Ann. Sc. Norm. Super. Pisa Cl. Sci. (5)
Vol. XVII (2017), 187-204

On a dynamical version of a theorem of Rosenlicht

JASON BELL, DRAGOS GHIOCA AND ZINOVY REICHSTEIN

Abstract. Consider the action of an algebraic group G on an irreducible algebraic variety X defined over a field k. M. Rosenlicht showed that orbits in general position in X can be separated by rational invariants. We prove a dynamical analogue of this theorem, where G is replaced by a semigroup of dominant rational maps $X \dashrightarrow X$. Our semigroup G is not required to have the structure of an algebraic variety and can be of arbitrary cardinality. This generalizes earlier work of E. Amerik and F. Campana, where $k = \mathbb{C}$ and the semigroup G is assumed to be generated by a single endomorphism.

Mathematics Subject Classification (2010): 14E05 (primary); 14C05, 37F10 (secondary).

1. Introduction

Throughout this paper we will work over a base field k. By a *k-variety* we will mean a separated reduced scheme of finite type over k. By an irreducible k-variety we shall mean "irreducible over k", not necessarily absolutely irreducible.

Our starting point is the following classical theorem of M. Rosenlicht [15, Theorem 2], [16].

Theorem 1.1. *Consider the action of a smooth algebraic group G on an irreducible algebraic variety X defined over a field k.*

(a) *There exists a G-invariant dense open subvariety $X_0 \subset X$ and a G-equivariant morphism $\phi\colon X_0 \to Z$ (where G acts trivially on Z), with the following property. For any field extension K/k and any K-point $x \in X_0(K)$, the orbit $G \cdot x$ equals the fiber $\phi^{-1}(\phi(x))$;*
(b) *moreover, for any Z as in part* (a)*, the field of invariants $k(X)^G$ is a purely inseparable extension of $\phi^* k(Z)$, and one can choose Z and ϕ so that $\phi^* k(Z) = k(X)^G$ (in characteristic zero, this is automatic).*

The authors have been partially supported by Discovery Grants from the National Science and Engineering Research Council of Canada.
Received January 11, 2015; accepted in revised form October 20, 2015.
Published online March 2017.

In short, *for points x, y in general position in X, distinct G-orbits $G \cdot x$ and $G \cdot y$ can be separated by rational G-invariant functions*. In particular, G-orbits in X_0 are closed in X_0. The rational map $\phi\colon X \dashrightarrow Z$, with $k(Z) = k(X)^G$, is unique up to birational isomorphism. It is called *the rational quotient* for the G-action on X. See [14, Chapter 2] for details on this construction and its applications, [11] for computational aspects, and [13] for a generalization to actions of infinite-dimensional algebraic groups.

The purpose of this note is to prove a dynamical version of this result, where the algebraic group G is replaced by a semigroup of dominant rational maps $X \dashrightarrow X$. Here the semigroup G is not required to have the structure of an algebraic variety, and can be of arbitrary cardinality. Our main result is Theorem 1.2 below.

Theorem 1.2. *Let k be a field, X be an irreducible quasi-projective k-variety, and G be a semigroup of dominant rational k-maps $X \dashrightarrow X$.*

Then there exists a dense open subvariety X_0, a countable collection of closed G-invariant subvarieties $Y_1, Y_2, \dots \subsetneq X_0$ and a dominant morphism $\phi\colon X_0 \to Z$ with the following properties:

(a) *let K/k be a field extension and $x, y \in X_0(K)$ be K-points which do not lie in the indeterminacy locus of any $g \in G$, or on Y_i for any $i \geqslant 1$. Then $\phi(x) = \phi(y)$ if and only if $\overline{G \cdot x} = \overline{G \cdot y}$ in X_K;*
(b) *$\phi \circ g = \phi$, as rational maps $X \dashrightarrow Z$, for any $g \in G$;*
(c) *for any dense open subvariety $X_0 \subset X$ and a dominant morphism $\phi\colon X_0 \to Z$ satisfying* (a) *and* (b), *the field of invariants $k(X)^G$ is a purely inseparable extension of $\phi^* k(Z)$. Moreover, one can choose X_0, Z and ϕ so that $\phi^* k(Z) = k(X)^G$ (in characteristic zero, this is automatic);*
(d) *each Y_d is G-invariant in the following sense. Suppose K/k is a field extension and $x \in Y_d(K)$ does not lie in the indeterminacy locus of any $g \in G$. Then $g(x) \in Y_d(K)$ for every $g \in G$.*

Furthermore, if G is a monoid (i.e., *contains the identity morphism $X \to X$) then:*

(e) *X_0 can be chosen to be g-invariant for every $g \in G$ which is an automorphism of X* (i.e., *g^{-1} exists in G, and g, g^{-1} are both regular);*
(f) *if $x \in X_0(K)$ is as in part* (a), *then the fiber $\phi^{-1}(\phi(x))$ of x in X_0 equals $\overline{(G \cdot x)} \cap X_0$.*

In short, *for points x, y in very general position in X, distinct orbit closures $\overline{G \cdot x}$ and $\overline{G \cdot y}$ can be separated by rational G-invariant functions*. Here, as usual, "very general position" means "away from a countable union of proper subvarieties".

Note that the G-orbit of $x \in X(K)$ in the setting of Theorem 1.1 is a K-subvariety of X; it is defined as the image of the orbit map $G \to X$, taking $g \in G$ to $g \cdot x$. In the dynamical setting of Theorem 1.2 the orbit $G \cdot x$ is just a collection of K-points $g \cdot x$, as g ranges over G. The closure $\overline{G \cdot x}$ is a K-subvariety of X in both cases. Several remarks are in order:

(1) in the case where $k = \mathbb{C}$ and the semigroup G is generated by a single dominant rational map $X \dashrightarrow X$, Theorem 1.2 was proved by E. Amerik and F. Campana. Their result, [1, Theorem 4.1], is stated more generally, in the setting of Kähler manifolds. Theorem 1.2 was motivated by our attempt to find a purely algebraic characteristic-free proof of this result;

(2) the idea behind our construction of the map ϕ in Theorem 1.2 is as follows. Assuming $X \subset \mathbb{P}^n$, we set $\phi(x)$ to be to the class of the orbit closure $\overline{G \cdot x} \subset \mathbb{P}^n$ in the Hilbert scheme $\mathrm{Hilb}(n)$ of subschemes of $\mathbb{P}^n$. The challenge is to show that this defines a rational map

$$\phi\colon X \dashrightarrow \mathrm{Hilb}(n).$$

The "quotient variety" Z will then be defined as the closure of the image of this map in $\mathrm{Hilb}(n)$. This argument is in the same spirit as the proofs of Theorem 1.1 in [15] and of [1, Theorem 4.1], with Hilbert schemes replacing Chow varieties or Barlet spaces used in these earlier proofs. To further illustrate our approach, we give a short proof of Theorem 1.1(a) using the Hilbert scheme in the last section;

(3) a conjecture of A. Medvedev and T. Scanlon [12, Conjecture 7.14] asserts that in the case where k is algebraically closed of characteristic 0, G is generated by a single regular endomorphism $X \to X$ and $k(X)^G = k$ (*i.e.*, Z is a point), X has a k-point with a dense G-orbit. (See also [2, Conjecture 7].) Over $\mathbb{C}$ the Medvedev-Scanlon conjecture follows from the above-mentioned [1, Theorem 4.1]. In the case where k is an algebraically closed uncountable field of arbitrary characteristic, it was proved by the first author, D. Rogalski and S. Sierra [4, Theorem 1.2]. Corollary 6.1 below may be viewed as a strengthening of [4, Theorem 1.2].
Over a countable field, the Medvedev-Scanlon conjecture (which was, in turn, motivated by an earlier related conjecture of S.-W. Zhang [18, Conjecture 4.1.6]) remains largely open. It has been established only in a small number of special instances (see, in particular, [12, Theorem 7.16] and [3, Theorem 1.3]), and no counterexamples are known;

