3.3 The Uniform Boundedness Conjecture

Northcott's Theorem 3.12 says that a morphism $\phi : \mathbb{P}^N \to \mathbb{P}^N$ has only finitely many K-rational preperiodic points. It is even effective in the sense that we can, in principle, find an explicit constant $C(\phi)$ in terms of the coefficients of ϕ such that every point $P \in \mathrm{PrePer}(\phi)$ satisfies $h(P) \le C(\phi)$. This also allows us to compute an upper bound for $\#\,\mathrm{PrePer}(\phi, \mathbb{P}^N(K))$, but the bound grows extremely rapidly as the coefficients of ϕ become large. A better bound, at least for periodic points, may be derived from the local estimates in Chapter 2 as described in Corollary 2.26. However, even that estimate depends on the coefficients of ϕ, since it is in terms of the two smallest primes for which ϕ has good reduction. The following uniformity conjecture says that there should be a bound for the size of $\mathrm{PrePer}(\phi, \mathbb{P}^N(K))$ that depends in only a minimal fashion on ϕ and K.

Conjecture 3.15. (Morton–Silverman [290]) *Fix integers $d \ge 2$, $N \ge 1$, and $D \ge 1$. There is a constant $C(d, N, D)$ such that for all number fields $K/\mathbb{Q}$ of degree at most D and all finite morphisms $\phi : \mathbb{P}^N \to \mathbb{P}^N$ of degree d defined over K,*

$$\#\,\mathrm{PrePer}(\phi, \mathbb{P}^N(K)) \le C(d, N, D).$$

Remark 3.16. There are many results in the literature giving explicit bounds for the size of the sets $\mathrm{PrePer}(\phi, \mathbb{P}^N(K))$ or $\mathrm{Per}(\phi, \mathbb{P}^N(K))$ in terms of ϕ, especially in the case $N = 1$. Some of these results use global methods, while others use a small prime of good (or at least not too bad) reduction for ϕ. For example, we used local methods in Corollary 2.26 to give a weak bound for $\#\,\mathrm{Per}(\phi, \mathbb{P}^1(K))$. For further results, see [47, 81, 84, 85, 92, 93, 126, 150, 158, 174, 175, 176, 177, 178, 206, 306, 243, 290, 302, 303, 305, 307, 327, 329, 332, 333, 335, 428].

Remark 3.17. Very little is known about Conjecture 3.15. Indeed, it is not known even in the simplest case $(d, N, D) = (2, 1, 1)$, that is, for $\mathbb{Q}$-rational points and degree 2 maps on $\mathbb{P}^1$. Specializing further, if we let $\phi_c : \mathbb{P}^1 \to \mathbb{P}^1$ denote the quadratic map $\phi_c(z) = z^2 + c$, then the conjecture implies that

$$\sup_{c \in \mathbb{Q}} \#\,\mathrm{Per}(\phi_c, \mathbb{P}^1(\mathbb{Q})) < \infty,$$

but the best known upper bounds for $\#\,\mathrm{Per}(\phi_c, \mathbb{P}^1(\mathbb{Q}))$ depend on c.

There are one-parameter families of c-values for which $\phi_c(z)$ has a $\mathbb{Q}$-rational periodic point of exact period 1, 2, or 3, see Exercise 3.9 and Example 4.9, and one can show that ϕ_c cannot have $\mathbb{Q}$-rational periodic points of exact period 4 or 5, see [158, 287]. Poonen has conjectured that ϕ_c cannot have any $\mathbb{Q}$-rational periodic points of period greater than 3. Assuming this conjecture, he gives a complete description of all possible rational preperiodic structures for ϕ_c, see [335].

Remark 3.18. Another interesting collection of rational maps is the family

$$\phi_{a,b}(z) = az + \frac{b}{z}.$$

These maps have the symmetry property $\phi_{a,b}(-z) = -\phi_{a,b}(z)$, i.e., conjugation by the map $f(z) = -z$ leaves them invariant. It is known that there are one-parameter families of these maps with a $\mathbb{Q}$-rational periodic point of exact period 1 (in addition to the obvious fixed point at ∞), 2, or 4, and that none of the maps $\phi_{a,b}(z)$ has a $\mathbb{Q}$-rational periodic point of exact period 3. See [264] for details, and Examples 4.69 and 4.71 and Exercises 4.1, 4.40 and 4.41 for additional properties of these maps.

Remark 3.19. Conjecture 3.15 is an extremely strong uniformity conjecture. For example, if we consider only maps $\phi : \mathbb{P}^1 \to \mathbb{P}^1$ of degree 4 defined over $\mathbb{Q}$, then the assertion that $\#\operatorname{PrePer}(\phi, \mathbb{P}^1(\mathbb{Q})) \le C$ for an absolute constant C immediately implies Mazur's theorem [270] that the torsion subgroup of an elliptic curve $E/\mathbb{Q}$ is bounded independently of E. To see this, we observe that Proposition 0.3 tells us that

$$E_{\text{tors}} = \operatorname{PrePer}\bigl([2], E\bigr),$$

and hence the associated Lattès map $\phi_{E,2}$ described in Section 1.6.3 satisfies

$$x\left(E_{\text{tors}}\right) = \operatorname{PrePer}(\phi_{E,2}, \mathbb{P}^1).$$

Note that $\phi_{E,2}$ has degree 4.

In a similar manner, Conjecture 3.15 for maps of degree 4 on $\mathbb{P}^1$ over number fields implies Merel's Theorem [275] that the size of the torsion subgroup of an elliptic curve over a number field is bounded solely in terms of the degree of the number field. Turning this argument around, Merel's theorem implies the uniform boundedness conjecture for Lattès maps, i.e., for rational maps associated to elliptic curves, see Theorem 6.65. Lattès maps are the only nontrivial family of rational maps for which the uniform boundedness conjecture is currently known.

In higher dimension, Fakhruddin [150] has shown that Conjecture 3.15 implies that there is a constant $C(N,D)$ such that if K is a number field of degree at most D and if A/K is an abelian variety of dimension N, then

$$\#A(K)_{\text{tors}} \le C(N,D).$$

He also shows that if Conjecture 3.15 is true over $\mathbb{Q}$, then it is true for all number fields.

3.4 Canonical Heights and Dynamical Systems

It is obvious from the definition of the height that

$$h(\alpha^d) = dh(\alpha) \qquad \text{for all } \alpha \in \bar{\mathbb{Q}}. \tag{3.11}$$

Notice that Theorem 3.11 applied to the particular map $\phi(z) = z^d$ gives the less precise statement

$$h(\phi(P)) = dh(P) + O(1). \tag{3.12}$$

Clearly the exact formula (3.11) is more attractive than the approximation (3.12). It would be nice if we could modify the height h in some way so that the general formula (3.12) from Theorem 3.11 is true without the $O(1)$. It turns out that this can be done for each morphism ϕ. To create these special heights, we follow a construction due originally to Tate.

Theorem 3.20. *Let S be a set, let $d > 1$ be a real number, and let*

$$\phi : S \to S \qquad \textit{and} \qquad h : S \to \mathbb{R}$$

be functions satisfying

$$h\bigl(\phi(P)\bigr) = dh(P) + O(1) \qquad \textit{for all } P \in S.$$

Then the limit

$$\hat{h}(P) = \lim_{n\to\infty} \frac{1}{d^n} h\bigl(\phi^n(P)\bigr) \tag{3.13}$$

exists and satisfies:

(a) $\hat{h}(P) = h(P) + O(1)$.

(b) $\hat{h}\bigl(\phi(P)\bigr) = d\hat{h}(P)$.

The function $\hat{h} : S \to \mathbb{R}$ uniquely determined by the properties (a) *and* (b).

Proof. We prove that the limit (3.13) exists by proving that the sequence is Cauchy. Let $n > m \geq 0$ be integers. We are given that there is a constant C so that

$$\bigl|h(\phi(Q)) - dh(Q)\bigr| \leq C \qquad \text{for all } Q \in S. \tag{3.14}$$

We apply inequality (3.14) with $Q = \phi^{i-1}(P)$ to the telescoping sum

$$\begin{aligned} \left|\frac{1}{d^n}h(\phi^n(P)) - \frac{1}{d^m}h(\phi^m(P))\right| &= \left|\sum_{i=m+1}^{n} \frac{1}{d^i}\Bigl(h(\phi^i(P)) - dh(\phi^{i-1}(P))\Bigr)\right| \\ &\leq \sum_{i=m+1}^{n} \frac{1}{d^i}\bigl|h(\phi^i(P)) - dh(\phi^{i-1}(P))\bigr| \\ &\leq \sum_{i=m+1}^{n} \frac{C}{d^i} \leq \sum_{i=m+1}^{\infty} \frac{C}{d^i} = \frac{C}{(d-1)d^m}. \end{aligned} \tag{3.15}$$

The inequality (3.15) clearly implies that

$$\left|\frac{1}{d^n}h(\phi^n(P)) - \frac{1}{d^m}h(\phi^m(P))\right| \to 0 \qquad \text{as } m, n \to \infty,$$

so the sequence $d^{-n}h(\phi^n(P))$ is Cauchy and the limit (3.13) exists.

