large classes of rational functions. In particular, their results show that for sufficiently large primes p, these classes of functions will have a very low proportion of periodic points when taken mod p.

Throughout the paper, the notation $\varphi^n(x)$ for a function φ will be used exclusively to mean applying φ n times to x. Also, the variable p will exclusively refer to a prime in $\mathbb{Z}$.

This paper will focus on specifically the case of rational periodic points; note that the original paper generalizes most of its results to (in most cases) number fields or Noetherian integral domains. We now state (simplified versions of) the main results of [JKMT16], beginning with a very general claim that "most" functions have very few periodic points for sufficiently large primes.

Date: May 14, 2025.

THEOREM 1.1 ([JKMT16, Theorem 1.2]). *Let $\epsilon > 0$ and $d > 1$. For almost all*[1] *rational functions $\varphi \in \mathbb{Q}(x)$ of degree d,*

$$\limsup_{p\to\infty} \frac{\#\mathrm{Per}(\varphi_p)}{p+1} \leq \epsilon,$$

where φ_p is the reduction of φ mod p.

The second theorem gives some sufficient conditions on rather general rational functions over number fields to have decreasing density of periodic points; we again hone in on the case of rational coefficients for simplicity.

THEOREM 1.2 ([JKMT16, Theorem 1.3]). *Let $\varphi \in \mathbb{Q}(x)$ be a rational function of degree $d > 1$ such that for any two distinct critical points α_1, α_2 of φ and any positive integers m, n, we have $\varphi^m(\alpha_1) \neq \varphi^n(\alpha_2)$ unless $m = n$ and $\alpha_1 = \alpha_2$. Then*

(1)

$$\liminf_{p\to\infty} \frac{\#\mathrm{Per}(\varphi_p)}{p+1} = 0$$

(2) if $\mathbb{Q}$ is algebraically closed in the splitting field of $\varphi(x) - t$ over $\mathbb{Q}(t)$, then we have

$$\lim_{p\to\infty} \frac{\#\mathrm{Per}(\varphi_p)}{p+1} = 0.$$

Finally, their third main theorem focuses on the quadratic case, where polynomials can be conjugated nicely into the form $x^2 + c$ and so total classification as to whether the density of periodic points approaches 0 is possible. Theorem 1.3 can then be applied to Section 2, our motivating example of a method of factorization, by taking $f(x) = x^2 - 1$.

THEOREM 1.3 ([JKMT16, Theorem 1.5]). *Let $f \in \mathbb{Q}[x]$ be a quadratic polynomial. Then*

$$\lim_{p\to\infty} \frac{\#\mathrm{Per}(\varphi_p)}{p+1} = 0$$

unless there is a linear polynomial $\sigma = ax + b \in \mathbb{Q}[x]$ such that $\sigma^{-1} f \sigma$ is equal to the Chebyshev polynomial $x^2 - 2$.

Now, we give an outline of the paper. First, we give an overview of Pollard's rho method of factorization [Pol75] for finding a prime factor p of a number in $O(\sqrt{p})$ time, a significant improvement over the classical $O(p)$ bound. This method works well as a motivating example as it effectively uses the theory from [JKMT16] for a seemingly unrelated but useful application. Next, we explain the concept of wreath products and their relevance to the proofs of the above theorems. Finally, we sketch the proofs of the main theorems, citing the intermediate lemmas used.

[1] For a given d, one can express the set of rational functions with degree d as points in $(2d+2)$-dimensional space such that the leading coordinates are nonzero and the polynomials are relatively prime (this condition comes from computing the resultant). The theorem statement claims that a dense subset of these polynomials, under the Zariski topology relating to prime ideals, satisfy the supremum property.

2. Motivating example: the rho factorization method

We begin with a motivating example, due to Pollard, for finding a prime factor p of a given n in $O(\sqrt{p})$ time assuming that the function $f(x) = x^2 - 1$ has very few periodic points for large p. This assumption was not formally proven in [JKMT16], but a similar idea comes from applying Theorem 1.3 to f.