(4) if the semigroup G is not assumed to be countable, then the points $x, y \in X_0(K)$ in part (a) are not truly in very general position, since they are required to lie away from the indeterminacy loci of (possibly uncountably many) elements $g \in G$. This requirement is imposed to make sure that the orbits $G \cdot x$ and $G \cdot y$ are well defined. In Section 5 we will prove a variant of Theorem 1.2, where the orbit $G \cdot x$ is defined more generally, as

$$G \cdot x := \{g(x) \mid g \in G \text{ is defined at } x\}. \tag{1.1}$$

With this definition, we will show that rational G-invariant functions on X separate orbit closures in very general position, even if G is uncountable; see Proposition 5.1. Our proof is based on replacing G by a suitable countable subsemigroup H. Note that the new "exceptional subvarieties" of X_0, resulting from replacing G by H (which we denote by W_i in the statement of

Proposition 5.1) are generally bigger than the subvarieties Y_i in the statement of Theorem 1.2 and are only H-invariant, not necessarily G-invariant;

(5) in the case where X is not quasi-projective, Theorem 1.2 and Proposition 5.1 can be applied to a quasi-projective dense open subvariety $X' \subset X$. Note however that replacing X by X', and thus viewing elements of G as dominant rational maps $X' \dashrightarrow X'$, may make the condition on $x, y \in X_0(K)$ in part (a) more stringent by enlarging the indeterminacy loci of these rational maps;

(6) examples 6.4 and 6.5 show that if we replace the countable collection of $\{Y_i,\ i \geq 1\}$ of proper subvarieties of X by a finite collection, Theorem 1.2 will fail, even in the simplest case, where the semigroup G is generated by a single dominant morphism $\sigma\colon X \to X$. Note that in Example 6.4, σ is an automorphism.

ACKNOWLEDGEMENTS. The authors are grateful to E. Amerik, M. Borovoi, B. Poonen, T. Scanlon and T. Tucker for helpful comments.

2. A dense collection of rational sections

In this section we will consider the following situation. Let X be an irreducible k-variety, V be a closed subvariety of $X \times \mathbb{P}^n$, and $\pi\colon V \to X$ be the projection to the first factor, and $s_\lambda\colon X \dashrightarrow V$ be a collection of rational sections $X \dashrightarrow V$, indexed by a set Λ. We will denote the scheme-theoretic fiber $\pi^{-1}(x)$ of a point $x \in X$ by V_x.

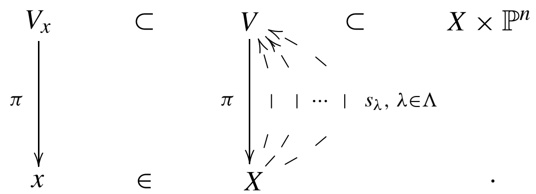

Note that we do not assume that V is irreducible and do not impose any restrictions on the cardinality of Λ.

For notational simplicity, we will sometimes identify $\{x\} \times \mathbb{P}^n$ with $\mathbb{P}^n$ and thus think of V_x as a closed subscheme of $\mathbb{P}^n$. Similarly, since each s_λ is of the form $x \mapsto (x, s'_\lambda(x))$ for some rational map $s'_\lambda\colon X \to \mathbb{P}^n$, we will sometimes, by a slight abuse of notation, identify s_λ with s'_λ and view s_λ as a rational map $X \dashrightarrow \mathbb{P}^n$.

If K/k is a field extension, we will denote by $X(K)'$ the collection of K-points of X lying away from the indeterminacy locus of s_λ, for every $\lambda \in \Lambda$. In other words, for $x \in X(K)'$, $s_\lambda(x)$ is defined for every $\lambda \in \Lambda$. Note that if Λ is large enough, $X(K)'$ may be empty for some fields K/k, even if K is algebraically closed. On the other hand, the generic point η of X lies in $X(K_{\text{gen}})'$, where $K_{\text{gen}} = k(X)$ is the function field of X.

Proposition 2.1. *Assume that the union of $s_\lambda(X)$ over all $\lambda \in \Lambda$ is dense in V. Then there exists a countable collection $\{Y_i,\ i \geqslant 1\}$ of proper subvarieties of X with the following property: for any field extension K/k and $x \in X(K)'$ away from $\bigcup_{i=1}^{\infty} Y_i$, the set*

$$\{s_\lambda(x) \mid \lambda \in \Lambda\}$$

is Zariski dense in the fiber $V_x := \pi^{-1}(x)$.

Proof. By generic flatness (see, *e.g.*, [5, Theorem 14.4]), after replacing X by a dense open subvariety, we may assume that π is flat. Let us denote the Hilbert polynomial of the fiber V_x by p_{V_x}. Since π is flat, p_{V_x} is independent of the choice of $x \in X$. In particular, if η is the generic point of X, then

$$p_{V_x} = p_{V_\eta} \tag{2.1}$$

Note that if $I(V_x) \subset K[t_0, \dots, t_n]$ is the homogeneous ideal of V_x in $\mathbb{P}^n$, and $I(V_x)[d]$ is the K-vector space of homogeneous polynomials of degree d in $I(V_x)$, then

$$p_{V_x}(d) := \binom{n+d}{d} - \dim_K(I(V_x)[d]) \text{ for } d \gg 0.$$

Thus for $d \gg 0$, $\dim_K(I(V_x)[d])$ depends only on d and not on the choice of a field extension K/k or a point $x \in X(K)$.

Let K/k be a field and $x \in X(K)'$. Denote by W_x the closure of $\{s_\lambda(x) \mid \lambda \in \Lambda\}$ in V_x. Clearly $W_x \subset V_x \subset \mathbb{P}^n_K$ and thus $I(V_x) \subset I(W_x)$. We want to show that for $x \in X(K)'$ in very general position, $W_x = V_x$. Our first step towards this goal is the following simple lemma.

Lemma 2.2. *Let A and B be closed subschemes of the projective space $\mathbb{P}^n$. If $B \subset A$ and A and B have the same Hilbert polynomial, then $A = B$.*

Proof. Assume the contrary. Then there exists a homogeneous polynomial $r(t_0, \dots, t_n)$ such that r is identically 0 on B but not on A. Let $d := \deg(r)$. Choose a linear form $l(t_0, \dots, t_n)$ such that no power of l is identically 0 on $A \setminus Z$, where Z is the hypersurface in $\mathbb{P}^n$ cut out by r. (Note that we can always choose $l = t_j$ for some $j = 0, \dots, n$.) Then $l^i r$ lies in $I(B)[d+i]$ but not in $I(A)[d+i]$ for every $i \geqslant 0$. Hence $\dim(I(B)[d+i]) > \dim(I(A)[d+i])$ for every $i \geqslant 0$, contradicting our assumption that A and B have the same Hilbert polynomial. □

Proposition 2.1 now reduces to the following:

Claim 2.3. For every $d \geqslant 1$ there exists a proper closed subvariety $Y_d \subset X$ such that $\dim(I(W_x)[d]) = \dim(I(V_\eta))[d]$ for any field K/k and any $x \in X(K)'$ away from Y_d.