In order to prove (a), we take $m = 0$ in (3.15), which yields

$$\left|\frac{1}{d^n}h(\phi^n(P)) - h(P)\right| \leq \frac{C}{d-1}.$$

Next we let $n \to \infty$ to obtain

$$\left|\hat{h}_\phi(P) - h(P)\right| \le \frac{C}{d-1},$$

which is (a) with an explicit value for the $O(1)$ constant.

The proof of (b) is a simple computation using the definition of $\hat{h}$,

$$\hat{h}_\phi(\phi(P)) = \lim_{n\to\infty} \frac{1}{d^n} h\bigl(\phi^n(\phi(P))\bigr) = \lim_{n\to\infty} \frac{d}{d^{n+1}} h\bigl(\phi^{n+1}(P))\bigr) = d\hat{h}_\phi(P).$$

Finally, to prove uniqueness, suppose that $\hat{h}' : S \to \mathbb{R}$ also has properties (a) and (b). Then the difference $g = \hat{h} - \hat{h}'$ satisfies

$$g(P) = O(1) \qquad \text{and} \qquad g(\phi(P)) = dg(P).$$

These formulæ hold for all elements $P \in S$, so

$$d^n g(P) = g(\phi^n(P)) = O(1) \qquad \text{for all } n \ge 0.$$

In other words, the quantity $d^n g(P)$ is bounded as $n \to \infty$, which can only happen if $g(P) = 0$. This proves that $\hat{h}(P) = \hat{h}'(P)$, so $\hat{h}$ is unique. □

Definition. Let $\phi : \mathbb{P}^N \to \mathbb{P}^N$ be a morphism of degree $d \ge 2$. The *canonical height function* (*associated to* ϕ) is the unique function

$$\hat{h}_\phi : \mathbb{P}^N(\bar{\mathbb{Q}}) \longrightarrow \mathbb{R}$$

satisfying

$$\hat{h}_\phi(P) = h(P) + O(1) \qquad \text{and} \qquad \hat{h}_\phi\bigl(\phi(P)\bigr) = d\hat{h}_\phi(P).$$

The existence and uniqueness of $\hat{h}_\phi$ follows from Theorem 3.20 applied to the maps

$$\phi : \mathbb{P}^N(\bar{\mathbb{Q}}) \longrightarrow \mathbb{P}^N(\bar{\mathbb{Q}}) \qquad \text{and} \qquad h : \mathbb{P}^N(\bar{\mathbb{Q}}) \longrightarrow \mathbb{R},$$

since Theorem 3.11 tells us that ϕ and h satisfy

$$h\bigl(\phi(P)\bigr) = dh(P) + O(1) \qquad \text{for all } P \in \mathbb{P}^N(\bar{\mathbb{Q}}).$$

Remark 3.21. The definition $\hat{h}_\phi(P) = \lim_{n\to\infty} d^{-n} h\bigl(\phi^n(P)\bigr)$ is not practical for accurate numerical calculations. Thus even for $P \in \mathbb{P}^1(\mathbb{Q})$, one would need to compute the exact value of $\phi^n(P)$ whose coordinates have $O(d^n)$ digits. A practical method for the numerical computation of $\hat{h}_\phi(P)$ to high accuracy uses the decomposition of $\hat{h}_\phi$ as a sum of local heights or Green functions. This decomposition is described in Sections 3.5 and 5.9. See in particular Exercise 5.29 for a detailed description of the algorithm.

The canonical height provides a useful arithmetic characterization of the preperiodic points of ϕ.

Theorem 3.22. *Let $\phi : \mathbb{P}^N \to \mathbb{P}^N$ be a morphism of degree $d \geq 2$ defined over $\bar{\mathbb{Q}}$ and let $P \in \mathbb{P}^N(\bar{\mathbb{Q}})$. Then*

$$P \in \operatorname{PrePer}(\phi) \quad \textit{if and only if} \quad \hat{h}_\phi(P) = 0.$$

Proof. If P is preperiodic, then the quantity $h(\phi^n(P))$ takes on only finitely many values, so it is clear that $d^{-n}h(\phi^n(P)) \to 0$ as $n \to \infty$.

Now suppose that $\hat{h}_\phi(P) = 0$. Let K be a number field containing the coordinates of P and the coefficients of ϕ, i.e., $P \in \mathbb{P}^N(K)$ and ϕ is defined over K. Theorem 3.20 and the assumption $\hat{h}_\phi(P) = 0$ imply that

$$h(\phi^n(P)) = \hat{h}_\phi(\phi^n(P)) + O(1) = d^n \hat{h}_\phi(P) + O(1) = O(1) \quad \text{for all } n \geq 0.$$

Thus the orbit

$$\mathcal{O}_\phi(P) = \{P, \phi(P), \phi^2(P), \phi^3(P), \dots\} \subset \mathbb{P}^N(K)$$

is a set of bounded height, so it is finite from Theorem 3.7. Therefore P is a preperiodic point for ϕ. □

Remark 3.23. Further material on canonical heights in dynamics may be found in Sections 3.5, 5.9, and 7.4, as well as [13, 18, 21, 34, 36, 37, 81, 82, 83, 135, 147, 207, 208, 210, 211, 212, 381, 384, 420, 427]

Theorem 3.22 is a generalization of Kronecker's theorem (Theorem 3.8), which says that $h(\alpha) = 0$ if and only if α is a root of unity. Thus Kronecker's theorem follows by applying Theorem 3.22 to the d^{th}-power map $\phi(z) = z^d$ whose canonical height is the ordinary height h.

The fact that only roots of unity have height 0 leads naturally to the question of how small a nonzero height can be. If we take the relation $h(\alpha^d) = dh(\alpha)$ and substitute in $\alpha = 2^{1/d}$, we find that

$$h(2^{1/d}) = \frac{1}{d}h(2) = \frac{\log 2}{d},$$

so the height can become arbitrarily small. However, this is only possible by taking numbers lying in fields of higher and higher degree. For any algebraic number α, let

$$D(\alpha) = [\mathbb{Q}(\alpha) : \mathbb{Q}]$$

denote the degree of its minimal polynomial over $\mathbb{Q}$.

Conjecture 3.24. (Lehmer's Conjecture [242]) *There is an absolute constant $\kappa > 0$ such that*

$$h(\alpha) \geq \kappa / D(\alpha)$$

for every nonzero algebraic number α that is not a root of unity.

There has been a great deal of work on Lehmer's conjecture by many mathematicians, see for example [87, 7, 6, 70, 242, 340, 396, 398, 419]. The best result currently known, which is due to Dobrowolski [127], says that

$$h(\alpha) \geq \frac{\kappa}{D(\alpha)} \left(\frac{\log \log D(\alpha)}{\log D(\alpha)} \right)^3 .$$

The smallest known nonzero value of $D(\alpha)h(\alpha)$ is

$$D(\beta_0)h(\beta_0) = 0.1623576\ldots,$$

where $\beta_0 = 1.17628\ldots$ is a real root of

$$x^{10} + x^9 - x^7 - x^6 - x^5 - x^4 - x^3 + x + 1.$$

Theorem 3.22 tells us that $\hat{h}_\phi(P) = 0$ if and only if P is a preperiodic point for ϕ. This suggests a natural generalization of Lehmer's conjecture to the dynamical setting. (See [295] for an early version of this conjecture in a special case.)

Conjecture 3.25. (Dynamical Lehmer Conjecture) *Let $\phi : \mathbb{P}^N \to \mathbb{P}^N$ be a morphism defined over a number field K, and for any point $P \in \mathbb{P}^N(\bar{K})$, let $D(P) = [K(P) : K]$. Then there is a constant $\kappa = \kappa(K, \phi) > 0$ such that*

$$\hat{h}_\phi(P) \geq \frac{\kappa}{D(P)} \qquad \textit{for all } P \in \mathbb{P}^N(\bar{K}) \textit{ with } P \notin \operatorname{PrePer}(\phi).$$

There has been considerable work on this conjecture for maps $\phi : \mathbb{P}^1 \to \mathbb{P}^1$ that are associated to groups as described in Section 1.6. For example, in the case that ϕ is attached to an elliptic curve E, it is known that

$$\hat{h}_\phi(P) \geq \begin{cases} \dfrac{\kappa}{D(P)^3 \log^2 D(P)} & \text{in general [269],} \\ \dfrac{\kappa}{D(P)^2} & \text{if } j(E) \text{ is nonintegral [185],} \\ \dfrac{\kappa}{D(P)} \left(\dfrac{\log \log D(P)}{\log D(P)} \right)^3 & \text{if } E \text{ has complex multiplication [241].} \end{cases}$$

Aside from maps associated to groups, there does not appear to be a single example where it is known that $\hat{h}_\phi(P)$ is always greater than a constant over a fixed power of $D(P)$. Using trivial estimates based on the number of points of bounded height in projective space, it is easy to prove a lower bound that decreases faster than exponentially in $D(P)$, see Exercise 3.17.

