ALGEBRAIC RELATIONS AMONGST PERIODIC POINTS OF A LIFTING OF FROBENIUS

THOMAS SCANLON

ABSTRACT. Let p be a prime number and $f \in \mathbb{Z}_p[x]$ a polynomial over the p-adic integers lifting the Frobenius in the sense that $f(x) \equiv x^p \pmod{p}$ and $\deg(f) = p$. Let $\Pi_f := \{\zeta \in \mathbb{Q}_p^{\mathrm{alg}} : f^{\circ n}(\zeta) = \zeta \text{ for some } n \in \mathbb{Z}_+\}$ be the set of f-periodic points. We show that an irreducible algebraic varieties $X \subseteq \mathbb{A}^m$ which meet $\Gamma_f^{\times m}$ in a Zariski dense set only in the case that X is defined by equations of the form $x_i = a$ for some f-periodic point a and $x_i = f^{\circ n}(x_j)$ for some natural number n.

1. INTRODUCTION

Recall that for $f : X \to X$ a function from a set back to itself a point $x \in X$ is called *periodic* (respectively, *pre-periodic*) if $f^{\circ n}(x) = x$ for some $n \in \mathbb{Z}_+$ (respectively, $f^{\circ(n+m)}(x) = f^{\circ m}(x)$ for some $n, m \in \mathbb{Z}_+$). When X is an abelian variety and $f : X \to X$ is given by multiplication by some integer $m \geq 2$, then f-periodic points are the torsion points of order prime to m while the pre-periodic points are the torsion points without qualification. The Manin-Mumford conjecture (Raynaud's theorem) asserts that in this case an irreducible subvariety of X contains a dense set of periodic points just in case X is a translate of an abelian subvariety by a torsion point. It is not hard to see that such varieties are precisely the f-periodic varieties.

Zhang has conjectured that the same holds for more general rational dynamical systems [9]. Specifically, he considers algebraic varieties X over $\mathbb{C}$ given together with a morphism $f : X \to X$ which is polarizable in the sense that there is a line bundle $\mathcal{L}$ on X for which $f^*\mathcal{L} \approx \mathcal{L}^q$ for some $q > 1$ and conjectures that an irreducible subvariety $Y \subseteq X$ contains a Zariski dense set of f-(pre-)periodic points just in case Y is itself (pre-)periodic.

In this note we study a class of rational dynamical systems for which it is fairly easy to verify that Zhang's conjecture for the periodic points holds but unlike in the case of abelian varieties substantial work is required to thoroughly describe the periodic varieties. We consider the case that $X = \mathbb{A}^n$ is affine space over some p-adic field K and $f : X \to X$ is given as $(x_1, \ldots, x_n) \mapsto (g(x_1), \ldots, g(x_n))$ where $g \in \mathbb{Z}_p[x]$ lifts the Frobenius in the sense that $g(x)$ induces the map $x \mapsto x^p$ on the residue field and $\deg(g) = p$. For such functions, except in the case that $g(x)$ is a linear transform of either the monomial x^p itself or the p^{th} Chebyshev polynomial, then an irreducible variety contains a dense set of f-periodic points only if it is

Partially supported by NSF CAREER grant DMS-0450010. The main ideas and results from this document will be included in a joint paper with Alice Medvedev currently in progress and will not be published independently.

defined by equations of the form $x_i = f^{\circ n}(x_j)$ for $n \geq 0$ and $x_i = a$ for a an f-periodic point.

Our method of proof applies well beyond the simple case studied here, but we shall return to this issue in a later paper.

2. Notation and conventions

In this section we establish the notation and basic definations to be used throughout the rest of this paper.

Recall that a *difference field* is a field L given together with a distinguished field automorphism $\sigma : L \to L$. If $f : X \to Y$ is a morphism of varieties over L, then $f^\sigma : X^\sigma \to Y^\sigma$ is the σ-transform of f from the σ-transform X to the σ-transform. Concretely, the σ-transform is obtained by applying σ to the coëfficients of the defining equations of the relevant objects. If one likes, X^σ is the fibre product of X with $\operatorname{Spec}(L)$ over $\operatorname{Spec}(L)$ via the map $\sigma^* : \operatorname{Spec}(L) \to \operatorname{Spec}(L)$.