Indeed, Claim 2.3 tells us that, for $x \in X(K)'$ away from $Y_1 \cup Y_2 \cup \dots$, we have $p_{W_x} = p_{V_\eta}$. Combining this with (2.1), we obtain $p_{W_x} = p_{V_x}$. Since $W_x \subset V_x$, Lemma 2.2 now tells us that $W_x = V_x$, as desired.

The rest of the proof will be devoted to establishing Claim 2.3. We begin by observing that it suffices to prove this claim for $d = 1$. Indeed, if we settle this case, we will be able to deduce Claim 2.3 for any $d \geqslant 1$, after replacing V by its image under

$$(\mathrm{id}, \mathrm{Ver}_d)\colon X \times \mathbb{P}^n \to X \times \mathbb{P}^N ,$$

where $N := \binom{n+d}{d}$ and Ver_d is the d-fold Veronese embedding.

In the case where $d = 1$, set $r := p_{V_\eta}(1) = n + 1 - \dim(I(V_\eta)[1])$. Claim 2.3 now reduces to

Claim 2.4. Let $\eta \in X(k(X))$ be the generic point of X. Suppose the linear span of the $k(X)$-points $\{s_\lambda(\eta) \mid \lambda \in \Lambda\}$ in $\mathbb{P}^n_{k(X)}$ is of dimension r. Then there exists a closed subvariety $Y_1 \subset X$ such that for any field extension K/k and any $x \in X(K)'$ away from Y_1, the linear span of the K-points $\{s_\lambda(x) \mid \lambda \in \Lambda\}$ is of dimension r.

To prove Claim 2.4, write $\mathbb{P}^n$ as $\mathbb{P}(\mathcal{V})$, where $\mathcal{V}$ is the underlying $(n + 1)$-dimensional vector space and let

$$\Delta := \left\{ ([v_1], \dots, [v_r]) \in (\mathbb{P}^n)^r \mid v_1 \wedge \cdots \wedge v_r = 0 \text{ in } \bigwedge^r(\mathcal{V}) \right\}.$$

For each $(\lambda_1, \dots, \lambda_r) \in \Lambda^r$, let $s_{\lambda_1,\dots,\lambda_r} \colon X \dashrightarrow (\mathbb{P}^n)^r$ be given by

$$x \mapsto (s_{\lambda_1}(x), \dots, s_{\lambda_r}(x)).$$

Let $Z_{\lambda_1,\dots,\lambda_r} \subset X$ be the union of the indeterminacy loci of $s_{\lambda_1}, \dots, s_{\lambda_r}$ and

$$U_{\lambda_1,\dots,\lambda_r} := X \setminus Z_{\lambda_1,\dots,\lambda_r} .$$

Then $s_{\lambda_1,\dots,\lambda_r} \colon U_{\lambda_1,\dots,\lambda_r} \to (\mathbb{P}^n)^r$ is a regular map, and $s^{-1}_{\lambda_1,\dots,\lambda_r}(\Delta)$ is a closed subvariety of $U_{\lambda_1,\dots,\lambda_r}$. Note that in some cases $s^{-1}_{\lambda_1,\dots,\lambda_r}(\Delta) = U_{\lambda_1,\dots,\lambda_r}$. This will happen if and only if $s_{\lambda_1,\dots,\lambda_r}(\eta) \in \Delta$. For example, $s^{-1}_{\lambda_1,\dots,\lambda_r}(\Delta) = U_{\lambda_1,\dots,\lambda_r}$ whenever $\lambda_i = \lambda_j$ for some $i \neq j$.

By our assumption there exist $\alpha_1, \dots, \alpha_r \in \Lambda$ such that $s_{\alpha_1,\dots,\alpha_r}(\eta) \notin \Delta$. Then

$$Y_1 := Z_{\alpha_1,\dots,\alpha_r} \cup s^{-1}_{\alpha_1,\dots,\alpha_r}(\Delta)$$

is a proper closed subvariety of X with the desired property. This completes the proof of Claim 2.4 and thus of Claim 2.3 and Proposition 2.1. □

Remark 2.5. In the sequel it will be convenient for us to define Y_1 more symmetrically as follows:

$$Y_1 := \bigcap_{(\lambda_1,\dots,\lambda_r) \in \Lambda^r} (Z_{\lambda_1,\dots,\lambda_r} \cup s^{-1}_{\lambda_1,\dots,\lambda_r}(\Delta)) .$$

Note also that Claim 2.4 concerns only field-valued points of Y_1. Thus the scheme structure of $s^{-1}_{\lambda_1,\dots,\lambda_r}(\Delta)$ does not make a difference here; we are only interested in the underlying variety.

3. The Hilbert scheme and proof of Theorem 1.2(a), (b), (d), (e), (f)

The Hilbert scheme Hilb(n), constructed by A. Grothendieck [10], classifies closed subschemes of $\mathbb{P}^n$ in the following sense. A family of subschemes of $\mathbb{P}^n$ parametrized by a scheme X is, by definition, a closed subscheme

$$V \subset X \times \mathbb{P}^n$$

such that the projection $\pi\colon V \to X$ to the first factor is flat. As we mentioned in the previous section, if X is irreducible, then every fiber of π has the same Hilbert polynomial. Families of subschemes of $\mathbb{P}^n$ parametrized by X are in a natural (functorial in X) bijective correspondence with morphisms $X \to \mathrm{Hilb}(n)$. Note that Hilb(n) is not a Noetherian scheme; it is a disjoint union of (infinitely many) schemes of the form Hilb(n, p), where p is a fixed Hilbert polynomial. Each Hilb(n, p) is a projective variety defined over $\mathbb{Z}$; it parametrizes families of subschemes of $\mathbb{P}^n$ with Hilbert polynomial p.

We are now ready to proceed with the proof of Theorem 1.2. In this section we will construct a dense open subvariety $X_0 \subset X$, a dominant morphism $\phi\colon X_0 \to Z$, and a countable collection of proper k-subvarieties $Y_d \subset X_0$. We will prove parts (a), (b), (d), (e) and (f) of Theorem 1.2 and defer the proof of part (c) to the next section.

By our assumption X is a quasi-projective variety. In other words, X is a locally closed subvariety of some projective space $\mathbb{P}^n_k$. Let $V \subset X \times \mathbb{P}^n$ be the Zariski closure of the union of the graphs of $g\colon X \dashrightarrow X \subset \mathbb{P}^n$, as g ranges over G. Let $\pi\colon V \to X$ be the projection

$$\begin{array}{ccc} V := \overline{\{(x, g(x)) \mid x \in X,\ g \in G\}} & \subset X \times \mathbb{P}^n \\ \big\downarrow \pi & \\ X & \end{array} \tag{3.1}$$

to the first factor and $X_0 \subset X$ be the flat locus of π, *i.e.* the largest dense open subvariety of X over which π is flat, and $V_0 := \pi^{-1}(X_0)$. (Recall that X_0 is dense in X by generic flatness.) We now view $V_0 \subset X_0 \times \mathbb{P}^n$ as a family of subschemes of $\mathbb{P}^n$ parametrized by X_0. By the universal property of the Hilbert scheme this family induces a morphism $\phi\colon X_0 \to \mathrm{Hilb}(n)$. Denote the closure of the image of this morphism by Z. If K/k is a field extension and $x, y \in X_0(K)$ then by our construction

$$\phi(x) = \phi(y) \text{ if and only if } V_x = V_y. \tag{3.2}$$

Here we identify $\{x\} \times \mathbb{P}^n_K$ and $\{y\} \times \mathbb{P}^n_K$ with $\mathbb{P}^n_K$.