(paraphrased from [Pol75, Section 1]) Let n be an integer and consider the sequence $(x_i)_{i\in\mathbb{Z}} : 0 \leq x_i \leq n-1$ such that $x_i \equiv \begin{cases} 2 & i = 0 \\ x_{i-1}^2 - 1 & i \geq 1 \end{cases} \pmod{n}$. We can also form the sequence $Q_i \equiv \prod_{j=1}^{i}(x_{2j} - x_j) \pmod{n}$. Note that calculating the triple (x_i, x_{2i}, Q_i) takes 4 mod n multiplications from the previous triple: one step taking x_{i-1} to x_i, two steps taking x_{2i-2} to x_{2i}, and one step taking Q_{i-1} to Q_i. Then, for some fixed integers m and S (e.g. $m = 100, S = 10^4$), compute $d_i = \gcd(Q_i, n)$ when $m|i$ with some relatively quick algorithm such as the Euclidean algorithm. If d_i is a proper divisor of n greater than 1, this gives a partial factorization of n and the process can then be continued taking n/d_i as the new number to factor.

Pollard used this process in Section 2 of [Pol75] in order to compute the expected amount of calculations for Q_i to share a common factor with n. If r is the least integer such that $x_r \equiv x_{2r} \pmod{p}$ for some prime factor $p|n$, all Q_i for $i > r$ will share the common factor p with n, so it suffices to compute the expected value of r. Since every point in $\mathbb{F}_p$ is necessarily preperiodic for any function by the Pigeonhole Principle, there will exist some t and c such that it takes t terms for the sequence (x_i) to reach a periodic sequence of period c. Now, we have an immediate bound $r < t + c$ if $t > 0$ as r will always be the first multiple of c that is at least t.

Pollard relied on the assumption that $f(x) = x^2 - 1$ is a "random mapping" in the sense that cycle length is extraordinarily unlikely, if not impossible, to be a substantial (linear) portion of the whole field. With this assumption, he calculated that the expectation of $c(p)$ and $t(p)$ would each be approximately $0.6267\sqrt{p}$ and the expectation of $r(p)$ would be approximately $1.0308\sqrt{p}$; he also calculated some worst-case bounds of $\max_{p<n} r(p)$ for fixed n, which still gave results on the order of $\sqrt{p}$. These calculations all occurred in 1975, so computers have improved significantly since then for checking bounds, but the algorithm remains relevant.

The key assumption of randomness here relies on the polynomial $x^2 - 1$ and its behavior in $\mathbb{F}_p$; Juul Kurlberg Madhu Tucker [JKMT16] showed that the density of periodic points of this polynomial approaches 0 for large p in Theorem 1.3, so cycle length must be $o(p)$ for p sufficiently large. While their result does not directly show the $O(\sqrt{p})$ estimate, it still shows improvement over more naive methods of finding a prime factor.

3. Wreath products and number theoretic preliminaries

Now, we move on from factoring to talk about wreath products. In general, wreath products are a construction taking two subgroups of permutation groups on a finite number of variables to another subgroup of a permutation group. We start by defining the wreath product of two groups G and H acting on the sets $\{1, \ldots, a\}$ and $\{1, \ldots, b\}$, respectively.

Definition 3.1. For G and H as above, let $\pi \in G$ and $\tau_1, \ldots, \tau_a \in H$. Then, we can define the *wreath product* $\sigma = (\pi; \tau_1, \ldots, \tau_a)$ acting on $\{1, \ldots, a\} \times \{1, \ldots, b\}$

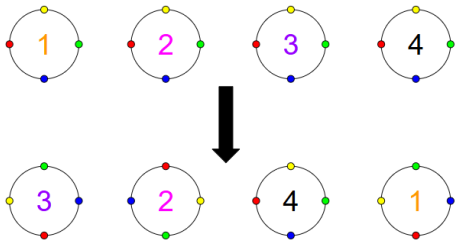

FIGURE 1. Let $G = S_4$ be the symmetric group on 4 elements and let $H = C_4$ be the cyclic group on 4 elements. Then, the element $\sigma := ((134); (1234), (1432), 1, (1234))$ of the wreath product $G[H]$ would act on $\{1, 2, 3, 4\} \times \{1, 2, 3, 4\}$ as in the figure. Here, the first coordinate is the large colored numbers (modified by $(134) \in G$) and the second coordinate is the sets of smaller colored circles (modified by the respective elements of H).