If $f : X \to X$ is a morphism (in any category) and $n \in \mathbb{N}$ is a natural number, then we define $f^{\circ n} : X \to X$, the n^{th} iterate of f by recursion with $f^{\circ 0} := \mathrm{id}_X$ and $f^{\circ(n+1)} := f \circ f^{\circ n}$. If $f : X \to X^\sigma$, then we define the n^{th} *skew iterate* of f, $f^{\diamond n}$ by $f^{\diamond 0} := \mathrm{id}_X$ and $f^{\diamond(n+1)} := f^{\sigma^n} \circ f^{\diamond n} : X \to X^{\sigma^{n+1}}$. In the special case that $f = f^\sigma$ we have $f^{\diamond n} = f^{\circ n}$. If we wish to consider the coördinate-wise action of f on the n^{th} Cartesian power of X we write $f^{\times n} : X^n \to (X^\sigma)^n$.

In what follows, (K, v) is a valued field with valuation ring $R = \mathcal{O}_{K,v} := \{x \in K : v(x) \geq 0\}$ having the maximal ideal $\mathfrak{m} = \mathfrak{m}_{K,v} := \{x \in K : v(x) > 0\}$ and algebraically closed residue field $k = R/\mathfrak{m}$ of characteristic $p > 0$. We assume that K is maximally complete of characteristic zero. Maximal completeness is not a serious restriction as most of the results we prove in this context could be deduced for more general fields by first passing to a maximal completion and then restricting. Much of what we say remains valid in positive characteristic, but there are nontrivial complications involving additive polynomials. We fix an automorphism $\rho : K \to K$ which preserves the valuation in the sense that $v(\rho(x)) = v(x)$ for every $x \in K^\times$ and which lifts the Frobenius in the sense that $\rho(x) \equiv x^p \mod \mathfrak{m}$ for $x \in R$.

The reader may wish to specialize to the case that $k = \mathbb{F}_p^{\text{alg}}$, $K = \widehat{\mathbb{Q}_p^{\text{unr}}}$, the completion of the maximal unramified extension of the p-adics, and ρ being the unique lifting of the Frobenius automorphism to K or perhaps to take for K the maximal completion of the algebraic closure of $\mathbb{Q}_p$.

For a polynomial $f \in R[x]$ we say that f *lifts the Frobenius* if $f(x) \equiv x^p \mod \mathfrak{m}$ and that f is a *good lifting of the Frobenius* if in addition $\deg(f) = p$. One might wish to generalize this notion to say that a morphism $f : X \to Y$ of schemes over R *lifts the Frobenius* if on the special fibre $Y_k = X_k^{(p)}$ and $f_k = F_{X_k}$, the Frobenius morphism on X_k and that f is a *good lifting of the Frobenius* if in addition the map f is finite and $\deg(f) = \deg(F_{X_k})$.

3. σ-varieties

In this section we recall some of the formalism of σ-varieties (see [7, 6]) and express some of the results from the model theory of difference fields (see [3, 4]) in terms of σ-varieties.

In this section, we fix a field L with a distinguished automorphism $\sigma : L \to L$. In our applications, $L = K$ and $\rho = \sigma$. Moreover, we work over L so that, for example, the word *variety* means *variety over* L.

Definition 3.1. A *σ-variety* is a pair (X, f) where X is a variety and $f : X \to X^\sigma$ is a morphism of varieties from X to its σ-transform. A morphism of σ-varieties $\gamma : (X, f) \to (Y, g)$ is given by a morphism of varieties $\gamma : X \to Y$ for which the following square commutes.

$$\begin{array}{ccc} X & \xrightarrow{f} & X^\sigma \\ {\scriptstyle\gamma}\downarrow & & \downarrow{\scriptstyle\gamma^\sigma} \\ Y & \xrightarrow{g} & Y^\sigma \end{array}$$

If (X, f) is a σ-variety and $Y \subseteq X$ is a subvariety, we say that Y is a *σ-subvariety* if $f(Y) \subseteq Y^\sigma$. That is, if the restriction of f to Y gives Y the structure of a σ-variety and the inclusion map $Y \hookrightarrow X$ is a map of σ-varieties.