We may view each $g \in G$ as a rational section $X_0 \to V_0$ given by $x \mapsto (x, g(x))$, as in the previous section. By the definition of V, the union of the images of these sections is dense in V_0. Thus by Proposition 2.1 there exists a countable collection of proper k-subvarieties $Y_i \subset X_0$, such that for any field extension K/k and any $x \in X_0(K)$ away from the union of these subvarieties,

$$V_x = \overline{G \cdot x} \text{ in } \mathbb{P}^n_K. \tag{3.3}$$

Choose $Y_1, Y_2, \ldots$ using the formula in Remark 2.5.

(a) By (3.2), $\phi(x) = \phi(y)$ if and only if $V_x = V_y$ in $\mathbb{P}^n_K$. By (3.3), and $V_x = \overline{G \cdot x}$, $V_y = \overline{G \cdot y}$, where the closure is taken in $\mathbb{P}^n_K$. This shows that $\phi(x) = \phi(y)$ if and only if $G \cdot x$ and $G \cdot y \subset X_K$ have the same closure in $\mathbb{P}^n_K$. On the other hand, $G \cdot x$ and $G \cdot y$ have the same closure in $\mathbb{P}^n_K$ if and only if they have the same closure in X_K.

(b) It suffices to show that the rational maps $\phi \circ g$ and $\phi\colon X \dashrightarrow \mathrm{Hilb}(n)$ agree on the generic point η of X for every $g \in G$. Choose $g \in G$ and fix it for the rest of the proof. Then η and $\mu := g(\eta)$ are K_{gen}-points of X, where $K_{\text{gen}} := k(X)$. Since g is dominant, neither η nor μ lie on any proper subvariety of X defined over k. In particular, they do not lie in the indeterminacy locus of any $h \in G$ or on Y_i for any $i \geqslant 1$. By (3.2), proving that $\phi(\eta) = \phi(\mu)$ is equivalent to showing that

$$V_\eta = V_\mu \text{ in } \mathbb{P}^n_{K_{\text{gen}}}, \tag{3.4}$$

where $V_\eta, V_\mu \in \mathbb{P}^n_{K_{\text{gen}}}$ are the fibers of η and μ, respectively, under $\phi\colon V \to X$ in $\mathbb{P}^n_{K_{\text{gen}}}$.

Since $\eta, \mu \in X_0(K_{\text{gen}})$ do not lie on Y_i for any $i \geqslant 1$, part (a) tells us that $V_\eta = \overline{G \cdot \eta}$ and $V_\mu = \overline{G \cdot \mu}$. Now recall that $\mu = g(\eta)$. Thus $G \cdot \mu \subset G \cdot \eta$ and consequently, $V_\mu \subset V_\eta$. Since $\pi\colon V \to X$ is flat over X_0, the fibers V_η and V_μ have the same Hilbert polynomial, and (3.4) follows from Lemma 2.2.

(d) Once again, after replacing X by its image under the Veronese embedding Ver_d, we may assume without loss of generality that $d = 1$, as in the proof of Proposition 2.1. By our assumption g is regular at $x \in Y_1(K)$ for every $g \in G$. Thus the condition that $x \in Y_1$ is equivalent to $(g_1(x), \ldots, g_r(x)) \in \Delta$ for every $g_1, \ldots, g_r \in G$; see the formula for Y_1 in Remark 2.5. We want to show that $g(x) \in Y_1$, *i.e.*,

$$g(x) \in Z_{g_1,\ldots,g_r} \cup s^{-1}_{g_1,\ldots,g_r}(\Delta)$$

for every $g_1, \ldots, g_r \in G$ (once again, see the formula in Remark 2.5).

Choose a particular r-tuple $g_1, \ldots, g_r \in G$. If one of the endomorphisms $g_1, \ldots, g_r$ is not defined at $g(x)$, then $g(x) \in Z_{g_1,\ldots,g_r}$, and we are done. On the other hand, if $g(x)$ lies in the domain of every g_i, then $g_i(g(x)) = (g_i \circ g)(x)$ for $i = 1, \ldots, r$. Since $x \in Y_1$, we have $(g_1 \circ g, \ldots, g_r \circ g)(x) \in \Delta$. In other words, $g(x) \in s^{-1}_{g_1,\ldots,g_r}(\Delta)$, as desired.

(e) If g is an automorphism of X, then the variety $V \subset X \times \mathbb{P}^n$ defined in (3.1), is g-invariant, where g acts on $X \times \mathbb{P}^n$ via the first factor. Consequently, the projection $\pi\colon V \to X$ is g-equivariant, and the flat locus $X_0 \subset X$ of π is invariant under g.

(f) By part (b), $\phi((G \cdot x) \cap X_0) = \phi(x)$ and thus

$$\overline{(G \cdot x) \cap X_0} \subset \phi^{-1}(\phi(x)), \tag{3.5}$$

where the closure is taken in X_0. If $y \in X_0$ and $\phi(y) = \phi(x)$, then by (3.2), $V_y = V_x$. Since G is a monoid, $y \in V_y$ and thus $y \in V_x$. In other words, $\phi^{-1}(\phi(x)) \cap X_0 \subset V_x \cap X_0$. Combining this with (3.5), we obtain

$$\overline{(G \cdot x) \cap X_0} \subset \phi^{-1}(\phi(x)) \subset V_x \cap X_0\,. \tag{3.6}$$

On the other hand, by (3.3), $G \cdot x$ is dense in V_x. Thus $(G \cdot x) \cap X_0$ is dense in $V_x \cap X_0$, *i.e.*,

$$\overline{(G \cdot x) \cap X_0} = V_x \cap X_0\,.$$

We conclude that both inclusions in (3.6) are equalities. In particular, $\overline{(G \cdot x) \cap X_0} = \phi^{-1}(\phi(x))$, as desired.

4. Proof of Theorem 1.2(c)

By part (b), $\phi^*(k(Z)) \subset k(X)^G$. Let Y be a k-variety whose function field $k(Y)$ is $k(X)^G$. The inclusions $k(Z) \overset{i^*}{\hookrightarrow} k(Y) = k(X)^G \overset{\psi^*}{\hookrightarrow} k(X) = k(X_0)$ induce dominant rational maps

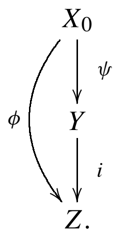

After replacing X_0, Z, and Y by suitable dense open subvarieties, we may assume that all three maps in the above diagram are regular.

Let K/k be a field and $x, y \in X_0(K)$ be as in Theorem 1.2(a). We claim that if $\phi(x) = \phi(y)$ then $\psi(x) = \psi(y)$. Indeed, by Theorem 1.2(a), $\phi(x) = \phi(y)$ implies that $\overline{G \cdot x} = \overline{G \cdot y}$ in X_K. Then

$$\overline{(G \cdot x) \cap X_0} = \overline{(G \cdot y) \cap X_0} \quad \text{in} \quad X_0. \tag{4.1}$$

By our construction, ψ is G-equivariant, where G acts trivially on Y. Thus ψ sends all of $\overline{(G \cdot x) \cap X_0}$ to the point $\psi(x)$ and all of $\overline{(G \cdot y) \cap X_0}$ to the point $\psi(y)$ in Y. By (4.1), $\psi(x) = \psi(y)$, as claimed.