by sending (i, j) to $(\pi(i), \tau_i(j))$. The group of such wreath products across all possible $\pi \in G$ and $\tau_i \in H$ (note the τ_i do not have to be distinct) is denoted $G[H]$. Since there is a natural isomorphism $\{1, \dots, a\} \times \{1, \dots, b\} \to \{1, \dots, ab\}$ sending $(i, j) \mapsto ib + j$, we have a natural action of $G[H]$ on $\{1, \dots, ab\}$ and can therefore construct further wreath products involving $G[H]$. The nth *wreath power* of a group G, denoted $[G]^n$, is computed as $[G]^1 = G$ and $[G]^n = [G]^{n-1}[G]$.

Wreath products could be imagined to be named as such because real wreaths are generally formed as intertwined circles of sprigs from trees. In a similar vein, one can imagine a wreath product as a set of a interwoven circles (each representing the numbers $\{1, \dots, b\}$, with the ith circle reordered by the action of τ_i) woven together with the action of π. For example, consider Figure 1.

Why are wreath products, seemingly a construction of permutations, relevant in arithmetic dynamics? The answer lies within the Galois group of the splitting field of a polynomial, which acts as a permutation. To get there, we need a bit of number theory.

DEFINITION 3.2. For a polynomial $f(x) \in k[x]$ for some field k, the *splitting field* of f over k is the smallest field over which f can be decomposed into a product of linear factors.

EXAMPLE 3.3. The splitting field of $f(x) = x^3 - 2$ over $\mathbb{Q}$ is $\mathbb{Q}(\sqrt[3]{2}, \omega\sqrt[3]{2}, \omega^2\sqrt[3]{2}) = \mathbb{Q}(\sqrt[3]{2}, \omega)$, where $\omega = e^{2i\pi/3}$. To construct this, we adjoin the roots of f over $\overline{\mathbb{Q}}$ to the base field, $\mathbb{Q}$, and simplify roots that generate each other.

An important number theoretic fact is that the Galois group of the splitting field of a polynomial over a field acts by permuting the roots (although not necessarily all of them). To apply this, Juul Kurlberg Madhu Tucker [JKMT16] noted that in general, composing two polynomials of degrees a and b over a field K such that their composition has ab distinct roots in $\overline{K}$ means that the Galois group of the composition will be a subgroup of the wreath product of the original Galois groups, generalizing a result of Odoni [Odo85, Lemma 4.1]. Looking to generalize to rational functions, they examined the Galois groups of the splitting fields of multiple applications of φ in order to find periodic points in the following theorem.

THEOREM 3.4 ([JKMT16, Theorem 3.1]). *Let k be a field. Suppose that $\varphi(x) \in k(x)$ is a rational function of degree $d \geq 2$ such that $\varphi'(x) \neq 0$. Let K_n be the splitting field of $\varphi^n(x) - t$ over $k(t)$, $E = K_1 \cap \overline{k}$, and $G_n = \mathrm{Gal}(K_n/E(t))$. Fix $N \in \mathbb{N}$ and let $\varphi_c := \{z : \varphi'(z) = 0\}$. If for any $a, b \in \varphi_c$ and $r, s \leq N$, we have $\varphi^r(a) \neq \varphi^s(b)$ unless $a = b$ and $r = s$, then $G_N \cong [G]^N$.*

To work towards the proof of Theorem 3.4, we first need to define ramification of a prime ideal.

DEFINITION 3.5. Let L and K be number fields with L a finite extension of K. Primes $\mathfrak{p}$ of $\mathfrak{o}_K$ will factor as products of primes $\mathfrak{q}$ of $\mathfrak{o}_L$; say $\mathfrak{p} = \prod_{i=1}^{k} \mathfrak{q}_i^{a_i}$. If any of the a_i is greater than 1, $\mathfrak{p}$ is said to *ramify*. Otherwise, it is *unramified*.

With ramification, we can now sketch the proof of Theorem 3.4. (add in lemmas individually?)

Proof. The main idea of the proof is strong induction; the base case $N = 1$ is immediate. Let $\alpha_1, \ldots, \alpha_{d^n}$ be the (necessarily distinct) roots of $\varphi^n(x) - t$ in $\overline{k(t)}$, let M_i be the splitting field of $\varphi(x) - \alpha_i$ over $E(\alpha_i) = E(t, \alpha_i)$, and let $\hat{M}_i = K_n[\prod_{j \neq i} M_j]$.