Remark 3.26. The Lehmer conjecture involves a single map ϕ and points from number fields of increasing size. Another natural question to ask about lower bounds for the canonical height involves fixing the field K and letting the map ϕ vary. For example, consider quadratic polynomials $\phi_c(z) = z^2 + c$ as c varies over $\mathbb{Q}$. Is it true that $\hat{h}_{\phi_c}(\alpha)$ is uniformly bounded away from 0 for all $c \in \mathbb{Q}$ and all nonpreperiodic α? In other words, does there exist a constant $\kappa > 0$ such that

$$\hat{h}_{\phi_c}(\alpha) \ge \kappa \quad \text{for all } c \in \mathbb{Q} \text{ and all } \alpha \notin \operatorname{PrePer}(\phi_c)?$$

We might even ask that the lower bound grow as c becomes larger (in an arithmetic sense). Thus is there a constant $\kappa > 0$ so that

$$\hat{h}_{\phi_c}(\alpha) \ge \kappa h(c) \quad \text{for all } c \in \mathbb{Q} \text{ and all } \alpha \notin \operatorname{PrePer}(\phi_c)?$$

This is a dynamical analog of a conjecture for elliptic curve that is due to Serge Lang, see [184], [232, page 92], or [385, VIII.9.9].

For the quadratic map $z^2 + c$, the height of the parameter c provides a natural measure of its size, but the situation for general rational maps $\phi(z) \in K(z)$ is more complicated. We cannot simply use the height of the coefficients of ϕ, because the canonical height is invariant under conjugation (see Exercise 3.11), while the height of the coefficients is not conjugation invariant. We return to this question in Section 4.11 after we have developed a way to measure the size of the conjugacy class of a rational map.

3.5 Local Canonical Heights

The canonical height $\hat{h}_\phi$ attached to a rational map ϕ is a useful tool in studying the arithmetic dynamics of ϕ. For more refined analyses, it is helpful to decompose the canonical height as a sum of local canonical heights, one for each absolute value on K. In this section we briefly summarize the basic properties of local canonical heights, but we defer the proofs until Section 5.9. The reader wishing to proceed more rapidly to the main arithmetic results of this chapter may safely omit this section on first reading, since the material covered is not used elsewhere in this book.

The construction of the canonical height relies on the fact that the ordinary height satisfies $h(\phi(P)) = dh(P) + O(1)$, so it is "almost canonical." The ordinary height of a point $P = [\alpha, 1]$ is equal to the sum

$$h(P) = h(\alpha) = \sum_{v \in M_K} n_v \log \max\{|\alpha|_v, 1\},$$

so for each $v \in M_K$ it is natural to define a local height function

$$\lambda_v(\alpha) = \log \max\{|\alpha|_v, 1\}.$$

We can understand λ_v geometrically by observing that for $v \in M_K^0$,

$$\lambda_v(\alpha) = -\log \rho_v(\alpha, \infty),$$

where ρ_v is the nonarchimedean chordal metric defined in Section 2.1. One says that $\lambda_v(\alpha)$ is the *logarithmic distance* from α to ∞.

Unfortunately, the function λ_v does not transform canonically, since $\lambda_v(\phi(\alpha))$ is not equal to $d\lambda_v(\alpha) + O(1)$. To see why, note that $\lambda_v(\phi(\alpha))$ is large if α is close to a pole of ϕ, while $\lambda_v(\alpha)$ is large if α is close to the point $\infty \in \mathbb{P}^1$. (Here the word

7.1 Dynamics Of Rational Maps On Projective Space

Recall that a *rational map* $\phi : \mathbb{P}^N \to \mathbb{P}^N$ is described by homogeneous polynomials with no common factor, and that ϕ is a *morphism* if the polynomials have no common root in $\mathbb{P}^N(\bar{K})$. (See page 88 in Chapter 3 for the precise definition.) As noted in the introduction to this chapter, height functions are a powerful tool for studying the arithmetic of morphisms $\phi : \mathbb{P}^N \to \mathbb{P}^N$. The situation is considerably more complicated if the map $\phi : \mathbb{P}^N \to \mathbb{P}^N$ is only required to be a rational map. Notice that we did not run into this situation when studying rational functions $\phi(z) \in K(z)$ of one variable, since every rational map $\phi : \mathbb{P}^1 \to \mathbb{P}^1$ is automatically a morphism. But in dimensions 2 and higher, there are many rational maps that are not morphisms.

Example 7.1. The rational map

$$\phi : \mathbb{P}^2 \longrightarrow \mathbb{P}^2, \qquad \phi\big([X, Y, Z]\big) = [X_0^2, X_0X_1, X_2^2], \tag{7.1}$$

is not a morphism, since it is not defined at the point $[0, 1, 0]$. Notice that if we discard $[0, 1, 0]$, then ϕ fixes every point on the line $X_0 = X_2$, and ϕ sends every point on the line $X_0 = 0$ to the single point $[0, 0, 1]$. This sort of behavior is not possible for morphisms $\mathbb{P}^2 \to \mathbb{P}^2$.

Continuing with this example, recall that if ϕ were a morphism, then Theorem 3.11 would tell us that $h\big(\phi(P)\big) = 2h(P) + O(1)$ for all $P \in \mathbb{P}^2(\bar{\mathbb{Q}})$. But this is clearly false for the map (7.1), since for all $a, b \in \bar{\mathbb{Q}}^*$ we have

$$\phi\big([a, b, a]\big) = [a^2, ab, a^2] = [a, b, a].$$

Thus

$$h\big(\phi\big([a, b, a]\big)\big) = h\big([a, b, a]\big),$$

so we cannot use Theorem 3.7 to conclude that ϕ has only finitely many $\mathbb{Q}$-rational periodic points. Of course, that's good, since in fact ϕ has infinitely many $\mathbb{Q}$-rational fixed points!

An initial difficulty in studying the dynamics of a rational map $\phi : \mathbb{P}^N \to \mathbb{P}^N$ arises from the fact that the orbit $\mathcal{O}_\phi(P)$ of a point may "terminate" if some iterate $\phi^n(P)$ arrives at a point where ϕ is not defined. This suggests looking first at maps ϕ for which there is a large uncomplicated (e.g., quasiprojective, or even affine) subset $U \subset \mathbb{P}^N$ with the property that $\phi(U) \subset U$ and studying the dynamics of ϕ on U. As a further simplication, we might require that ϕ be an automorphism of U, since quasiprojective varieties often allow interesting automorphisms.

7.1.1 Affine Morphisms and the Locus of Indeterminacy

In this section we study rational maps $\mathbb{P}^N \to \mathbb{P}^N$ with the property that they induce morphisms of affine space $\mathbb{A}^N \to \mathbb{A}^N$. Concretely, an *affine morphism*

$$\phi : \mathbb{A}^N \to \mathbb{A}^N$$

is a map of the form

$$\phi = (F_1, \cdots, F_N) \qquad \text{with} \quad F_1, \ldots, F_N \in K[z_1, \ldots, z_N].$$

To avoid trivial cases, we generally assume that at least one of the F_i is not the 0-polynomial.

Definition. The degree of a polynomial

$$F(z_1, \ldots, z_N) = \sum_{i_1, \ldots, i_N} a_{i_1 \cdots i_N} z_1^{i_1} \cdots z_N^{i_N} \in K[z_1, \ldots, z_N]$$

is defined to be

$$\deg F = \max\bigl\{ i_1 + \cdots + i_N : a_{i_1 \cdots i_N} \neq 0 \bigr\}.$$

In other words, the degree of F is the largest total degree of the monomials that appear in F. (By convention the 0-polynomial is assigned degree $-\infty$.) The *degree of a morphism* $\phi = (F_1, \ldots, F_N) : \mathbb{A}^N \to \mathbb{A}^N$ is defined to be

$$\deg \phi = \max\{\deg F_1, \ldots, \deg F_N\}.$$

Homogenization of the coordinates of an affine morphism $\phi : \mathbb{A}^N \to \mathbb{A}^N$ of degree d yields a rational map $\bar{\phi} : \mathbb{P}^N \to \mathbb{P}^N$ of degree d. For each coordinate function F_i of ϕ, we let

$$\bar{F}_i(X_0, X_1, \ldots, X_N) = X_0^d F_i \left(\frac{X_1}{X_0}, \frac{X_2}{X_0}, \ldots, \frac{X_N}{X_0} \right).$$

Notice that each $\bar{F}_i$ is a homogeneous polynomial of degree d (or the 0-polynomial), so the map

$$\bar{\phi} = [X_0^d, \bar{F}_1, \bar{F}_2, \ldots, \bar{F}_N] : \mathbb{P}^N \to \mathbb{P}^N$$

is a rational map of degree d. We call $\bar{\phi}$ the *rational map induced by* ϕ. A rational map need not be everywhere defined.