In particular, if K/k is field extension and $x, y \colon \operatorname{Spec}(K) \to X$ are dominant points, then x and y satisfy the conditions of Theorem 1.2(a) and thus

$$\phi(x) = \phi(y) \text{ if and only if } \psi(x) = \psi(y).$$

By Lemma 4.1 below, i is purely inseparable. This proves the first assertion of Theorem 1.2(c). To prove the second assertion of part (c), we simply replace Z by Y and ϕ by ψ.

It remains to prove:

Lemma 4.1. *Let* $\phi\colon X \xrightarrow{\psi} Y \xrightarrow{i} Z$ *be dominant maps of irreducible k-varieties. Suppose that for any pair of dominant points* $x, x'\colon \operatorname{Spec}(K) \to X$*, where* K/k *is a field extension,*

$$\phi(x) = \phi(x') \text{ if and only if } \psi(x) = \psi(x'). \tag{4.2}$$

Then the field extension $k(Y)/i^*(k(Z))$ *is purely inseparable.*

Proof. Let F be the algebraic closure of $k(X)$ and $x\colon \operatorname{Spec}(F) \to X$ be the dominant F-point of X obtained by composing the natural projection $\operatorname{Spec}(F) \to \operatorname{Spec}(k(X))$ with the generic point $\operatorname{Spec}(k(X)) \to X$. Set $z := \phi(x)$. The fiber $\phi^{-1}(z)$ is an F-subvariety of X_F. Denote its irreducible components by $X_1, \dots, X_n$.

By the fiber dimension theorem [9, Théorème 13.1.3], the generic point $x_i\colon \operatorname{Spec}(F(X_i)) \to X_i \hookrightarrow X$ is dominant for every $i = 1, \dots, n$. If K/F is a compositum of $F(X_1), \dots, F(X_n)$ over F and $(x_i)_K$ is the composition of the projection $\operatorname{Spec}(K) \to \operatorname{Spec}(F(X_i))$ with x_i, then

$$\phi((x_1)_K) = \cdots = \phi((x_n)_K) = z_K = \phi(x_K).$$

Our assumption (4.2) now tells us that $\psi((x_1)_K) = \cdots = \psi((x_n)_K) = \psi(x_K)$. Since x is, by definition, an F-point of X, we see that $\psi((x_1)_K) = \ldots = \psi((x_n)_K) = \psi(x)$ descends to a dominant F-point $y\colon \operatorname{Spec}(F) \to Y$, where $i(y) = z$. In other words, ψ maps each X_i to the single point $y \in Y(F)$, as depicted in the following diagram:

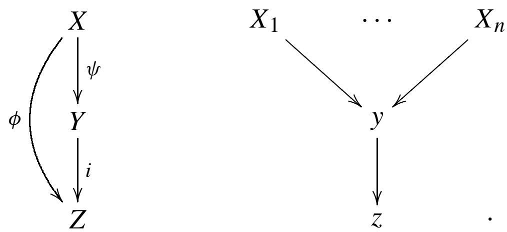

Thus $\psi(\phi^{-1}(z)) = y$. Equivalently, $\psi(\psi^{-1}(i^{-1}(z))) = y$ or $i^{-1}(z) = y$. Applying the fiber dimension theorem one more time, we obtain $\dim(Y) = \dim(Z)$. Since $z\colon \operatorname{Spec}(F) \to Z$ is dominant, and F is algebraically closed, the number of preimages of z in Y equals the separability degree of $k(Y)$ over $i^*(k(Z))$. In our case the preimage of z is a single point $y \in Y(F)$; hence, $k(Y)$ is purely inseparable over $i^*k(Z)$. This completes the proof of Lemma 4.1 and thus of Theorem 1.2. □

Remark 4.2. We do not know whether or not $\phi^*k(Z)$ always coincides with $k(X)^G$, where Z is the closure of the image of the rational map $\phi\colon X \dashrightarrow \operatorname{Hilb}(n)$ we constructed. As we have just seen, this is always the case in characteristic zero, so the question is only of interest in prime characteristic. An analogous question in the context of Theorem 1.1 was left open in [15] and was subsequently settled in the positive by A. Seidenberg in [17].

5. Passing to a countable subsemigroup

The purpose of this section is to prove a refinement of Theorem 1.2, which shows that ϕ will separate orbit closures in very general position, even if G is uncountable; see remark (3) in the Introduction. The price we will pay for this strengthening of Theorem 1.2 is that the new "exceptional subvarieties" $W_i \subsetneq X_0$ may no longer be G-invariant in the sense of Theorem 1.2(d). Note also that Proposition 5.1 below is only of interest if G is uncountable. Otherwise, we can take $H = G$, and Proposition 5.1 will be subsumed by Theorem 1.2.

Proposition 5.1. *Let X be a quasi-projective irreducible k-variety, and G be a semigroup of dominant rational maps $X \dashrightarrow X$. Under the assumptions of Theorem* 1.2*, with $k(Z) = k(X)^G$, G has a countable subsemigroup H with the following properties. There exists a countable collection $\{W_i \mid i \geqslant 1\}$ of subvarieties of X_0, such that each $W_i \subsetneq X_0$ is H-invariant in the sense of Theorem* 1.2(d)*, and*

(a) $k(X)^G = k(X)^H$;
(b) *for any field extension F/k and any $x, y \in X_0(F)$ away from the (countably many) indeterminacy loci of elements of H and away from $\cup_{i\geqslant 1} W_i$, $\phi(x) = \phi(y)$ if and only if $\overline{H \cdot x} = \overline{H \cdot y}$ in X_F;*
(c) *moreover, if $x \in X_0(F)$ as in part (b), $\overline{G \cdot x} = \overline{H \cdot x}$, where $G \cdot x$ is defined as in* (1.1).

Our proof will rely on the following elementary lemma.

Lemma 5.2. *Let W be an algebraic variety (not necessarily irreducible), and S be a dense collection of points in W. Then there exists a countable subcollection $S' \subset S$ which is dense in W.*

Proof. After replacing W by a dense open subvariety, and removing the points of S that do not lie in this dense open subvariety we may assume without loss of generality that $W \subset \mathbb{A}^n$ is affine. Let $I(W)$ be the ideal of W in $k[x_1, \dots, x_n]$ and $I(W)[d]$ be the finite-dimensional vector space of polynomials of degree $\leq d$ vanishing on W. It is easy to see that for each d there is a finite subset $S_d \subset S$ such that $I(S_d)[d] = I(W)[d]$. Taking $S' = \cup_{d=1}^{\infty} S_d$, we obtain $I(S')[d] = I(W)[d]$ for every $d \geqslant 1$. Thus $I(S') = I(W)$ and S' is dense in W. □

Proof of Proposition 5.1. We will assume that $X \subset \mathbb{P}^n$ is a locally closed subvariety of $\mathbb{P}^n$ for some $n \geq 1$. Let $V \subset X \times \mathbb{P}^n$ be the closure of the union of the images of

$$s_g : X \dashrightarrow X \times \mathbb{P}^n ,$$

over all $g \in G$, as in (3.1). Applying Lemma 5.2 to the generic points of the closures of the images of s_g, as $g \in G$, we see that there exists a countable collection of elements $\{h_i \mid i \geqslant 1\}$ such that the images of s_{h_i} are dense in V. Let H be the countable subsemigroup of G generated by these h_i. Denote the flat locus of the projection $\pi: V \to X$ by $X_0 \subset X$ and the morphism associated to π, viewed as a family of subschemes over X_0, by $\phi: X_0 \to \mathrm{Hilb}(n)$, as before. Arguing as in Section 3, we see that there exists a countable collection $\{W_i \mid i \geqslant 1\}$ of H-invariant closed subvarieties of X_0 such that for any field extension K/k and any point $x \in X_0(K)$ away from $\cup_{i\geqslant 1} W_i$ and from the indeterminacy locus of every $h \in H$, $V_x = \overline{H \cdot x}$ in X; cf. (3.3). Now $\phi(x) = \phi(y)$ if and only if $V_x = V_y$ if and only if $\overline{H \cdot x} = \overline{H \cdot y}$. This proves part (b).