The proof uses multiple lemmas, the first of which states that G_n must be a subgroup of $[G]^n$, giving the easier direction of the final claim. This first lemma is proved with induction on n with the lemma that the composition of functions gives a subgroup of the corresponding wreath product of Galois groups. The next lemma states that the primes of $E(t)$ that ramify in K_n take the form $(\varphi^m(a) - t)$ for some critical point $a \in \varphi_c$ and $m \leq n$, which is proved with the help of a lemma of Cullinan and Hajir [CH12, Proposition 1]. Finally, the last lemma used states that for $n < N$ and $a \in \varphi_c$, the prime $(\varphi(a) - \alpha_i)$ in $E(\alpha_i)$ does not ramify in $\hat{M}_i$, relying on the assumption that critical points are not sent to the same value for $n < N$ and the previous lemma in the proof.

In the main proof, the goal is to bound the size of G_{N+1} with the strong induction hypothesis and then use that it must be a subgroup of $[G]^{N+1}$ to show that they are the same group. This is done by finding subgroups in $\mathrm{Gal}(M_i/E(\alpha_i))$ and $\mathrm{Gal}(K_{N+1}/\hat{M}_i)$ which are isomorphic. The former must be isomorphic to G by the assumption of critical points being sent to different values and the latter is isomorphic to a subgroup of G, so it must be the whole thing. The inductive hypothesis then finishes the proof. □

4. Main theorems and proof sketches

We now move on to proving the main theorems after having built up some preliminary tools for working with wreath products in Section 3. We begin with the claim that for most primes p, reduction mod p will be "good" in the sense of preserving roots and Galois groups of polynomials.

PROPOSITION 4.1 ([JKMT16, Proposition 4.1]). *Let $h(x) \in \mathbb{Q}[x]$ be a nonconstant irreducible polynomial, let $(\theta_1, \ldots, \theta_d)$ be roots of h in some splitting field for h over $\mathbb{Q}$, and let $B = A[\theta_1, \ldots, \theta_d]$. Then, for all but finitely many primes p, we have the following:*

We can reduce h mod p to get a rational function h_p. We can do the same for B to get B_p and a reduction map $r_p : B \to B_p$. There is also a homomorphism $\rho_p : \mathrm{Gal}(h(x)/\mathbb{Q}) \to \mathrm{Gal}(h_p(x)/\mathbb{F}_p)$ with the property that $\rho_p(\sigma)(r_p(\theta_i)) = r_p(\sigma(\theta_i))$

for all $\sigma \in \mathrm{Gal}(h(x)/\mathbb{Q})$ and roots θ_i. For such primes p, r_p induces a bijection between the roots of h and the roots of h_p and ρ_p is an isomorphism of groups.

Proof. Choose p to not divide any denominators of elements or differences between elements; then, everything other than the statement about the isomorphism of roots of h_p and ρ_p is clear since the roots will have distinct reductions. For nontrivial $\sigma \in \mathrm{Gal}(h(x)/\mathbb{Q})$, meaning $\sigma(\theta_i) = \theta_j$ for some $i \neq j$, we have that $\rho_p(\sigma)(r_p(\theta_i)) \neq r_p(\theta_i)$ by the property of ρ_p, so $\rho_p(\sigma)$ is not the identity and thus ρ_p is injective. The other direction comes from bounding the size of Galois groups with a result from Odoni [Odo85, Lemma 2.4]. □

We introduce a bit more notation specific to periodic points to help with the next proposition.

Definition 4.2 ([JKMT16, cf. Definition 1.6]). If H is a group acting on a set S, then we define $\mathrm{FPP}(H)$ to be the proportion of elements of H fixing at least one $s \in S$.

Example 4.3. Let $H = S_3$ be the symmetric group on 3 elements and $S = \{1, 2, 3\}$. Then, the only elements of H with any fixed points are the identity and the 3 transpositions, so $\mathrm{FPP}(H) = \frac{4}{6} = \frac{2}{3}$.

One can think of FPP as perhaps standing for "Fixed Point Proportion" for the proportion of elements of H which have a fixed point. We now use the FPP of the Galois group as a bound for the amount of periodic points for large p.