If X is a variety over L, then $\sigma : L \to L$ induces a function which we shall continue to denote by σ from the L-rational points of X to the L-rational points of X^σ. Given a σ-variety (X, f) we may consider its set of L-rational points $(X, f)^\sharp(L) = (X, f)^\sharp(L, \sigma) := \{a \in X(L) : f(a) = \sigma(a)\}$. If $Y \subseteq X$ is a subvariety for which $Y(L) \cap (X, f)^\sharp(L)$ is Zariski dense in Y, then Y is a σ-subvariety. (The converse is true only under the hypothesis that (L, σ) is a difference closed field, or in the language of model theory, a model of ACFA.)

Two contradictory versions of triviality for σ-varieties appear in the literature. Sometimes, one says that a σ-variety of the form (X, id_X) is trivial and a σ-variety which admits a finite-to-finite correspondence (possibly after base change) with such a σ-variety is isotrivial. Here, in saying that (X, f) and (Y, g) are in *finite-to-finite correspondence* we mean that there is a third σ-variety (Z, h) and generically finite dominant morphisms $(Z, h) \to (X, f)$ and $(Z, h) \to (Y, g)$.

For us, these so-called "trivial" or "iso-trivial" σ-varieties are as far from being trivial as possible as their Cartesian powers have very rich families of σ-subvarieties. Indeed, if X is defined over the fixed field of σ and $Y \subseteq X^n$ is any subvariety of the n^{th} Cartesian power of X also defined over the fixed field of σ, then Y is a σ-subvariety of (X^n, id_{X^n}). In some sense, most σ-varieties have very few σ-subvarieties. So, we shall refer to a difference variety of the form (X, id_X) as *constant* and those which are in finite-to-finite correspondence with constant σ-varieties as *iso-constant*.

It can happen that a σ-variety (X, f) is not iso-constant, but its structure still comes from a fixed field. For example, if L has characteristic p, then the σ-variety $(\mathbb{A}^1, x \mapsto x^p)$ is not iso-constant, but $(\mathbb{A}^1, x \mapsto x^p)^\sharp(L, \sigma) = (\mathbb{A}^1, \mathrm{id}_{\mathbb{A}^1})^\sharp(L, \eta)$ where $\eta(x) := \sigma(x^{\frac{1}{p}})$. We refer to a σ-variety which becomes isoconstant after σ is replaced by an automorphism of the form $\sigma^n\tau^m$ for some $n \in \mathbb{Z}_+$ and $m \in \mathbb{Z}$ where $\tau(x) = x^p$ as *weakly iso-constant*. We apologize for the imprecision of this definition, but as weakly iso-constant σ-varieties do not play any rôle in this paper, we do not pursue the matter.

The notion of the triviality of forking was isolated in the study of stable theories though it has meaning well beyond that context. Even in the context of σ-varieties one ought to consider more refined versions of triviality than what we describe with the next definition, but this suffices for the purposes of this paper.

Definition 3.2. We say that the σ-variety (X, f) is *trivial* if for each natural number n every irreducible σ-subvariety $Z \subseteq X^n$ (over any difference field extension

of (L,σ) is of the form $\bigcap_{1\leq i<j\leq n}\pi_{i,j}^{-1}Y_{i,j}$ where $Y_{i,j}\subseteq X^2$ is a σ-subvariety of $(X^2,f^{\times 2})$ and $\pi_{i,j}:X^n\to X^2$ is the coördinate projection $(x_1,\dots,x_n)\mapsto(x_i,x_j)$.