(c) Since $H \cdot x \subset G \cdot x \subset V_x$, we have $\overline{H \cdot x} = \overline{G \cdot x} = V_x$.

(a) By our construction, $\phi^* k(Z) \subset k(X)^G \subset k(X)^H$, and by Theorem 1.2(c), $k(X)^H$ is purely inseparable over $\phi^* k(Z)$. Thus $k(X)^H$ is purely inseparable over $k(X)^G$.

It remains to show that, in fact, $k(X)^H = k(X)^G$. Choose $a \in k(X)^H$. Then a satisfies some purely inseparable polynomial $p(t) \in k(X)^G[t]$. For any $g \in G$, $g(a)$ also satisfies $p(t)$. Since a is the only root of $p(t)$ in $k(X)$, we conclude that $g(a) = a$. In other words, $a \in k(X)^G$, *i.e.*, $k(X)^H = k(X)^G$, as desired. □

6. A variant of the Medvedev-Scanlon conjecture and examples

The following corollary of Proposition 5.1 is a generalization of the Medvedev-Scanlon conjecture [12, Conjecture 7.14] in the case where the base field k is uncountable; cf. Remark (3) in the Introduction.

Corollary 6.1. *Let k be an uncountable algebraically closed field, X be an irreducible k-variety, and G be a semigroup of dominant rational maps $X \dashrightarrow X$. If $k(X)^G = k$ then $G \cdot x$ is dense in X for some $x \in X(k)$.*

Here $G \cdot x := \{g(x) \mid g \in G$ is defined at $x\}$, as in (1.1).

Proof. After replacing X by a dense open subvariety, we may assume that X is quasi-projective. Let $X_0 \subset X$ be a dense open subvariety, and $\phi: X_0 \to Z$ be a morphism such that $\phi^* k(Z) = k(X)^G$, and H be countable subsemigroup of G such that $k(X)^H = k(X)^G$, as in Proposition 5.1.

Since $k(X)^H = k(X)^G = k$, the variety Z is a single point. Let $S \subset X_0(k)$ be the set of k-points of X_0 away from the exceptional subvarieties W_i for every $i \geqslant 1$, and away from the indeterminacy locus of every $h \in H$. Since k is algebraically closed and uncountable, S is Zariski dense in X_0. By Proposition 5.1(b),

$$Y = \overline{H \cdot x} \text{ is independent of the choice of } x \in S. \tag{6.1}$$

Since $h: X \dashrightarrow X$ is dominant for every $h \in H$, we see that the union $\bigcup_{x\in S} \overline{H \cdot x}$ is dense in X. On the other hand, by (6.1), this union equals Y, which is closed in X.

We conclude that $Y = X$, *i.e.*, $\overline{H \cdot x} = X$ and consequently, $\overline{G \cdot x} = X$ for every $x \in S$. □

We note that in [7], it is proved that if X is an Abelian variety, and G is a finitely generated, commutative semigroup of dominant rational maps $X \dashrightarrow X$, then Corollary 6.1 holds for any algebraically closed field k of characteristic 0 (regardless of the cardinality of k).

The following example shows that if we don't assume that k is uncountable, then the Medvedev-Scanlon's conjecture [12, Conjecture 7.14] can fail in prime characteristic.

Example 6.2. Let p be a prime number, and $X = \mathbb{A}^m$ be an m-dimensional affine space defined over an algebraically closed field k of characteristic p, and G be the semigroup generated by the Frobenius endomorphism $\sigma\colon X \to X$. Then $k(X)^{<\sigma>} = k$. On the other hand, there exists a point $x \in \mathbb{A}^m(k)$ with a Zariski dense orbit under σ if and only if $\mathrm{trdeg}_{\mathbb{F}_p}(k) \geq m$.

The following example shows that Proposition 5.1 will fail if we require H to be finitely generated, rather than just countable.

Example 6.3. Let X be a complex Abelian variety of dimension $\geqslant 1$, and G be the group of translations on X by torsion points of $X(\mathbb{C})$. Then G is countable and $\mathbb{C}(X)^G = \mathbb{C}$. On the other hand, any finitely generated subgroup H of G is finite, and $[\mathbb{C}(X) : \mathbb{C}(X)^H] < \infty$. Hence, $\mathbb{C}(C)^G \subsetneq \mathbb{C}(X)^H$ for any finitely generated H. □

The following examples show that the countable collection $\{Y_i,\ i \geq 1\}$ of proper "exceptional" subvarieties of X cannot be replaced by a finite collection in the statement of Theorem 1.2.

Example 6.4. Let E be an elliptic curve over k, let $X = E \times E$, let σ be an automorphism of X given by $\sigma(x, y) := (x, x + y)$, and let $G \xrightarrow{\sim} \mathbb{N}$ (or $G \xrightarrow{\sim} \mathbb{Z}$) be generated by σ as a semigroup (or as a group). In this case Z and ϕ are unique (up to birational isomorphism), and it is easy to see that ϕ is just projection to the first factor, $\phi\colon X \to Z := E$. The fiber $X_z = \{z\} \times E$ is the closure of a single orbit if and only if z is of infinite order in E. Thus the "exceptional set" $Y_1 \cup Y_2 \cup \ldots$ has to contain countably many "vertical" curves $\{z\} \times E$, as z ranges over the torsion points of E.

Example 6.5. Let $X = \mathbb{P}^n$ and $G \xrightarrow{\sim} \mathbb{N}$ be the cyclic semigroup generated by a single dominant morphism $\sigma : X \to X$ of degree ≥ 2. Assume that k is algebraically closed. Then the exceptional collection $Y_1, Y_2, \ldots$ is dense in X; in particular, it cannot be finite.

Proof. We claim that $\mathrm{trdeg}_k\, k(X)^G < n$. Indeed, assume the contrary. Then the field extension $k(X)/k(X)^G$ is algebraic and finitely generated; hence, it is finite. Now we can view $\sigma^*\colon k(X) \to k(X)$ as a $k(X)^G$-linear transformation of a finite-dimensional $k(X)^G$-vector space $k(X)$. Since σ^* is injective, we conclude that it is also surjective, and thus $\sigma : X \to X$ has degree 1, contradicting our choice of σ.