Proposition 4.4 ([JKMT16, Proposition 5.3]). *We use notation as in Theorem 3.4, taking $k = \mathbb{Q}$. Suppose that $\mathbb{Q}$ is algebraically closed in K_n and let $\delta > 0$. Then, there is a constant M_δ such that for all $p > M_\delta$, we have*

$$\frac{\#\mathrm{Per}(\varphi_p)}{p+1} \leq \mathrm{FPP}(G_n) + \delta.$$

Proof. Let p be a prime for which φ has good reduction. The main idea of the proof is to show that elements of $\varphi_p^n(\mathbb{F}_p)$ are in bijection with elements of G_n having a fixed point and then get an estimate of the amount of these with an effective version of the Chebotarev density theorem for number fields. The Chebotarev density theorem is a generalization of Dirichlet's theorem, stating that the primes which factor completely in a Galois extension of $\mathbb{Q}$ of degree n have density $1/n$. However, this application of the theorem takes an entirely different form, stated by Murty and Scherk [MS94, Theorem 1]. The proof then finishes with noting that all periodic points will always be elements of φ_p^n and then bounding the remaining terms other than $\mathrm{FPP}(G_n)$ for sufficiently large p, giving the delta-epsilon result. □

Combining Propositions 4.1 and 4.4 and some results by Odoni on $\mathrm{FPP}([S_d]^n)$, Juul Kurlberg Madhu Tucker [JKMT16] were able to prove Theorem 1.1.

Proof of Theorem 1.1. The first part of the proof uses three results of Odoni to first bound $\mathrm{FPP}([S_d]^n)$, a wreath product of Galois groups, and then show that the Galois group of the splitting field of $h_n(x) := \varphi^n(x) - t = 0$ is $[S_d]^n$ itself by bounding its size [Odo85, Theorem 1 and Lemmas 2.4 and 4.3]. Then, combining Proposition 4.1 and Proposition 4.4 gives a sufficiently tight bound on $\frac{\#\mathrm{Per}(\varphi_p)}{p+1}$ for large enough p. □

We can now continue towards the proof of the second main theorem, starting with an immediate corollary of Proposition 4.4 proven with delta-epsilon limit notation.

COROLLARY 4.5 ([JKMT16, Corollary 5.4]). *With notation as in Proposition 4.4, suppose that $\mathbb{Q}$ is algebraically closed in K_n for all n. Then, if*

$$\lim_{n\to\infty} \mathrm{FPP}(G_n) = 0,$$

we have

$$\lim_{p\to\infty} \frac{\#\mathrm{Per}(\varphi_p)}{p+1} = 0.$$

Next, we start working towards the ability to "pull back" primes from $\mathbb{Z}$ to a general ring of integers of a number field while keeping similar properties on the limits of proportions of periodic points.

LEMMA 4.6 ([JKMT16, Lemma 6.2]). *Let k' be a finite extension of $\mathbb{Q}$ and let $\varphi \in \mathbb{Q}[x]$ and let p be a prime of good reduction for φ. Let $\mathfrak{q}$ be a prime of k' such that $\mathfrak{q} \cap \mathbb{Z} = (p)$ and $[(\mathfrak{o}_{k'}/\mathfrak{q}) : \mathbb{F}_p] = 1$. Then φ induces a map $\tilde{\varphi}$ over k' such that $\tilde{\varphi}$ has good reduction at $\mathfrak{q}$ and we have $\#\mathrm{Per}(\tilde{\varphi}_{\mathfrak{q}}) = \#\mathrm{Per}(\varphi_p)$.*

Proof. The proof of Lemma 4.6 is more or less a direct check; the property that $\mathfrak{o}_{k'}/\mathfrak{q}$ has degree 1 over $\mathbb{F}_p$ means that for any $\beta \in \mathfrak{o}_{k'}/\mathfrak{q}$, we can find a unique $\alpha \in \mathbb{Z}/p\mathbb{Z}$ with $\beta \equiv \alpha \pmod{\mathfrak{q}}$. This correspondence between the quotients gives a natural bijection between $\mathbb{P}^1(\mathfrak{o}_{k'}/\mathfrak{q}) \to \mathbb{P}^1(\mathbb{F}_p)$, so one can see that an element of $\mathbb{P}^1_{\mathfrak{q}}$ is periodic under $\tilde{\varphi}_{\mathfrak{q}}$ exactly when the corresponding element of $\mathbb{P}^1_{\mathbb{F}_p}$ is periodic under φ_p. □

We use this relatively elementary lemma to prove another lemma about prime pullbacks. The goal of Lemma 4.7 is to allow us to work in some larger k' than $\mathbb{Q}$ for the proof of Theorem 1.2; if we can show the supremum is 0 there, the result will follow.