Definition. Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be an affine morphism of degree d and let

$$\bar{\phi} = [X_0^d, \bar{F}_1, \ldots, \bar{F}_N] : \mathbb{P}^N \to \mathbb{P}^N$$

be the rational map that it induces. The *locus of indeterminacy of* ϕ is the set

$$Z(\phi) = \bigl\{ P = [0, x_1, \ldots, x_N] \in \mathbb{P}^N : F_1(P) = \cdots = F_N(P) = 0 \bigr\}.$$

This is the set of points at which $\bar{\phi}$ is not defined. Notice that $Z(\phi)$ lies in the hyperplane $H_0 = \{X_0 = 0\}$ at infinity, since ϕ is well-defined on $\mathbb{A}^N$.

The polynomials $\bar{F}_1, \ldots, \bar{F}_N$ can be used to define a morphism

$$\Phi : \mathbb{A}^{N+1} \longrightarrow \mathbb{A}^{N+1}, \qquad \Phi = (X_0^d, \bar{F}_1, \ldots, \bar{F}_N).$$

The map Φ is called a *lift of* $\bar{\phi}$. If we let π be the natural projection map,

$$\pi : \mathbb{A}^{N+1} \smallsetminus \{0\} \longrightarrow \mathbb{P}^N, \qquad (x_0, \ldots, x_N) \longmapsto [x_0, \ldots, x_N],$$

then π, Φ, and $\bar{\phi}$ fit together into the commutative diagram

$$\begin{array}{ccc} \mathbb{A}^{N+1} \smallsetminus \{0\} & \xrightarrow{\Phi} & \mathbb{A}^{N+1} \smallsetminus \{0\} \\ \Big\downarrow \pi & & \Big\downarrow \pi \\ \mathbb{P}^N & \xrightarrow{\bar{\phi}} & \mathbb{P}^N \end{array}$$

Example 7.2. The map

$$\phi : \mathbb{A}^2 \longrightarrow \mathbb{A}^2, \qquad \phi(z_1, z_2) = (z_1 z_2, z_1^2)$$

induces the rational map

$$\phi : \mathbb{P}^2 \longrightarrow \mathbb{P}^2, \qquad \bar{\phi}\bigl([X_0, X_1, X_2]\bigr) = [X_0^2, X_1 X_2, X_1^2]$$

and has indeterminacy locus $Z(\phi) = \bigl\{[0,0,1]\bigr\}$ consisting of a single point.

7.1.2 Affine Automorphisms

Of particular interest are affine morphisms that admit an inverse.

Definition. An affine morphism $\phi : \mathbb{A}^N \to \mathbb{A}^N$ is an *automorphism* if it has an inverse morphism. In other words, ϕ is an affine automorphism if there is an affine morphism $\phi^{-1} : \mathbb{A}^N \to \mathbb{A}^N$ such that

$$\phi\bigl(\phi^{-1}(z_1, \ldots, z_N)\bigr) = (z_1, \ldots, z_N) \quad \text{and} \quad \phi^{-1}\bigl(\phi(z_1, \ldots, z_N)\bigr) = (z_1, \ldots, z_N).$$

Somewhat surprisingly, ϕ and ϕ^{-1} need not have the same degree, nor does $\deg(\phi^n)$ have to equal $(\deg \phi)^n$.

Example 7.3. Consider the map $\phi(x, y) = (x, y + x^2)$. It has degree 2 and is an automorphism, since it has the inverse $\phi^{-1}(x, y) = (x, y - x^2)$. The composition ϕ^2 is

$$\phi^2(x, y) = \phi(x, y + x^2) = (x, y + 2x^2),$$

so $\deg(\phi^2) = 2 = \deg(\phi)$. More generally, $\phi^n(x, y) = (x, y + nx^2)$ has degree 2, so the degree of ϕ^n does not grow. This contrasts sharply with what happens for morphisms of $\mathbb{P}^N$.

Example 7.4. Let $a \in K^*$ and let $f(y) \in K[y]$ be a polynomial of degree $d \geq 2$. The map

$$\phi : \mathbb{A}^2 \longrightarrow \mathbb{A}^2, \qquad \phi(x, y) = \bigl(y, ax + f(y)\bigr),$$

is called a *Hénon map*. It is an automorphism of $\mathbb{A}^2$, since one easily checks that it has an inverse ϕ^{-1} given by

$$\phi^{-1} : \mathbb{A}^2 \longrightarrow \mathbb{A}^2, \qquad \phi^{-1}(x, y) = \bigl(a^{-1}y - a^{-1}f(x), x\bigr).$$

Hénon maps, especially those with $\deg(f) = 2$, have been extensively studied since Hénon [182] introduced them as examples of maps $\mathbb{R}^2 \to \mathbb{R}^2$ having strange attractors. There are many open questions regarding the real and complex dynamics of Hénon maps, see for example [122, §2.9] or [193], as well as [194, 388] for a compactification of the Hénon map.

The rational maps $\mathbb{P}^2 \to \mathbb{P}^2$ induced by ϕ and ϕ^{-1} are

$$\begin{aligned}\bar{\phi}\bigl([X_0, X_1, X_2]\bigr) &= [X_0^d,\ X_0^{d-1}X_2,\ aX_0^{d-1}X_1 + \bar{f}(X_0, X_2)],\\ \bar{\phi}^{-1}\bigl([X_0, X_1, X_2]\bigr) &= [X_0^d,\ a^{-1}X_0^{d-1}X_2 - a^{-1}\bar{f}(X_0, X_1),\ X_0^{d-1}X_1],\end{aligned}$$

where we write $\bar{f}(u, v) = u^d f(v/u)$ for the homogenization of f. It is easy to see that the loci of indeterminacy of ϕ and ϕ^{-1} are

$$Z(\phi) = \bigl\{[0, 1, 0]\bigr\} \qquad \text{and} \qquad Z(\phi^{-1}) = \bigl\{[0, 0, 1]\bigr\}.$$

In particular, the locus of indeterminacy of ϕ is disjoint from the locus of indeterminacy of ϕ^{-1}. Maps with this property are called *regular*, see Section 7.1.3.

Example 7.5. Consider the very simple Hénon map

$$\phi(x, y) = (y, -x + y^2).$$

The extension $\bar{\phi} = [X_0^2, X_0X_2, -X_0X_1 + X_2^2]$ of ϕ to $\mathbb{P}^2$ has degree 2, but it is not a morphism, since it is not defined at the point $[0, 1, 0]$. And just as in Example 7.2, there is no height estimate of the form $h\bigl(\phi(P)\bigr) = 2h(P) + O(1)$ for $\bar{\phi}$. We can see this by noting that

$$\bar{\phi}\bigl([b, a, b]\bigr) = [b^2, b^2, -ab + b^2] = [b, b, -a + b],$$

so if $a, b, \in \mathbb{Z}$ with $\gcd(a, b) = 1$ and $b > a > 0$, then $[b, a, b]$ and $\bar{\phi}\bigl([b, a, b]\bigr)$ have the same height. Hence for every $\epsilon > 0$ even the weaker statement

$$h\bigl(\bar{\phi}(P)\bigr) \ge (1 + \epsilon)h(P) + O(1) \qquad \text{for all } P = (x, y) \in \mathbb{A}^2(\mathbb{Q})$$

is false. It turns out that ϕ has only finitely many $\mathbb{Q}$-rational periodic points (Theorem 7.18), but the proof does not follow directly from a simple height argument.

Example 7.6. More generally, if $\phi : \mathbb{A}^N \to \mathbb{A}^N$ is an affine automorphism, then it is not possible to have simultaneous estimates of the form

$$\begin{aligned}h\bigl(\phi(P)\bigr) &\ge (1 + \epsilon)h(P) + O(1),\\ h\bigl(\phi^{-1}(P)\bigr) &\ge (1 + \epsilon)h(P) + O(1),\end{aligned} \tag{7.2}$$

for some $\epsilon > 0$ and all $P \in \mathbb{A}^N(K)$. To see this, suppose that (7.2) were true. Then we would have for all $P \in \mathbb{A}^N(K)$,

$$h(P) = h\bigl(\phi(\phi^{-1}(P))\bigr) \ge (1 + \epsilon)h\bigl(\phi^{-1}(P)\bigr) + O(1) \ge (1 + \epsilon)^2 h(P) + O(1).$$

Thus $h(P)$ would be bounded, leading to the untenable conclusion that $\mathbb{A}^N(K)$ is finite. So it is too much to require that both $\phi(P)$ and $\phi^{-1}(P)$ have heights larger than the height of P. However, as we shall see, it is often possible to show that some combination of $h\bigl(\phi(P)\bigr)$ and $h\bigl(\phi^{-1}(P)\bigr)$ is large, which is then sufficent to prove that $\operatorname{Per}(\phi)$ is a set of bounded height.

We conclude this section with two useful geometric lemmas. The first relates the locus of indetermincay of an affine automorphism and its inverse, and the second characterizes when the degree of a composition is smaller than the product of the degrees.