The claim tells us that the general fiber (and hence, every non-empty fiber) of the map $\sigma\colon X \to Z$ has dimension $\geqslant 1$. Suppose $x \in X(k)$ is a periodic point of X, *i.e.*, $\sigma^n(x) = x$ for some $n \geqslant 1$. Then $G \cdot x$ is finite, and thus $\overline{G \cdot x}$ is 0-dimensional. Consequently, $G \cdot x$ cannot be dense in the fiber $\phi^{-1}(\phi(x))$, so x has to lie in the exceptional locus $Y_1 \cup Y_2 \cup \dots$. On the other hand, by a result of N. Fakhruddin [6, Corollary 5.3] periodic k-points for σ are dense in $X = \mathbb{P}^n$. □

7. Rosenlicht's theorem revisited

In this section we will give a proof of Theorem 1.1(a) under the assumption that X is quasi-projective, in the spirit of the arguments in this paper. The idea is to modify Rosenlicht's original proof by replacing the Chow variety with the Hilbert scheme. Note that part (b) of Theorem 1.1 follows from part (a) by the same argument we used in Section 4.

Denote our G-action on X by $\alpha\colon G \times X \to X$, and consider the graph of this action, *i.e.*, the k-morphism

$$\psi = (\mathrm{pr}_2, \alpha)\colon G \times X \to X \times X \subset X \times \mathbb{P}^n,$$

where $\mathrm{pr}_2\colon G \times X \to X$ is projection to the second factor.

Lemma 7.1. *There exist integers $d \geqslant 0$ and $e \geqslant 1$, and a G-invariant dense open subvariety $U \subset X$ defined over k, such that for every field extension K/k and every $x \in U(K)$, the orbit $G \cdot x$ has exactly e geometrically irreducible components, each of dimension d.*

Proof. The orbit $G \cdot x \simeq G/G_x$ has $\dfrac{\#\, G}{[G_x : (G^0)_x]} = \dfrac{\#\, G \cdot \#\,(G^0)_x}{\#\, G_x}$ geometrically irreducible components, each of dimension $\dim(G) - \dim(G_x)$. Here G^0 denotes the connected component of G, G_x the stabilizer of x in G, $(G^0)_x$ the stabilizer of x in G^0, and $\#\, H := [H : H^0]$ the number of geometrically irreducible components of an algebraic group H. Thus it suffices to choose U so that $\dim(G_x)$, $\#\,(G^0)_x$ and $\#\, G_x$ are independent of the choice of $x \in U$.

Let S and S_0 be the preimages of the diagonal $\Delta \in X \times X$ under ψ in $G \times X$ and $G^0 \times X$, respectively. The stabilizers G_x and $(G^0)_x$ are the fibers of $(x, x) \in \Delta$ under $\psi_{|S}$ and $\psi_{|S_0}$. Let η be the generic point of $\Delta \simeq X$. The points $x \in X$ such that $\psi_{|S}^{-1}(x)$ has the same dimension (respectively, the same number of geometrically irreducible components) as $\psi_{|S}^{-1}(\eta)$ lie in a dense open subvariety U_1 (respectively, U_2) of X; see [9, Théorème 13.1.3] (respectively, [9, Proposition 9.7.8]). Similarly, the points $x \in X$ such that $\psi_{|S_0}^{-1}(x)$ has the same number of geometrically irreducible components as $\psi_{|S_0}^{-1}(\eta)$ lie in a dense open subvariety $U_3 \subset X$. Taking $U = U_1 \cap U_2 \cap U_3$, we see that $\dim(G_x)$, $\#\,(G^0)_x$ and $\#\, G_x$ are independent of the choice of $x \in U$.

It remains to check that U_1, U_2 and U_3 are G-invariant (and hence, so is U). This follows from the fact that $\psi_{|S}\colon S \to \Delta \simeq X$ and $\psi_{|S_0}\colon S_0 \to \Delta$ are both G-equivariant. Here G acts on $G \times X$ via $c \times \alpha$, where c denotes the conjugation action of G on itself. The morphism $\psi\colon G \times X \to X \times X$ is equivariant under this action and hence, so are $\psi_{|S}$ and $\psi_{|S_0}$. □

From now on, we will replace X by the dense open subvariety U, as in Lemma 7.1, and thus assume that $G \cdot x$ has e geometrically irreducible components, each of dimension d, for every $x \in X$.

Denote the closure of the image of ψ in $X \times \mathbb{P}^n$ by V. Set $\pi : V \to X$ to be the restriction of the natural projection $X \times \mathbb{P}^n \to X$ to V, and $V_x := \pi^{-1}(x)$ to be the fiber of π over a point $x \in X$, as in (3.1).

Lemma 7.2. *Let K/k be a field and $x\colon$ Spec$(K) \to X$ be a dominant point of X. Then:*

(a) *$G \cdot x$ is dense in V_x;*
(b) *V_x is geometrically reduced and has e geometrically irreducible components, each of dimension d.*

Proof. (a) To check that $G \cdot x$ is dense in V_x, we may pass to the algebraic closure $\overline{K}$. For each $\overline{K}$-point $g \in G$, consider the section $s_g\colon X \to V$ taking x to $(x, g \cdot x)$, where $g \cdot x := \alpha(g, x)$. Since G is smooth, $\overline{K}$-points are dense in G, and thus the union of the images of the sections s_g, over all $g \in G(\overline{K})$, is dense in V. Since the point x is dominant, Proposition 2.1 now tells us that the points $g \cdot x$ are dense in V_x, as g ranges over $G(\overline{K})$. Thus $G \cdot x$ is dense in V_x.

(b) Since G is smooth, $G \cdot x$ is geometrically reduced. Moreover, by our assumption $G \cdot x$ has e geometrically irreducible components, each of dimension d. (b) now follows from (a). □

Lemma 7.3. *There is a dense open G-invariant subvariety $X_0 \subset X$ such that:*

(a) *π is flat over X_0;*
(b) *the G-action on X_0 lifts to a G-action on V_{X_0};*
(c) *the fiber V_x is geometrically reduced and has e geometrically irreducible components, each of dimension d for any $x \in X_0$;*
(d) *the orbit $G \cdot x$ contains a dense open subvariety of V_x, for any $x \in X_0$.*

In the case where X is closed in $\mathbb{P}^n$, parts (a) and (b) are obvious, because $V \subset X \times X$ is invariant under the diagonal G-action on $X \times X$, and consequently, the flat locus of the projection $\pi : V \to X$ is G-invariant. Since we do not assume that X is closed in $\mathbb{P}^n$, we need to work a bit harder to lift the G-action to V, over a dense open subvariety of X.

Proof of Lemma 7.3. By generic flatness, π is flat over some dense open (but not necessarily G-invariant) subvariety $U \subset X$. The flat family $\pi : V_U \to U$ induces a morphism $\phi\colon U \to$ Hilb(n). We claim that ϕ (viewed as a rational map $X \dashrightarrow$

Hilb(n)) is G-equivariant, where G acts trivially on Hilb(n). To prove this, let $K = k(X)$, $\eta \in X(K)$ be the generic point of X. We want to show that $\phi(g \cdot \eta) = \phi(\eta)$ for every $g \in G(\overline{K})$ or equivalently, $V_\eta = V_{g\cdot\eta}$. By Lemma 7.2(a), $V_{g\eta} = \overline{G \cdot (g\eta)} = \overline{G \cdot \eta} = V_\eta$, and the claim follows.

The domain $X_0 \subset X$ of the rational map $\phi\colon X \dashrightarrow \mathrm{Hilb}(n)$ is thus a G-invariant dense open subvariety X_0 of X containing U. The morphism $\phi\colon X_0 \to \mathrm{Hilb}(n)$ corresponds to a family $V' \subset X_0 \times \mathbb{P}^n$, which is flat over X_0. Over U, V and V' coincide. To prove part (a), it suffices to show that they coincide over X_0 (and thus V is flat over X_0). Since $G \times U$ is dense in $G \times X$, $\psi(G \times U)$ is dense in V. Hence, V_{X_0} is the Zariski closure of V_U in $X_0 \times \mathbb{P}^n$. On the other hand, by [8, Théorème 2.3.10], V' is the Zariski closure of V'_U. Since $V_U = V'_U$, we conclude that $V_{X_0} = V'$. This completes the proof of part (a).