While constant σ-varieties give an extreme version of nontrivial σ-varieties, an intermediate class of σ-varieties is provided by *modular groups*. If G is a positive dimensional connected commutative algebraic group over L and $\psi:G\to G^\sigma$ is an isogeny, then (G,ψ) is a nontrivial σ-variety, as, for instance, the graph of addition is a σ-subvariety of G^3 which cannot be obtained as an intersection of pullbacks of subvarieties of G^2. However, whenever the separable degree of ψ is greater than one, (G,ψ) is not iso-constant. In fact, when $\dim G=1$ and the separable degree of ψ is greater than one, then every irreducible σ-subvariety of any Cartesian power of G is a translate of an algebraic subgroup (see [2]). Moreover, there are only countably many σ-subgroup varieties. Such σ-algebraic groups are called *modular*.

Modular σ-varieties may leave a trace on σ-varieties which do not carry a group structure. If (G,ψ) is a σ-variety for which $\psi:G\to G^\sigma$ is an isogeny of algebraic groups and $\Gamma\leq G\rtimes\operatorname{Aut}(G)$ is a finite group of affine automorphisms of G for which for every $g\in G(L^{\mathrm{alg}})$ and $\gamma\in\Gamma$ there is some $\gamma'\in\Gamma^\sigma$ for which $\psi(\gamma g)=\gamma'\psi(g)$, then ψ descends to a rational map $\overline{\psi}:\Gamma\backslash G\to(\Gamma\backslash G)^\sigma$. If (G,ψ) is modular, then the irreducible σ-subvarieties of the Cartesian powers of $(\Gamma\backslash G,\overline{\psi})$ are simply the quotients by Γ of the translates of certain algebraic subgroups.

The three classes of isoconstant, modular group, and trivial σ-varieties are mutually exclusive even to the point that there cannot be finite-to-finite correspondences between σ-varieties from different classes. The main theorems of [3] and [4] give a classification of "minimal sets" defined by difference equations in terms of these three classes. When specialized to the case of σ-varieties of the form (C,f) where C is a curve, their trichotomy theorem says precisely that such a σ-variety must be isoconstant, trivial, or in finite-to-finite correspondence with a modular group. The main theorem of Medvedev's doctoral thesis [5] gives explicit criteria for determining into which of these three classes a given σ-variety structure on a curve belongs.

One sees easily that if C is a curve and $f:C\to C^\sigma$ has separable degree one, then (C,f) is weakly iso-constant. In particular, if C is a curve of genus greater than one, then every σ-variety structure on C is weakly isotrivial. If C has genus one and $f:C\to C^\sigma$ has separable degree greater than one, then as we have already discussed (C,f) is a modular group. As the trichotomy is insensitive to the choice of points on the curve, this leaves the case of analyzing σ-varieties of the form $(\mathbb{P}^1,f)$.

Theorem 3.3 (Medvedev)**.** *Let $f:\mathbb{P}^1\to\mathbb{P}^1$ be a nonconstant rational function. Unless there is a one-dimensional algebraic group G, an isogeny $\psi:G\to G^\sigma$ and a nonconstant map of τ-varieties $\gamma:(G,\psi)\to(\mathbb{P}^1,g)$ representing $\mathbb{P}^1$ (generically) as a quotient $\Gamma\backslash G$ of G by a finite group of affine automorphisms, the σ-variety $(\mathbb{P}^1,f)$ is trivial. Moreover, if the characteristic of L is zero and f is a polynomial, then $(\mathbb{P}^1,g)$ is only non-trivial in the cases where g is a linear transform of a monomial or a Chebyshev polynomial, a nonconstant polynomial $T_n(x)$ for which $T_n(x+\frac{1}{x})=x^n+\frac{1}{x^n}$ for some integer $n\geq 2$.*

Knowing that the σ-variety $(\mathbb{P}^1,f)$ is trivial allows one to describe σ-subvarieties of $((\mathbb{P}^1)^n,f^{\times n})$ in terms of curves, but it leaves open the problem of describing σ-subvarieties of $\mathbb{P}^1\times\mathbb{P}^1$. An old theorem of Ritt about the uniqueness of the

decomposition of a polynomial as a compositional product of other polynomials completes the picture [8].