Lemma 7.7. *Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be an affine automorphism of degree at least 2 and denote the hyperplane at infinity by $H_0 = \{X_0 = 0\} = \mathbb{P}^N \smallsetminus \mathbb{A}^N$. Then*

$$\bar\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\phi^{-1}).$$

Proof. Let

$$\Phi = (X_0^d, \bar F_1, \bar F_2, \dots, \bar F_N) \quad \text{and} \quad \Phi^{-1} = (X_0^e, \bar G_1, \bar G_2, \dots, \bar G_N)$$

be the lifts of $\bar\phi$ and $\bar\phi^{-1}$, respectively. The fact that ϕ and ϕ^{-1} are inverses of one another implies that there is a homogeneous polynomial f of degree $de-1$ with the property that

$$(\Phi^{-1} \circ \Phi)(X_0, \dots, X_N) = (f \cdot X_0, f \cdot X_1, \dots, f \cdot X_N).$$

But the first coordinate of the composition is X_0^{de}, so we see that $f = X_0^{de-1}$. Thus

$$(\Phi^{-1} \circ \Phi)(X_0, \dots, X_N) = (X_0^{de}, X_0^{de-1}X_1, X_0^{de-1}X_1, \dots, X_0^{de-1}X_N),$$

or equivalently,

$$\bar G_j(X_0^d, \bar F_1, \dots, \bar F_N) = X_0^{de-1}X_j \quad \text{for all } 1 \le j \le N. \tag{7.3}$$

Now let $P = [0, x_1, \dots, x_N] \in H_0 \smallsetminus Z(\phi)$, so $\phi(P) = \bigl[0, \bar F_1(P), \dots, \bar F_N(P)\bigr]$ with at least one $\bar F_i(P) \ne 0$. From (7.3) we see that

$$\bar G_j\bigl(\Phi(P)\bigr) = \bar G_j\bigl(0, \bar F_1(P), \dots, \bar F_N(P)\bigr) = 0^{de-1}x_j = 0 \quad \text{for all } 1 \le j \le N.$$

Hence

$$\Phi^{-1}\bigl(\Phi(P)\bigr) = \bigl(0, \bar G_1\bigl(\Phi(P)\bigr), \bar G_2\bigl(\Phi(P)\bigr), \dots, \bar G_N\bigl(\Phi(P)\bigr) = (0,0,0,\dots,0),$$

so ϕ^{-1} is not defined at $\phi(P)$. Therefore $\phi(P) \in Z(\phi^{-1})$. $\square$

Lemma 7.8. *Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ and $\psi : \mathbb{A}^N \to \mathbb{A}^N$ be affine morphisms, and let $H_0 = \{X_0 = 0\} = \mathbb{P}^N \smallsetminus \mathbb{A}^N$ be the usual hyperplane at infinity. Then*

$$\deg(\psi \circ \phi) < \deg(\psi)\deg(\phi) \quad \textit{if and only if} \quad \bar\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\psi).$$

Proof. Let $d = \deg(\phi)$, let $e = \deg(\psi)$, and let Φ and Ψ be lifts of $\bar\phi$ and $\bar\psi$, respectively. We write Φ explicitly as

$$\Phi = (X_0^d, \bar F_1, \bar F_2, \dots, \bar F_N).$$

The composition $\Psi \circ \Phi$ has the form

$$\Psi \circ \Phi = (X_0^{de}, \bar{E}_1, \bar{E}_2, \ldots, \bar{E}_N),$$

where $\bar{E}_1, \ldots, \bar{E}_N$ are homogeneous polynomials of degree de. The degree of $\psi \circ \phi$ will be strictly less than de if and only if there is some cancellation in the coordinate polynomials of $\Psi \circ \Phi$. Since the first coordinate is X_0^{de}, this shows that

$$\deg(\psi \circ \phi) < \deg(\psi)\deg(\phi) \quad \Longleftrightarrow \quad X_0 \text{ divides } \bar{E}_j \text{ for every } 1 \le j \le N.$$

Suppose now that $X_0 | \bar{E}_j$ for every j and let $P = [0, x_1, \ldots, x_N] \in H_0 \smallsetminus Z(\phi)$. Since ϕ is defined at P, some coordinate of

$$\Phi(P) = \bigl(0, \bar{F}_1(P), \ldots, \bar{F}_N(P)\bigr)$$

is nonzero. On the other hand, the assumption that $X_0 | \bar{E}_j$ implies that

$$(\Psi \circ \Phi)(P) = \bigl(0, \bar{E}_1(\Phi(P)), \bar{E}_2(\Phi(P)), \ldots, \bar{E}_N(\Phi(P))\bigr) = (0, 0, 0, \ldots, 0).$$

Hence ψ is not defined at $\phi(P)$, so $\phi(P) \in Z(\psi)$. This completes the proof that if $\deg(\psi \circ \phi) < de$, then $\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\psi)$.

For the other direction, suppose that $\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\psi)$. This implies that for (almost all) points of the form $(0, x_1, \ldots, x_N)$, the map ψ is not defined at the point $\phi\bigl([0, x_1, \ldots, x_N]\bigr)$. Hence

$$\Psi\bigl(\Phi(0, X_1, X_2, \ldots, X_N)\bigr) = (0, 0, 0, \ldots, 0),$$

so $\bar{E}_j(0, X_1, X_2, \ldots, X_N) = 0$ for all j. Therefore $X_0 | \bar{E}_j$ for all j. $\square$

Example 7.9. Let ϕ be the map $\phi(x, y) = (x, y + x^2)$ that we studied in Example 7.3. Dehomogenizing ϕ yields

$$\bar{\phi}\bigl([X_0, X_1, X_2]\bigr) = [X_0^2,\ X_0X_1,\ X_0X_2 + X_1^2],$$

so the locus of indeterminacy for ϕ is $Z(\phi) = \bigl\{[0, 0, 1]\bigr\}$. Notice that

$$\bar{\phi}\bigl([0, X_1, X_2]\bigr) = [0, 0, X_1^2] = [0, 0, 1] \in Z(\phi).$$

Hence $\bar{\phi}\bigl(H_0 \smallsetminus Z(\phi)\bigr) = Z(\phi)$, so Lemma 7.8 tells us that $\deg(\phi^2) < \deg(\phi)^2$. This is in agreement with Example 7.3, where we computed that $\deg(\phi^2) = 2$.

7.1.3 The Geometry of Regular Automorphisms of $\mathbb{A}^N$

In this section we briefly discuss the geometric properties of an important class of affine automorphisms.

Definition. An affine automorphism $\phi : \mathbb{A}^N \to \mathbb{A}^N$ is said to be *regular* if the indeterminacy loci of ϕ and ϕ^{-1} have no points in common,

$$Z(\phi) \cap Z(\phi^{-1}) = \emptyset.$$

The following theorem summarizes some of the geometric properties enjoyed by regular automorphisms of $\mathbb{A}^N$. We sketch the proof of (a) and refer the reader to [376] for (b) and (c).

Theorem 7.10. *Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be a regular affine automorphism.*

(a) *For all $n \ge 1$,*

$$\phi^n \textit{ is regular,} \qquad Z(\phi^n) = Z(\phi), \qquad \textit{and} \qquad \deg(\phi^n) = \deg(\phi)^n.$$

(b) *Let*

$$d_1 = \deg\phi, \quad d_2 = \deg\phi^{-1}, \quad \ell_1 = \dim Z(\phi)+1, \quad \ell_2 = \dim Z(\phi^{-1})+1.$$

Then

$$\ell_1 + \ell_2 = N \qquad \textit{and} \qquad d_2^{\ell_1} = d_1^{\ell_2}.$$

(c) *For all $n \ge 1$ the set of n-periodic points $\operatorname{Per}_n(\phi)$ is a discrete subset of $\mathbb{A}^N(\mathbb{C})$. Counted with appropriate multiplicities,*

$$\#\operatorname{Per}_{Nn}(\phi) = d_2^{\ell_1 Nn} = d_1^{\ell_2 Nn}.$$

Proof. (a) We first prove by induction on n that

$$Z(\phi^n) \subset Z(\phi) \quad \text{and} \quad Z(\phi^{-n}) \subset Z(\phi^{-1}) \quad \text{for all } n \ge 1.$$

This is trivially true for $n = 1$, so assume now that it is true for $n-1$. Let $P \in Z(\phi^n)$, so in particular $P \in H_0$. Suppose that $P \notin Z(\phi)$. The induction hypothesis tells us that $P \notin Z(\phi^{n-1})$, so applying Lemma 7.7 to the map ϕ^{n-1}, we deduce that

$$\phi^{n-1}(P) \in \phi^{n-1}\bigl(H_0 \smallsetminus Z(\phi^{n-1})\bigr) \subset Z(\phi^{-(n-1)}) \subset Z(\phi^{-1}).$$

(For the last equality we have again used the induction hypothesis.) On the other hand, we have that ϕ^{n-1} is defined at P and ϕ^n is not defined at P, which implies that $\phi^{n-1}(P) \in Z(\phi)$. This proves that $\phi^{n-1}(P)$ is in both $Z(\phi^{-1})$ and $Z(\phi)$, contradicting the assumption that ϕ is regular. Hence $P \in Z(\phi)$, which completes the proof that $Z(\phi^n) \subset Z(\phi)$. Similarly, we find that $Z(\phi^{-n}) \subset Z(\phi^{-1})$.