LEMMA 4.7 ([JKMT16, Lemma 6.3]). *Let k' be a finite extension of $\mathbb{Q}$, let $\varphi \in \mathbb{Q}[x]$, and let p be a prime of good reduction for φ. If $\tilde{\varphi}$ is the extension of φ to $\mathbb{P}^1_{k'}$, then*

$$\lim_{p\to\infty} \inf \frac{\#\mathrm{Per}(f_p)}{p+1} \leq \lim_{\substack{\text{primes } \mathfrak{q} \text{ of } k' \\ N(\mathfrak{q})\to\infty}} \sup \frac{\#\mathrm{Per}(f_{\mathfrak{q}})}{N(\mathfrak{q})+1}.$$

Proof. This proof relies on the Chebotarev density theorem for number fields in a more standard sense than in the proof of Proposition 4.4, returning to the sense of factoring integer primes into primes over an extension of $\mathbb{Q}$. If $\mathcal{P}$ is the set of primes for which φ has good reduction, then the infimum property in the limit is still upheld restricting to $p \in \mathcal{P}$. Then, the Chebotarev density theorem gives a positive density of primes $\mathfrak{q}$ for which the inequality is satisfied, so transferring from infimum to supremum and returning to the set of all primes of k' gives the desired result. □

Lemma 4.7 finally allows us to "pull back" to an integer prime, allowing us to complete the proof of Theorem 1.2.

Proof of Theorem 1.2. This proof mainly uses Theorem 3.4 to do the heavy lifting, with Corollary 4.5 then getting the desired result over E, the algebraic closure of $\mathbb{Q}$ in K_1. Finally, an application of Lemma 4.7 allows us to retract to get the

desired result over $\mathbb{Q}$ (and, in fact, any number field by generalizing the statement of Lemma 4.7). □

Moving to the third main theorem, the authors zoomed out beyond the quadratic case for their proof, showing a similar result for all polynomials conjugate to a polynomial of the form $x^d + c$ for some $d > 1$. Since all quadratic polynomials are conjugate to this, these results were sufficient for that case. Their first result uses the claim that 0, the only critical point of $x^d + c$, is not preperiodic, allowing applications of their previous theorems.

PROPOSITION 4.8 ([JKMT16, Proposition 6.4]). *Let $d > 1$ and let $f(x) = x^d + c \in \mathbb{Q}[x]$ be such that 0 is not preperiodic. Then*

(1)
$$\liminf_{p\to\infty} \frac{\#\mathrm{Per}(f_p)}{p+1} = 0;$$

(2) if $d = 2$, we have
$$\lim_{p\to\infty} \frac{\#\mathrm{Per}(f_p)}{p+1} = 0.$$

Proof. Let $k' = \mathbb{Q}(\zeta_d)$ for ζ_d a dth root of unity and let $\tilde{f}$ be the extension of f to $\mathbb{P}^1$. This means the splitting field of $\tilde{f}(x) - t$ over $\mathbb{Q}(t)$ is $k'(t)(\sqrt[d]{t-c})$ since these are exactly the roots of $\tilde{f}$ over $\mathbb{Q}(t)$. The Galois group is then the cyclic group C_d as a homomorphism on the splitting field will be uniquely determined by the image of $\zeta_d \sqrt[d]{t-c}$. The only critical point of f is 0, so since it is not preperiodic by assumption, we can apply Theorem 3.4 for all N and thus $\mathrm{Gal}((\tilde{f}^n(x) - t)/k'(t)) \cong [C_d]^n$. Applying Corollary 4.5 and [Odo85, Lemma 4.3] then gives

$$\lim_{\substack{\text{primes } \mathfrak{q} \text{ of } k' \\ N(\mathfrak{q})\to\infty}} \frac{\#\mathrm{Per}(\tilde{f}_{\mathfrak{q}})}{N(\mathfrak{q})+1} = 0.$$

Then, applying Lemma 4.7 finishes the first part. In the second part, we have $k' = \mathbb{Q}$, so this follows directly from Corollary 4.5. □

With the above proposition, the only barrier to proving Theorem 1.3 is to compute where 0 is a preperiodic point. Juul Kurlberg Madhu Tucker [JKMT16] handled that case in the following theorem.