If $f(x) \in L[x]$ is a polynomial and we write $f = f_1 \circ \cdots \circ f_n$ as a composition of polynomials all of which are compositionally indecomposable, then there are four basic ways to transform this decomposition into another decomposition.

(1) If $\lambda(x) = ax + b$ is a linear polynomial with $a \neq 0$, we may replace f_i with $f_i \circ \lambda$ and f_{i+1} with $\lambda^{-1} \circ f_{i+1}$.
(2) If $f_i(x) = x^n$ and $f_{i+1}(x) = x^m$ for integers n and m, we may interchange f_i and f_{i+1}.
(3) If f_i and f_{i+1} are both Chebyshev polynomials, we may interchange them.
(4) if $f_i(x) = x^n$ and $f_{i+1}(x) = x^m g(x^n)$ for some polynomial g, then we may replace f_i with $x^m g(x)^n$ and f_{i+1} with x^n (and we may reverse this process).

Ritt showed that when L has characteristic zero (to be honest, when $L = \mathbb{C}$, but the general characteristic zero case follows immediately) that if f has two different compositional decompositions, then the first may be moved to the second by a finite sequence of moves of the above form [8]. Using Ritt's theorem, we can completely describe the possible maps of σ-varieties between $(\mathbb{A}^1, f)$ and $(\mathbb{A}^1, g)$.

In what follows, we say that the polynomial $g(x) \in L[x]$ is a *linear transform* of the polynomial $f(x) \in L[x]$ if there is a linear polynomials $\lambda \in L[x]$ for which $g = \lambda^\sigma \circ f \circ \lambda^{-1}$.

Proposition 3.4. *Suppose that the characteristic of L is zero and that $g(x) \in L[x] \setminus L$ is an indecomposable nonlinear polynomial which is not a linear transform of a monomial, a Chebshev polynomial, or a polynomial of the form $x^n h(x)^m$ or $x^n h(x^m)$ for positive integers m and n and a polynomial h. If $f(x) \in L[x] \setminus L$ is any other nonconstant polynomial and $\gamma : (\mathbb{A}^1, f) \to (\mathbb{A}^1, g)$ is a nonconstant map of σ-varieties, then g is a linear transform of f^{σ^n} for some n and γ is a linear transform of $f^{\Diamond n}$. In fact, there is a linear μ for which $g = \mu^\sigma f^{\sigma^n} \mu^{-1}$ and $\gamma = \mu \circ f^{\Diamond n}$. Likewise, the same conclusion holds if we assume that f is an indecomposable nonlinear polynomial which is not a linear transform of one of Ritt's special polynomials.*

Proof. If γ is linear, then take $\mu = \gamma$ and $n = 0$. Write γ as $\gamma = \gamma_1 \circ \cdots \circ \gamma_n$ where each γ_i is indecomposable. We have the functional equation $\gamma_1^\sigma \circ \cdots \gamma_n^\sigma \circ f = \gamma^\sigma \circ f = g \circ \gamma = g \circ \gamma_1 \circ \cdots \circ \gamma_n$. It follows immediately from Ritt's theorem on the uniqueness of the number of indecomposable compositional factors that f is also indecomposable. Using Ritt's theorem again, Working by induction on $i \leq n$ we see that the only allowable transformations from the decomposition $\gamma_1^\sigma \circ \cdots \circ \gamma_n^\sigma \circ f$ to the decomposition $g \circ \gamma_1 \circ \cdots \circ \gamma_n$ are of the first kind, namely, the insertion of linear transformations. That is there are linear polynomials λ_i for which $g \circ \lambda_1 = \gamma_1^\sigma$, $\lambda_i^{-1} \circ \gamma_i \circ \lambda_{i+1}$ for $i < n$, and $f = \lambda_n^{-1} \circ \gamma_n$. Solving iteratively, we have $\gamma_{n-i} = \lambda_{n-i} \circ \lambda_{n+1-i}^\sigma \circ \cdots \circ \lambda_n^{\sigma^i} \circ f^{\sigma^i} \circ \lambda_n^{\sigma^{i-1}} \circ \cdots \circ \lambda_{n+1-i}^{-1}$ and $g = \mu^\sigma f^{\sigma^n} \mu^{-1}$ where $\mu = \lambda_1 \circ \lambda_2^\sigma \circ \cdots \circ \lambda_n^{\sigma^{n-1}}$.