Having shown that $Z(\phi^n) \subset Z(\phi)$ and $Z(\phi^{-n}) \subset Z(\phi^{-1})$, the regularity of ϕ implies that

$$Z(\phi^n) \cap Z(\phi^{-n}) \subset Z(\phi) \cap Z(\phi^{-1}) = \emptyset,$$

so ϕ^n is also regular.

Next suppose that $\deg(\phi^n) < \deg(\phi)^n$ for some $n \ge 2$. We take n to be the smallest value for which this is true, so in particular $\deg(\phi^{n-1}) = \deg(\phi)^{n-1}$, and hence

$$\deg(\phi^n) < \deg(\phi^{n-1})\deg(\phi).$$

We apply Lemma 7.8 with $\psi = \phi^{n-1}$ to conclude that

$$\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\phi^{n-1}) \subset Z(\phi),$$

where the last inclusion was proven earlier. On the other hand, Lemma 7.7 says that $\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\phi^{-1})$. Hence

$$\phi\bigl(H_0 \smallsetminus Z(\phi)\bigr) \subset Z(\phi) \cap Z(\phi^{-1}) = \emptyset.$$

This is a contradiction, which completes the proof that $\deg(\phi^n) = \deg(\phi)^n$.

It remains to show that $Z(\phi) \subset Z(\phi^n)$. Let

$$\Phi : \mathbb{A}^{N+1} \longrightarrow \mathbb{A}^{N+1}, \qquad \Phi = (X_0^d, F_1, F_2, \dots, F_N),$$

be a lift of ϕ, so

$$Z(\phi) = \bigl\{P \in H_0 : F_1(P) = \cdots = F_N(P) = 0\bigr\}.$$

By a slight abuse of notation, we say that $P \in Z(\phi)$ if and only if $\Phi(P) = 0$. (To be precise, we should lift P to $\mathbb{A}^{N+1}$.)

We proved that $\deg(\phi^n) = \deg(\phi)$, which implies that the coordinate functions of Φ^n have no common factor. Thus ϕ^n can be computed by evaluating Φ^n and mapping down to $\mathbb{P}^N$. Hence just as above we have $P \in Z(\phi^n)$ if and only if $\Phi^n(P) = 0$. Therefore

$$P \in Z(\phi) \quad\Longrightarrow\quad \Phi(P) = 0 \quad\Longrightarrow\quad \Phi^n(P) = 0 \quad\Longrightarrow\quad P \in Z(\phi^n).$$

This proves that $Z(\phi) \subset Z(\phi^n)$ and completes the proof of (a).
(b) See [376, Proposition 2.3.2].
(c) See [376, Theorem 2.3.4]. □

Remark 7.11. If $\phi : \mathbb{A}^2 \to \mathbb{A}^2$ is a regular automorphism of the affine plane, then Theorem 7.10(b) tells us that $\ell_1 = \ell_2 = 1$ (which is clear anyway since the indeterminacy locus of a rational map has codimension at least 2) and that $d_1 = d_2$. Thus planar regular automorphisms satisfy $\deg(\phi) = \deg(\phi^{-1})$. In the opposite direction, if $d_1 = d_2$, then Theorem 7.10(b) says that $\ell_1 = \ell_2$, and hence that $N = \ell_1 + \ell_2$ is even. In other words, a regular automorphism $\phi : \mathbb{A}^N \to \mathbb{A}^N$ with N odd always satisfies $\deg(\phi) \ne \deg(\phi^{-1})$.

Example 7.12. Let $\phi : \mathbb{A}^3 \to \mathbb{A}^3$ be given by

$$\phi(x, y, z) = (y, z + y^2, x + z^2).$$

One can check that the inverse of ϕ is

$$\phi^{-1}(x, y, z) = \bigl(z - (y - x^2)^2, x, y - x^2\bigr).$$

Homogenizing $x = X_1/X_0$, $y = X_2/X_0$, $z = X_3/X_0$, we have the formulæ

$$\begin{aligned} \bar{\phi} &= [X_0^2,\ X_0X_2,\ X_0X_3 + X_2^2,\ X_0X_1 + X_3^2], \\ \bar{\phi}^{-1} &= [X_0^4,\ X_0^3X_3 - (X_0X_2 - X_1^2)^2,\ X_0^3X_1,\ X_0^3X_2 - X_0^2X_1^2], \end{aligned}$$

from which it is easy to check that

$$\begin{aligned} Z(\phi) &= \{X_0 = X_2 = X_3 = 0\} = \{[0,1,0,0]\}, \\ Z(\phi^{-1}) &= \{X_0 = X_1 = 0\} \qquad\quad = \{[0,0,u,v]\}. \end{aligned}$$

Thus $Z(\phi)$ consists of a single point, while $Z(\phi^{-1})$ is a line. In the notation of Theorem 7.10, we have $N = 3$ and

$$\begin{aligned} d_1 &= \deg \phi = 2, & d_2 &= \deg \phi^{-1} = 4, \\ \ell_1 &= \dim Z(\phi) + 1 = 1, & \ell_2 &= \dim Z(\phi^{-1}) = 1 = 2. \end{aligned}$$

The map ϕ is regular, since $Z(\phi) \cap Z(\phi^{-1}) = \emptyset$.

Remark 7.13. Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be an affine morphism and let $\Phi : \mathbb{A}^{N+1} \to \mathbb{A}^{N+1}$ be a lift of ϕ. The map ϕ is called *algebraically stable* if

$$\Phi^n\bigl(\{X_0 = 0\}\bigr) \ne \{0\} \quad \text{for all } n \ge 1.$$

In other words, ϕ is algebraically stable if for every $n \ge 1$, some coordinate of $\Phi^n(X_0, \dots, X_N)$ is not divisible by X_0. Since the first coordinate of Φ^n is a power of X_0, this implies that there can be no cancellation among the coordinates, so an algebraically stable map ϕ satisfies

$$\deg(\phi^n) = \deg(\phi)^n.$$

Further, an adaptation of the proof of Theorem 7.10(a) shows that

$$Z(\phi^n) \subset Z(\phi^m) \quad \text{for all } n < m.$$

Regular automorphisms are algebraically stable, but there are algebraically stable automorphisms that are not regular. For a discussion of the complex dynamics of algebraically stable maps, see [161, 170, 376].

For arbitrary morphisms $\phi : \mathbb{A}^N \to \mathbb{A}^N$, we define the *dynamical degree of* ϕ to be the quantity

$$\operatorname{dyndeg}(\phi) = \lim_{n\to\infty} \deg(\phi^n)^{1/n}.$$

The dynamical degree provides a coarse measure of the stable complexity of the map ϕ, and presumably it has an impact on the arithmetic properties of ϕ. See [268] for an indication of this effect in certain cases. One can show that the dynamical degree is in fact the infimum of $\deg(\phi^n)^{1/n}$. The dynamical degree need not be an integer, nor even a rational number, see Exercise 7.4 for an example.

7.1.4 A Height Bound for Jointly Regular Affine Morphisms

In this section we prove a nontrivial lower bound for the height of points under regular affine automorphisms. The theorem is an amalgamation of results due to Denis [121], Kawaguchi [210, 207], Marcello [265, 266, 267, 268] and Silverman [388, 393]. Before stating the theorem, we need to define what is meant by the height of a point in affine space.

Definition. The height $h(P)$ of a point $P = (x_1, \ldots, x_N) \in \mathbb{A}^N(\bar{\mathbb{Q}})$ in affine space is defined to be the height of the associated point in projective space using the natural embedding $\mathbb{A}^N \to \mathbb{P}^N$,

$$h(P) = h\bigl([1, x_1, \ldots, x_N]\bigr).$$

Eventually we will apply the following height estimate to a regular affine automorphism ϕ and its inverse ϕ^{-1}, but it is no harder to prove the result for any pair of jointly regular maps, and working in a general setting helps clarify the underlying structure of the proof.