THEOREM 4.9 ([JKMT16, Theorem 6.5]). *Let $d > 1$ and let $f(x) = x^d + c \in \mathbb{Q}[x]$. Then*
$$\liminf_{p\to\infty} \frac{\#\mathrm{Per}(f_p)}{p+1} = 0$$
unless f is the Chebyshev polynomial $x^2 - 2$.

Proof. If 0 is not preperiodic in f, Proposition 4.8 immediately solves the problem. Otherwise, every critical point of f is preperiodic; this type of function is called *post-critically finite*. They then used a result of Jones [Jon15], which used an unrelated technique on monodromy groups, to show that unless f is conjugate to plus or minus a Chebyshev polynomial of degree d, $\lim_{n\to\infty} \mathrm{FPP}(\mathrm{Gal}((f^n(x) - t)/\mathbb{C}(t)) = 0$. The Chebyshev polynomials T_d are such that $T_d(x + \frac{1}{x}) = x^d + \frac{1}{x^d}$; therefore, the simplest such example is $T_2 = x^2 - 2$. Checking which Chebyshev polynomials can be conjugate to $x^d + c$ is relatively easy as the derivative would have to be a $(d-1)$st power, but computations on coefficients easily refute this, meaning $x^2 - 2$

is the only possibility as it is the only conjugate in this form of either plus or minus a Chebyshev polynomial of degree 2. Let k_n be the algebraic closure of $\mathbb{Q}$ in the splitting field of $f^n(x) - t$ over $\mathbb{Q}(t)$. In the other case, they use the Chebotarev density theorem to find a positive-density set of primes p such that the prime ideal $p\mathfrak{o}_{k_n}$ factors into a product of primes $\mathfrak{q}$ with $\mathfrak{o}_{k_n}/\mathfrak{q} = \mathbb{F}_p$. Then, applying Proposition 4.4 to find sufficiently large $\mathfrak{q}$ and then Lemma 4.6 to convert back to p gives the result. □

With this, we can complete the proof of the final main theorem and classify the density of periodic points of quadratic polynomials in all but one conjugacy class.

Proof of Theorem 1.3. As every quadratic polynomial is conjugate to a polynomial of the form $x^2 + c$ by a linear polynomial $\sigma = ax + b$, for all $p \nmid a$, the number of periodic points of f and $\sigma^{-1} f \sigma$ is the same mod p. Then, Theorem 4.9 immediately finishes the proof. □

References

[CH12] John Cullinan and Farshid Hajir. Ramification in iterated towers for rational functions. *Manuscripta Math.*, 137(3-4):273–286, 2012.

[JKMT16] Jamie Juul, Pär Kurlberg, Kalyani Madhu, and Tom J. Tucker. Wreath products and proportions of periodic points. *Int. Math. Res. Not. IMRN*, (13):3944–3969, 2016.

[Jon15] Rafe Jones. Fixed-point-free elements of iterated monodromy groups. *Trans. Amer. Math. Soc.*, 367(3):2023–2049, 2015.

[MS94] Vijaya Kumar Murty and John Scherk. Effective versions of the Chebotarev density theorem for function fields. *C. R. Acad. Sci. Paris Sér. I Math.*, 319(6):523–528, 1994.

[Odo85] R. W. K. Odoni. The Galois theory of iterates and composites of polynomials. *Proc. London Math. Soc. (3)*, 51(3):385–414, 1985.

[Pol75] J. M. Pollard. A Monte Carlo method for factorization. *Nordisk Tidskr. Informationsbehandling (BIT)*, 15(3):331–334, 1975.

Massachusetts Institute of Technology
Email address: lucyeps@mit.edu