The "likewise" clause is proven in precisely the same way. □

4. Main theorem

In this section we prove our main theorem on periodic points for a lifting of the Frobenius.

Proposition 4.1. *If $f \in R[x]$ lifts the Frobenius and $n \in \mathbb{Z}_+$ is any positive integer, then the reduction map establishes a bijection between $\{a \in R : \rho^n(a) = f^{\Diamond n}(a)\}$ and the residue field k.*

Proof. This follows immediately from the ρ-Hensel's Lemma of [1]. As the argument is elementary, we repeat it here.

Let us check injectivity first. If $\rho^n(a) = f^{\Diamond n}(a)$ and $\epsilon \in \mathfrak{m}$, then we have $\rho^n(a+\epsilon) - f^{\Diamond n}(a+\epsilon) = \rho^n(\epsilon) - (f^{\Diamond n})'(a)\epsilon + c\epsilon^2$ for some $c \in R$. As $f^{\Diamond n}(x) \equiv x^{p^n} \mod \mathfrak{m}$, we see that $v((f^{\Diamond n})'(a)) > 0$. Hence, the ultrametric inequality yields $v(\rho^n(a+\epsilon) - f^{\Diamond n}(a+\epsilon)) = v(\epsilon)$. So that we must have $\epsilon = 0$ if we wish for $a+\epsilon$ to be a solution to the equation.

For sujectivity, we apply Newton's method to the difference polynomial $Q(x) := \rho^n(x) - f^{\Diamond n}(x)$. Let $\bar{a} \in k$ be any element of the residue field. For any $a \in R$ lifting $\bar{a}$ we have $Q(a) \equiv 0 \mod \mathfrak{m}$. By maximal completeness of K, it suffices to show that if a is not already a root of Q, we can find another b lifting $\bar{a}$ with $v(Q(b)) > v(Q(a))$. Let $\epsilon := Q(a)$, which as we have observed is in $\mathfrak{m}$. Then for any $Y \in R^\times$ we have $Q(a+\epsilon Y) = Q(a) + \rho(\epsilon)\rho(Y) + b$ for $c \in R$ with $v(c) > v(\epsilon)$. As k is perfect, we can find some $Y \in R^\times$ for which $Y^{p^n} \equiv -\frac{Q(a)}{\rho(\epsilon)} \mod \mathfrak{m}$. Set $b := a + \epsilon Y$. □

Corollary 4.2. *If $f(x) \in R[x]$ lifts the Frobenius, $n \in \mathbb{Z}_+$ is any positive integer and $a \in R$, then $\rho^n(a) = f^{\Diamond n}(a)$ if and only if $\rho(a) = f(a)$.*

Proof. If $\rho(a) = f(a)$ then one sees by induction that $\rho^n(a) = f^{\Diamond n}(a)$. Indeed, $\rho(a) = f(a)$ by hypothesis, and $\rho^{n+1}(a) = \rho(\rho^n(a)) = \rho(f^{\Diamond n}(a)) = (f^{\Diamond n})^\rho(\rho(a)) = f^{\rho^n} \circ \cdots \circ f^\rho(f(a)) = f^{\Diamond(n+1)}(a)$.

Conversely, suppose that $a \in R$ and that $\rho^n(a) = f^{\Diamond n}(a)$. By Proposition 4.1, there is a unique solution b to $\rho(x) = f(x)$ and $x \equiv a \mod \mathfrak{m}$. By the previous paragraph, b is also a solution to $\rho^n(x) = f^{\Diamond n}(x)$ and $x \equiv a \mod \mathfrak{m}$ to which a is the unique solution. Hence, $b = a$. □

Remark 4.3. If we assume in addition that $f \in R[x]$ is a *good* lifting of the Frobenius, then every solution to $\rho^n(x) = f^{\Diamond n}(x)$ in K actually belongs to R. Indeed, if $v(x) < 0$, then $v(f^{\Diamond n}(x)) = p^n v(x) < v(x) = v(\rho^n(x))$.