Theorem 7.14. *Let $\phi_1 : \mathbb{A}^N \to \mathbb{A}^N$ and $\phi_2 : \mathbb{A}^N \to \mathbb{A}^N$ be affine morphisms with the property that*

$$Z(\phi_1) \cap Z(\phi_2) = \emptyset.$$

(*We say that ϕ_1 and ϕ_2 are* jointly regular.) *Let*

$$d_1 = \deg \phi_1 \qquad \textit{and} \qquad d_2 = \deg \phi_2.$$

There is a constant $C = C(\phi_1, \phi_2)$ so that for all $P \in \mathbb{A}^N(\bar{\mathbb{Q}})$,

$$\frac{1}{d_1} h\bigl(\phi_1(P)\bigr) + \frac{1}{d_2} h\bigl(\phi_2(P)\bigr) \ge h(P) - C. \tag{7.4}$$

Remark 7.15. We recall that the upper bound

$$h\bigl(\psi(P)\bigr) \le (\deg \psi) h(P) + O(1) \tag{7.5}$$

is valid even for rational maps $\psi : \mathbb{P}^N \to \mathbb{P}^N$ (see Theorem 3.11), since the proof of (7.5) uses only the triangle inequality. Thus Theorem 7.14 may be viewed as providing a nontrivial lower bound complementary to the elementary upper bound

$$\frac{1}{d_1} h\bigl(\phi_1(P)\bigr) + \frac{1}{d_2} h\bigl(\phi_2(P)\bigr) \le 2h(P) + O(1).$$

Proof of Theorem 7.14. Write the rational functions $\mathbb{P}^N \to \mathbb{P}^N$ induced by ϕ_1 and ϕ_2 as

$$\bar{\phi}_2 = [X_0^{d_1}, \bar{F}_1, \bar{F}_2, \ldots, \bar{F}_N] \qquad \text{and} \qquad \bar{\phi}_2 = [X_0^{d_2}, \bar{G}_1, \bar{G}_2, \ldots, \bar{G}_N],$$

where the $\bar{F}_i$ are homogeneous polynomials of degree d_1 and the $\bar{G}_i$ are homogeneous polynomials of degree d_2. The loci of indeterminacy of ϕ_1 and ϕ_2 are given by

$$Z(\phi_1) = \{X_0 = \bar{F}_1 = \cdots = \bar{F}_N = 0\},$$
$$Z(\phi_2) = \{X_0 = \bar{G}_1 = \cdots = \bar{G}_N = 0\}.$$

We define a rational map $\psi : \mathbb{P}^{2N} \to \mathbb{P}^{2N}$ of degree $d_1 d_2$ by

$$\psi = \left[X_0^{d_1 d_2}, \bar{F}_1^{d_2}, \ldots, \bar{F}_N^{d_2}, \bar{G}_1^{d_1}, \ldots, \bar{G}_N^{d_1}\right].$$

The locus of indeterminacy of ψ is the set

$$Z(\psi) = \{X_0 = \bar{F}_1 = \cdots = \bar{F}_N = \bar{G}_1 = \cdots = \bar{G}_N = 0\} = Z(\phi_1) \cap Z(\phi_2) = \emptyset,$$

since by assumption $Z(\phi_1)$ and $Z(\phi_2)$ are disjoint. Hence ψ is a morphism, so we can apply the fundamental height estimate for morphisms (Theorem 3.11) to deduce that

$$h\big(\psi(P)\big) = d_1 d_2 h(P) + O(1) \qquad \text{for all } P \in \mathbb{P}^{2N}(\bar{\mathbb{Q}}). \tag{7.6}$$

The following lemma will give us an upper bound for the height of $\psi(P)$.

Lemma 7.16. *Let* $u, a_1, \ldots, a_N, b_1, \ldots, b_N \in \bar{\mathbb{Q}}$ *with* $u \neq 0$. *Then*

$$h\big([u, a_1, \ldots, a_N, b_1, \ldots, b_N]\big) \le h\big([u, a_1, \ldots, a_N]\big) + h\big([u, b_1, \ldots, b_N]\big).$$

Proof. Let $\alpha_i = a_i/u$ and $\beta_i = b_i/u$ for $1 \le i \le N$. Then for any absolute value v we have the trivial estimate

$$\begin{aligned} &\max\big\{1, |\alpha_1|_v, \ldots, |\alpha_N|_v, |\beta_1|_v, \ldots, |\beta_N|_v\big\} \\ &\qquad\qquad \le \max\big\{1, |\alpha_1|_v, \ldots, |\alpha_N|_v\big\} \cdot \max\big\{1, |\beta_1|_v, \ldots, |\beta_N|_v\big\}. \end{aligned}$$

Raising to an appropriate power, multiplying over all absolute values, and taking logarithms yields

$$h\big([1, \alpha_1, \ldots, \alpha_N, \beta_1, \ldots, \beta_N]\big) \le h\big([1, \alpha_1, \ldots, \alpha_N]\big) + h\big([1, \beta_1, \ldots, \beta_N]\big).$$

This is the desired result, since the height does not depend on the choice of homogeneous coordinates of a point. $\square$

We apply Lemma 7.16 to the point

$$\psi(P) = \big[X_0(P)^{d_1 d_2}, \bar{F}_1(P)^{d_2}, \ldots, \bar{F}_N(P)^{d_2}, \bar{G}_1(P)^{d_1}, \ldots, \bar{G}_N(P)^{d_1}\big]$$

with $P \in \mathbb{A}^N(\bar{\mathbb{Q}})$, which ensures that $X_0(P) \neq 0$. The lemma tells us that

$$\begin{aligned} h\big(\psi(P)\big) &\le h\left(\big[X_0(P)^{d_1 d_2}, \bar{F}_1(P)^{d_2}, \ldots, \bar{F}_N(P)^{d_2}\big]\right) \\ &\qquad\qquad + h\left(\big[X_0(P)^{d_1 d_2}, \bar{G}_1(P)^{d_1}, \ldots, \bar{G}_N(P)^{d_1}\big]\right) \\ &= d_2 h\left(\big[X_0(P)^{d_1}, \bar{F}_1(P), \ldots, \bar{F}_N(P)\big]\right) \\ &\qquad\qquad + \ d_1 h\left(\big[X_0(P)^{d_2}, \bar{G}_1(P), \ldots, \bar{G}_N(P)\big]\right) \\ &= d_2 h\big(\phi_1(P)\big) + d_1 h\big(\phi_2(P)\big). \end{aligned}$$

We combine this with (7.6) to obtain

$$d_1 d_2 h(P) + O(1) = h\big(\psi(P)\big) \le d_2 h\big(\phi_1(P)\big) + d_1 h\big(\phi_2(P)\big).$$

Dividing both sides by $d_1 d_2$ completes the proof of Theorem 7.14. $\square$

For regular affine automorphism, it is conjectured that the height inequality (7.4) in Theorem 7.14 may be replaced by a stronger estimate.

Conjecture 7.17. *Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be a regular affine automorphism. Then there is a constant $C = C(\phi)$ so that for all $P \in \mathbb{A}^N(\bar{\mathbb{Q}})$,*

$$\frac{1}{d_1}h\bigl(\phi(P)\bigr) + \frac{1}{d_2}h\bigl(\phi^{-1}(P)\bigr) \ge \left(1 + \frac{1}{d_1 d_2}\right) h(P) - C. \tag{7.7}$$

Kawaguchi [210] proves Conjecture 7.17 in dimension 2, i.e., for regular affine automorphisms $\phi : \mathbb{A}^2 \to \mathbb{A}^2$, see also [388]. However, for general jointly regular affine morphisms, it is easy to see that (7.4) cannot be improved, see Exercise 7.8. Kawaguchi also constructs canonical heights for maps that satisfy (7.7), see [210] and Exercises 7.17–7.22.

7.1.5 Boundedness of Periodic Points for Regular Automorphisms of $\mathbb{A}^N$

Theorem 7.14 applied to a regular affine automorphism ϕ and its inverse implies that at least one of $\phi(P)$ and $\phi^{-1}(P)$ has reasonably large height. This suffices to prove that the periodic points of ϕ form a set of bounded height, a result first demonstrated by Marcello [265, 266] (see also [121, 393]) using a height bound slightly weaker than the one in Theorem 7.14.

Theorem 7.18. (Marcello) *Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be a regular affine automorphism of degree at least 2 defined over $\bar{\mathbb{Q}}$. Then $\operatorname{Per}(\phi)$ is a set of bounded height in $\mathbb{A}^N(\bar{\mathbb{Q}})$. In particular,*

$$\operatorname{Per}(\phi) \cap \mathbb{A}^N(K) \textit{ is finite for all number fields } K.$$

Proof. Let

$$d_1 = \deg \phi \qquad \text{and} \qquad d_2 = \deg \phi^{-1}.$$

Applying Theorem 7.14 with $\phi_1 = \phi$ and $\phi_2 = \phi^{-1}$ yields the basic inequality

$$\frac{1}{d_1}h\bigl(\phi(P)\bigr) + \frac{1}{d_2}h\bigl(\phi^{-1}(P)\bigr) \ge h(P) - C, \tag{7.8}$$

where C is a constant depending on ϕ, but not on $P \in \mathbb{A}^N(\bar{\mathbb{Q}})$.