For notational convenience, we will replace X by X_0 and thus assume that π is flat over X.

(b) Since ϕ is G-equivariant, $\phi \circ \alpha = \phi \circ \mathrm{pr}_2\colon G \times X \to X \to \mathrm{Hilb}(n)$, and thus the flat family of subvarieties of $\mathbb{P}^n$ over $G \times X$ corresponding to $\phi \circ \alpha$ is $G \times V$. By the universal property of Hilb(n), α lifts to β such that the following diagram commutes:

$$\begin{array}{ccccc} G \times V & \xrightarrow{\ \beta\ } & V & & \\ {\scriptstyle \mathrm{id}\times\pi}\downarrow & & \downarrow{\scriptstyle \pi} & & \\ G \times X & \xrightarrow[\ \alpha\ ]{} & X & \xrightarrow{\ \phi\ } & \mathrm{Hilb}(n). \end{array}$$

Lifting the morphisms $G \times (G \times X) \to X$ given by $(g, h, x) \to \alpha(g, \alpha(h, x))$ and $(g, h, x) \to \alpha(gh, x)$ to $G \times G \times V \to V$, and using the universal property of Hilb(n), we see that β is a G-action on V.

(c) Let η be the generic point of X. By Lemma 7.2(b), V_η is geometrically reduced and has e geometrically irreducible components, each of dimension d. By [8, Corollaire 2.3.5], V has e geometrically irreducible components, each of dimension d. In particular, for any $x \in X$, each geometrically irreducible component of the fiber V_x is of dimension $\geqslant d$.

By [9, Théorème 12.2.4], the points $x \in X$, such that V_x is geometrically reduced and has exactly e geometrically irreducible components lie in a dense open subvariety $X_0 \subset X$. By part (b), π is G-equivariant morphism $V \to X$; hence, X_0 is G-invariant.

Since $\pi\colon V \to X$ is flat, $\dim(V_x)$ is constant, as x ranges over X. Since $\dim(V_\eta) = d$, we conclude that $\dim(V_x) = d$ for every $x \in X$. As we saw above, each geometrically irreducible component of V_x has dimension $\geqslant d$. Thus each geometrically irreducible component of V_x has dimension d, and part (c) follows.

(d) By our assumption, for any $x \in X_0$, the orbit $G \cdot x$ has exactly e geometrically irreducible components, each of dimension d. By part (c), the same is true of V_x, and V_x is reduced. Thus $G \cdot x$ is Zariski dense in V_x for every $x \in X_0$. □

We are now ready to complete the proof of Theorem 1.1(a). Let $\phi\colon X_0 \to \mathrm{Hilb}(n)$ be the morphism induced by the flat family $\pi : V_{X_0} \to X_0$, as in Lemma 7.3. Set Z to be the image of ϕ. It remains to show that $\phi(x) = \phi(y)$ if and only if $G \cdot x = G \cdot y$ for any $x, y \in X_0$. By our construction, $\phi(x) = \phi(y)$ if and only $V_x = V_y$. We thus need to check that $V_x = V_y$ if and only if $G \cdot x = G \cdot y$.

Indeed, suppose $V_x = V_y$. By Lemma 7.3(d), each orbit $G \cdot x$, $G \cdot y$ contains a dense open subvariety of $V_x = V_y$. Thus $G \cdot x \cap G \cdot y \neq \emptyset$ and consequently, $G \cdot x = G \cdot y$. Conversely, suppose $G \cdot x = G \cdot y$. Using Lemma 7.3(d) once again, we see that $V_x = \overline{G \cdot x} = \overline{G \cdot y} = V_y$, as desired. □

References

[1] E. AMERIK and F. CAMPANA, *Fibrations méromorphes sur certain variétés à fibré canonique trivial*, Pure Appl. Math. Q. **4** (2008), 1–37.

[2] E. AMERIK, *Some applications of p-adic uniformization to algebraic dynamics*, arXiv: 1407.1558, 2014, 21 pages.

[3] J. P. BELL, D. GHIOCA and T. J. TUCKER, *Applications of p-adic analysis for bounding periods for subvarieties under étale maps*, Int. Math. Res. Not. IMRN **2015**, 3576–3597.

[4] J. P. BELL, D. ROGALSKI and S. J. SIERRA, *The Dixmier-Moeglin equivalence for twisted homogeneous coordinate rings*, Israel J. Math. **180** (2010), 461–507.

[5] D. EISENBUD, "Commutative Algebra", Graduate Texts in Mathematics, Vol. 150, Springer-Verlag, New York, 1995.

[6] N. FAKHRUDDIN, *Questions on self maps of algebraic varieties*, J. Ramanujan Math. Soc. **18** (2003), 109–122.

[7] D. GHIOCA and T. SCANLON, *Density of orbits of endomorphisms of Abelian varieties*, Trans. Amer. Math. Soc. **369** (2017), 447-466.

[8] A. GROTHENDIECK, *Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. II*, Publ. Math. Inst. Hautes Études Sci. (1965), 231.

[9] A. GROTHENDIECK, *Éléments de géométrie algébrique. IV. Étude locale des schémas et des morphismes de schémas. III*, Publ. Math. Inst. Hautes Études Sci. (1966), 255.

[10] A. GROTHENDIECK, *Techniques de construction et théorèmes d'existence en géométrie algébrique. IV. Les schémas de Hilbert*, In: "Séminaire Bourbaki", Vol. 6, Soc. Math. France, Paris, 1995, pp. Exp. No. 221, 249–276.

[11] G. KEMPER, *The computation of invariant fields and a constructive version of a theorem by Rosenlicht*, Transform. Groups **12** (2007), 657–670.

[12] A. MEDVEDEV and T. SCANLON, *Invariant varieties for polynomial dynamical systems*, Ann. of Math. (2) **179** (2014), 81–177.

[13] V. L. POPOV, *On infinite dimensional algebraic transformation groups*, Transform. Groups **19** (2014), 549–568.

[14] V. L. POPOV and E. B. VINBERG, "Invariant Theory, Algebraic Geometry, IV", Encyclopaedia of Mathematical Sciences, Vol. 55, Springer-Verlag, Berlin, 1994, pp. 284.

[15] M. ROSENLICHT, *Some basic theorems on algebraic groups*, Amer. J. Math. **78** (1956), 401–443.

[16] M. ROSENLICHT, *A remark on quotient spaces*, An. Acad. Brasil. Ciênc. **35** (1963), 487–489.

[17] A. SEIDENBERG, *On the functions invariant under the action of an algebraic group on an algebraic variety*, Rend. Semin. Mat. Fis. Milano **49** (1979), 69–77 (1981).

[18] S.-W. ZHANG, *Distributions in algebraic dynamics*, In: "Surveys in Differential Geometry", Vol. X, Int. Press, Somerville, MA, 2006, 381–430.

Department of Pure Mathematics
University of Waterloo
Waterloo, ON N2L 3G1, Canada
jpbell@uwaterloo.ca

Department of Mathematics
University of British Columbia
Vancouver, BC V6T 1Z2, Canada
dghioca@math.ubc.ca
reichst@math.ubc.ca