Proposition 4.4. *If $f(x) \in R[x]$ lifts the Frobenius and $f = f^\rho$, then any f-periodic point in R is a solution to $\rho(x) = f(x)$.*

Proof. Let $\zeta \in R$ be an f-periodic point of order $n \in \mathbb{Z}_+$ meaning that $f^{\circ n}(\zeta) = \zeta$. As ζ is a root of the nontrivial polynomial $f^{\circ n}(X) - X$ which has coëfficients in the fixed field of ρ, ζ itself is fixed by ρ^m for some $m \in \mathbb{Z}_+$. So, $f^{\Diamond mn}(\zeta) = f^{\circ mn}(\zeta) = \zeta = \rho^{mn}(\zeta)$. By Corollary 4.2, $f(\zeta) = \rho(\zeta)$. □

Corollary 4.5. *If $f(x) \in R[x]$ lifts the Frobenius, $f = f^\rho$, and $Z \subseteq \mathbb{A}^m$ is an irreducible variety containing a dense set of R-rational $f^{\times m}$-periodic points, then Z is $f^{\times m}$-periodic.*

Proof. By Proposition 4.4, Z contains a dense set of points satisfying the difference equation $\rho(x) = f^{\times m}(x)$. Hence, $f^{\times m}(Z) = Z^\rho$. As Z contains a dense set of points algebraic over the fixed field of ρ, Z itself is fixed by some power of ρ so that $(f^{\circ n})^{\times m}(Z) = Z$ for some $n \in \mathbb{Z}_+$. □

The Medvedev's Theorem 3.3 applies directly to describe the ρ-variety $(\mathbb{A}^1, f)$ where $f \in R[x]$ is a lifting of the Frobenius.

Proposition 4.6. *If $f(x) \in R[x]$ is a lifting of the Frobenius, then the difference variety $(\mathbb{A}^1, f)$ is trivial except in the cases that f is a linear transform of the monomial x^p or of the p^{th} Chebyshev polynomial.*

Remark 4.7. It bears noting that we have a complete description of the ρ-subvarieties of $(\mathbb{A}^n, f^{\times n})$ in the excluded cases as well. Applying the linear transformation, we may take f to be x^p or the p^{th} Chebyschev polynomial. In the former case, after removing 0 from $\mathbb{A}^1$, every irreducible ρ-subvariety of a Cartesian power of $(\mathbb{G}_m, x \mapsto x^p)$ is a multiplicative translate of a an algebraic subgroup. The difference variety $(\mathbb{A}^1, C_p)$, where C_p is the p^{th} Chebyschev polynomial, is the image of the difference variety $(\mathbb{G}_m, x \mapsto x^p)$ under the morphism $x \mapsto x + \frac{1}{x}$ and the irreducible ρ-subvarieties for the Chebyshev polynomial are precisely the images of the translates of algebaic tori.

Once we know that a ρ-variety is trivial to describe the possible ρ-subvarieties of all Cartesian powers it suffices to describe those in the second Cartesian power.

Proposition 4.8. *If $f(x) \in R[x]$ lifts the Frobenius but is not a linear transform of x^p or of the p^{th} Chebyshev polynomial and $Z \subseteq \mathbb{A}^2$ is an irreducible curve which is a ρ-subvariety of $(\mathbb{A}^2, f^{\times 2})$ which projects dominantly onto both coördinates, then Z or its converse is a graph of some skew iterate of f.*