We prove the theorem initially under the assumption that $d_1 d_2 > 4$. Define a function

$$f(P) = \frac{1}{d_1}h(P) - \frac{1}{\alpha d_2}h\bigl(\phi^{-1}(P)\bigr) - \frac{C}{\alpha - 1}, \tag{7.9}$$

where the real number $\alpha > 1$ will be specified later. Then f satisfies

$$\begin{aligned} f\bigl(\phi(P)\bigr) - \alpha f(P) &= \left(\frac{1}{d_1}h\bigl(\phi(P)\bigr) - \frac{1}{\alpha d_2}h(P) - \frac{C}{\alpha-1}\right) \\ &\qquad\qquad - \alpha\left(\frac{1}{d_1}h(P) - \frac{1}{\alpha d_2}h\bigl(\phi^{-1}(P)\bigr) - \frac{C}{\alpha-1}\right) \\ &= \left(\frac{1}{d_1}h\bigl(\phi(P)\bigr) + \frac{1}{d_2}h\bigl(\phi^{-1}(P)\bigr)\right) - \left(\frac{\alpha}{d_1} + \frac{1}{\alpha d_2}\right)h(P) + C \\ &\ge \left(1 - \frac{\alpha}{d_1} - \frac{1}{\alpha d_2}\right)h(P) \qquad \text{from (7.8).} \end{aligned}$$

Hence if we take

$$\alpha = \frac{d_1 d_2 + \sqrt{(d_1 d_2)^2 - 4 d_1 d_2}}{2 d_2},$$

then

$$1 - \frac{\alpha}{d_1} - \frac{1}{\alpha d_2} = 0,$$

and our assumption that $d_1 d_2 > 4$ ensures that $\alpha > 1$, so for this choice of α we conclude that

$$f\bigl(\phi(P)\bigr) \ge \alpha f(P) \qquad \text{for all } P \in \mathbb{A}^N(\bar{\mathbb{Q}}).$$

Applying this estimate to the points $P, \phi(P), \phi^2(P), \dots, \phi^{n-1}(P)$, we obtain the fundamental inequality

$$f\bigl(\phi^n(P)\bigr) \ge \alpha^n f(P) \qquad \text{for all } P \in \mathbb{A}^N(\bar{\mathbb{Q}}) \text{ and all } n \ge 0. \tag{7.10}$$

Similarly, we define

$$g(P) = \frac{1}{d_2}h(P) - \frac{1}{\beta d_1}h\bigl(\phi(P)\bigr) - \frac{C}{\beta - 1} \tag{7.11}$$

and take

$$\beta = \frac{d_1 d_2 + \sqrt{(d_1 d_2)^2 - 4 d_1 d_2}}{2 d_1}.$$

Then an analogous calculation, which we leave to the reader, shows that g satisfies

$$g\bigl(\phi^{-1}(P)\bigr) \ge \beta g(P) \qquad \text{for all } P \in \mathbb{A}^N(\bar{\mathbb{Q}}),$$

from which we deduce that

$$g\bigl(\phi^{-n}(P)\bigr) \ge \beta^n g(P) \qquad \text{for all } P \in \mathbb{A}^N(\bar{\mathbb{Q}}) \text{ and all } n \ge 0. \tag{7.12}$$

We compute

$$\begin{aligned}
&\alpha^{-n} f\bigl(\phi^{n+1}(P)\bigr) + \beta^{-n} g\bigl(\phi^{-n-1}(P)\bigr) \\
&\qquad \ge f\bigl(\phi(P)\bigr) + g\bigl(\phi^{-1}(P)\bigr) \qquad \text{from (7.10) and (7.12),} \\
&\qquad = \left(\frac{1}{d_1} h\bigl(\phi(P)\bigr) - \frac{1}{\alpha d_2} h(P) - \frac{C}{\alpha - 1}\right) \\
&\qquad\qquad + \left(\frac{1}{d_2} h\bigl(\phi^{-1}(P)\bigr) - \frac{1}{\beta d_1} h(P) - \frac{C}{\beta - 1}\right) \\
&\qquad\qquad\qquad\qquad \text{from the definition (7.9) and (7.11) of } f \text{ and } g, \\
&\qquad \ge \left(1 - \frac{1}{\alpha d_2} - \frac{1}{\beta d_1}\right) h(P) - \left(1 + \frac{1}{\alpha - 1} + \frac{1}{\beta - 1}\right) C \qquad \text{from (7.8).}
\end{aligned}$$

Using the definition of f and g and rearranging the terms, we have proven the inequality

$$\begin{aligned}
&\frac{h\bigl(\phi^{n+1}(P)\bigr)}{\alpha^n d_1} + \frac{h\bigl(\phi^{-n-1}(P)\bigr)}{\beta^n d_2} + \frac{(\alpha\beta - 1)C}{(\alpha - 1)(\beta - 1)} \\
&\qquad \ge \left(1 - \frac{1}{\alpha d_2} - \frac{1}{\beta d_1}\right) h(P) + \frac{h\bigl(\phi^n(P)\bigr)}{\alpha^{n+1} d_2} + \frac{h\bigl(\phi^{-n}(P)\bigr)}{\beta^{n+1} d_1}. \qquad (7.13)
\end{aligned}$$

Now suppose that $P \in \mathbb{A}^N(\bar{\mathbb{Q}})$ is a periodic point for ϕ. Then $h\bigl(\phi^k(P)\bigr)$ is bounded independently of k, so letting $n \to \infty$ in (7.13) yields

$$\frac{(\alpha\beta - 1)C}{(\alpha - 1)(\beta - 1)} \ge \left(1 - \frac{1}{\alpha d_2} - \frac{1}{\beta d_1}\right) h(P),$$

where we are using the fact that $\alpha > 1$ and $\beta > 1$. Our assumption that $d_1 d_2 > 4$ also ensures that

$$1 - \frac{1}{\alpha d_2} - \frac{1}{\beta d_1} = \sqrt{1 - \frac{4}{d_1 d_2}} > 0,$$

so the height of P is bounded by a constant depending only on ϕ. This completes the proof of the first assertion of Theorem 7.18 under the assumption that $d_1 d_2 > 4$, and the second is immediate from Theorem 7.28(f), which says that for any given number field, $\mathbb{P}^N(K)$ contains only finitely many points of bounded height.

In order to deal with the case $d_1 d_2 \le 4$, i.e., $d_1 = d_2 = 2$, we use Theorem 7.10, which tells us that ϕ^2 is regular and has degree d_1^2. Similarly $\deg(\phi^{-2}) = d_2^2$. Hence from what we have already proven, the periodic points of ϕ^2 form a set of bounded height, and since it is easy to see that $\operatorname{Per}(\phi) = \operatorname{Per}(\phi^2)$, this completes the proof in all cases. $\square$

Remark 7.19. We observe that Theorem 7.18 applies only to regular maps. It cannot be true for all affine automorphisms, since there are affine automorphisms whose fixed (or periodic) points include components of positive dimension. For example, the affine automorphism $\phi(x, y) = (x, y + f(x))$ fixes all points of the form (a, b) satisfying $f(a) = 0$. Of course, this map ϕ is not regular, since one easily checks that

$$Z(\phi) = Z(\phi^{-1}) = \bigl\{[0, 1, 0]\bigr\}.$$

Definition. Let $\phi : V \to V$ be a morphism of a (not necessarily projective) variety V. A point $P \in \mathrm{Per}(\phi)$ is *isolated* if P is not in the closure of $\mathrm{Per}_n(\phi) \smallsetminus \{P\}$ for all $n \geq 0$. In particular, if $\mathrm{Per}_n(\phi)$ is finite for all n, then every periodic point is isolated.

Conjecture 7.20. *Let $\phi : \mathbb{A}^N \to \mathbb{A}^N$ be an affine automorphism of degree at least 2 defined over $\bar{\mathbb{Q}}$. Then the set of isolated periodic points of ϕ is a set of bounded height in $\mathbb{A}^N(\bar{\mathbb{Q}})$.*

A classification theorem of Friedland and Milnor [162] says that every automorphism $\phi : \mathbb{A}^2 \to \mathbb{A}^2$ of the affine plane is conjugate to a composition of elementary maps and Hénon maps. Using this classification, Denis [121] proved Conjecture 7.20 in dimension 2. (See also [265, 266].)

7.2 Primer on Algebraic Geometry

In this section we summarize basic material from algebraic geometry, primarily having to do with the theory of divisors, linear equivalence, and the divisor class group (Picard group). This theory is used to describe the geometry of algebraic varieties and the geometry of the maps between them. We assume that the reader is familiar with basic material on algebraic varieties as may be found in any standard textbook, such as [169, 180, 181, 187].

This section deals with geometry, so we work over an algebraically closed field. Let

$$\begin{aligned} K &= \text{an algebraically closed field,} \\ V &= \text{a nonsingular irreducible projective variety defined over } K, \\ K(V) &= \text{the field of rational functions on } V. \end{aligned}$$

7.2.1 Divisors, Linear Equivalence and the Picard Group

In this section we recall the theory of divisors, linear equivalence, and the divisor class group (Picard group).

Definition. A *prime divisor* on V is an irreducible subvariety $W \subset V$ of codimension 1. The *divisor group of* V, denoted $\mathrm{Div}(V)$, is the free abelian group generated by the prime divisors on V. Thus $\mathrm{Div}(V)$ consists of all formal sums

$$\sum_W n_W W,$$

where the sum is over prime divisors $W \subset V$, the coefficients n_W are integers, and only finitely many n_W are nonzero. The *support* of a divisor $D = \sum n_W W$ is

$$|D| = \bigcup_{\substack{W \text{ with} \\ n_W \neq 0}} W.$$