Proof. Let $g := f^{\times 2} \restriction Z$ be the restriction of $f^{\times 2}$ to Z. By hypothesis, $g(Z) \subseteq Z^\rho$. As each of the projection maps restricted to Z give dominant finite-to-one maps of ρ-varieties from (Z, g) to $(\mathbb{A}^1, f)$ and the latter ρ-variety is trivial, (Z, g) must itself be trivial. Hence, by the discussion preceeding Theorem 3.3, Z is a genus zero curve. Hence, we have a rational function $h : \mathbb{P}^1 \to \mathbb{P}^1$ and a degree one map of ρ-varieties $\gamma : (\mathbb{P}^1, h) \to (\overline{Z}, g)$ where $\overline{Z}$ is the closure of Z in $\mathbb{P}^1 \times \mathbb{P}^1 \supseteq \mathbb{A}^2$. Moreover, we may assume that $\gamma^{-1}(Z) \subseteq \mathbb{A}^1$. Considering ramification at ∞, we see that h and γ are both given by polynomials. Let $\delta_i = \pi_i \circ \gamma$ where $\pi_i : \overline{Z} \to \mathbb{P}^1$ is the projection map to the i^{th} coördinate for $i = 1$ or 2. We have functional equations $f \circ \delta_i = \delta_i^\rho \circ g$. By Proposition 3.4, applied to the equation $f \circ \delta_1 = \delta_1^\rho \circ g$, g is a linear transform of $f^{\sigma^{-n}}$ for some $n \in \mathbb{Z}_+$. applying this linear transform on the domain, we may assume that $g = f^{\sigma^{-n}}$ and then Proposition 3.4 says that $\delta_1 = g^{\Diamond n}$ and that $\delta_2 = g^{\Diamond m}$ for some other $m \in \mathbb{Z}_+$. If $m \geq n$, then Z is the graph of $f^{\Diamond m-n}$ and if $n > m$, then Z is the converse relation of the graph of $f^{\Diamond n-m}$. □

Putting these results together we have our main theorem.

Theorem 4.9. *If $f(x) \in R[x]$ is a lifting of the Frobenius and $f = f^\rho$ and $X \subseteq \mathbb{A}^n$ is an irreducible variety containing a Zariski dense set of R-rational f-periodic points, then X is defined by equations of the form $x_i = a$ for a an f-periodic point and $x_i = f^{\circ m}(x_j)$.*

Proof. By Proposition 4.1 X is a ρ-subvariety of $(\mathbb{A}^n, f^{\times n})$. By Theorem 3.3, $(\mathbb{A}^1, f)$ is trivial. Hence, X is an intersection of pullbacks by coördinate projections of curves which are ρ-subvarieties of $(\mathbb{A}^2, f^{\times 2})$. By Proposition 4.8, these are just horizontal and vertical lines and graphs or converse graphs of iterates of f. □

References

[1] L. Bélair, A. Macintyre, and T. Scanlon, Model theory of the Frobenius on the Witt vectors, *Amer. J. of Math.*, to appear.
[2] Z. Chatzidakis, Groups definable in ACFA. Algebraic model theory (Toronto, ON, 1996), 25–52, NATO Adv. Sci. Inst. Ser. C Math. Phys. Sci., 496, Kluwer Acad. Publ., Dordrecht, 1997.
[3] Z. Chatzidakis and E. Hrushovski, Model theory of difference fields, *Trans. AMS* **351** (1999), no. 8, 2997–3071.
[4] Z. Chatzidakis, E. Hrushovski and Y. Peterzil, Model theory of difference fields. II. Periodic ideals and the trichotomy in all characteristics, *Proc. LMS* (3) **85** (2002), no. 2, 257–311.
[5] A. Medvedev, Group-like minimal sets in ACFA, Ph.D. thesis, University of California, Berkeley, May 2007.
[6] A. Pillay and M. Ziegler, Jet spaces of varieties over differential and difference fields, *Selecta Math. (N.S.)* **9** (2003), no. 4, 579–599.
[7] R. Pink and D. Rössler, On ψ-invariant subvarieties of semiabelian varieties and the Manin-Mumford conjecture, *J. Algebraic Geom.* **13** (2004), no. 4, 771 – 798.
[8] J. F. Ritt, Prime and composite polynomials, *Trans. AMS* **23** (1922), no. 1, 51–66.
[9] S. Zhang, Distributions in algebraic dynamics, Survey in Differential Geometry **10**, 381-430, International Press, 2006

University of California, Berkeley, Department of Mathematics, Evans Hall, Berkeley, California 94720-3840, USA