

Minimal sets in ACFA

by

Alice Medvedev

B.A. (California Institute of Technology) 2001

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:

Professor Thomas W. Scanlon, Chair

Professor Bjorn Poonen

Professor John MacFarlane

Spring 2007

The dissertation of Alice Medvedev is approved:

Chair

Date

Date

Date

University of California, Berkeley

Spring 2007

Minimal sets in ACFA

Copyright 2007

by

Alice Medvedev

Abstract

Minimal sets in ACFA

by

Alice Medvedev

Doctor of Philosophy in Mathematics

University of California, Berkeley

Professor Thomas W. Scanlon, Chair

A *difference field* is a field with a distinguished automorphism. ACFA is the model-companion of the theory of difference fields, in the natural language $L = \{+, \cdot, -, /, 0, 1, \sigma\}$. ACFA is a supersimple theory whose minimal types satisfy a version of the Zilber trichotomy: each is exactly one of field-like (non-orthogonal to a fixed field of a definable field automorphism), group-like (non-orthogonal to a minimal modular definable group), or trivial.

The only definable automorphisms are powers of σ , powers of the Frobenius automorphism, and compositions of these. Definable minimal modular groups are characterized in [CH99] as certain subgroups of algebraic groups. Very little is known at present about the trivial case.

We characterize certain group-like σ -degree 1 minimal sets, thereby providing many new explicit examples of trivial minimal sets. We show that a minimal set defined by $\sigma(x) = f(x)$ for a rational function f in one variable is group-like if and only if f is not purely inseparable (otherwise $(\mathbb{P}^1, f)^\sharp$ is field-like), and there exist an algebraic group G , a morphism of algebraic groups ϕ , and a morphism of curves π with $\deg(\pi) \leq \max(24, \deg(f))$ making the following diagram commute:

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \pi & & \downarrow \pi^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

The bound on the degree of π comes from a surprising observation that π must be the

quotient by a group of automorphisms of the curve G . This bound shows that each case of the trichotomy is definable for minimal sets of this form. Such definability results may lead to new model-complete theories of difference fields.

The model-theoretic notion “ $(\mathbb{P}^1, f)^\sharp$ is group-like” translated into algebraic geometry becomes a commutative diagram of algebraic curves containing f , which can be simplified using the theory of algebraic curves. We analyze jet spaces in that diagram to find the group of automorphisms of G that acts transitively on generic fibers of π . Jet spaces are one of two different higher-dimensional generalizations of tangent spaces; their construction is not found in the literature in sufficient detail, so we recount it in chapter 3.

Professor Thomas W. Scanlon
Dissertation Committee Chair

To friends and colleagues
who have the patience to deal with me.

Contents

1	Introduction	1
1.1	Background	1
1.1.1	Model-theoretic conventions	1
1.1.2	Difference algebra: the quantifier-free theory	2
1.1.3	Ranks in ACFA	3
1.1.4	Genericity	6
1.1.5	Non-orthogonality	7
1.2	Previous foundational work	8
1.3	Questions and answers	10
1.3.1	Motivation	12
1.4	Algebraic geometry notation, conventions and lemmas	13
1.4.1	Algebraic geometry: definitions	13
1.4.2	Definability in algebraic geometry	15
1.4.3	Funny notions	18
2	Encoding	22
2.1	A collection of well-known facts about sharps.	23
2.1.1	Sharps	23
2.1.2	Definability discussion	27
2.1.3	Finite rank	28
2.2	Group-like minimal types	32
2.3	The diagram appears	34
2.4	Stronger conclusions under extra assumptions	35
2.4.1	A small reduction	36
3	Construction of jet spaces	38
3.1	k -algebra functors	39
3.1.1	One ring	40
3.1.2	Commutative diagram for jets	41
3.1.3	Two rings	42
3.1.4	Three rings	43
3.2	Generators	43
3.3	Functors and (pre)sheaves	45

3.3.1	Localizations and quasi-coherence	47
3.3.2	Conclusions	48
3.4	Jet Spaces	49
3.4.1	<i>Spec</i> Sym	49
3.4.2	Definability	49
3.4.3	Jet spaces: fibers of jet bundles	50
3.4.4	Linear maps induced on jet spaces	53
3.5	Information encoded in jet spaces	54
3.5.1	Constant morphisms crush jet spaces	54
3.5.2	Relations between jet bundles	55
3.5.3	A morphism is determined by its jets at a point	56
3.5.4	An étale morphism induces an isomorphism on jet spaces.	57
3.6	Commutative algebra lemmata	58
4	Main story	60
4.1	Part 1, in which the group correspondence and the intermediate correspondence are replaced by morphisms.	61
4.1.1	Summary	61
4.1.2	The first diagram chase, to get (\star)	62
4.1.3	The purpose of (\star)	69
4.1.4	The second diagram chase	70
4.1.5	After the chases	72
4.2	Part 2, in which the intermediate minimal set turns out to be a group.	73
4.2.1	Setup	73
4.2.2	Elliptic curves	75
4.2.3	\mathbb{G}_a	75
4.2.4	\mathbb{G}_m	78
4.2.5	Conclusion and theorem	79
4.3	Part 3, in which jet spaces are analyzed and the group of automorphisms of G is found.	80
4.3.1	Outline	80
4.3.2	Jets	81
4.3.3	ν_{ab}^m	82
4.3.4	N_a^m	84
4.3.5	Another reduction, towards uniqueness	86
4.3.6	Automorphisms	89
4.4	Part 4, in which the degree of the cover is bounded.	91
4.4.1	Elliptic curves	91
4.4.2	\mathbb{G}_m	92
4.4.3	\mathbb{G}_a	93
4.4.4	The End	93
	Bibliography	95

Acknowledgments

I would like to thank many people for their comments, questions, and answers about this work. I would like to particularly thank Zoé Chatzidakis for explaining her published and unpublished results on the model theory of difference fields.

I would like to thank Bjorn Poonen for his meticulous reading of this thesis and his numerous suggestions simplifying the proofs and clarifying the exposition.

I would like to thank all my friends for their support and encouragement.

And above all I would like to thank my advisor, Thomas Scanlon, for all of the above.

Chapter 1

Introduction

1.1 Background

A *difference field* is a field with a distinguished automorphism. We work in $L = \{+, \cdot, -, /, 0, 1, \sigma\}$, the first-order language of fields with an additional unary function symbol σ . Then a difference field is an L -structure on which the language of rings gives a field structure, and σ acts as a field automorphism. ACFA (an acronym for Algebraically Closed Field with an Automorphism) is the model-companion of the theory of difference fields. The name is poorly chosen, because not every algebraically closed field with an automorphism is a model of ACFA: the automorphism should be generic in a certain sense, so that as many difference equations as possible have solutions. Our precise meaning of “generic” boils down to the definition of model-companion.

The model-companion of a theory is the first-order axiomatization of the class of existentially closed models of the theory, if this class is axiomatizable. ACFA becomes a complete theory once the characteristic of the field and the action of σ on the algebraic closure of the prime field are specified.

1.1.1 Model-theoretic conventions

We work in a monster model K of ACFA, over a model k of ACFA: unless otherwise specified “definable” means “definable with parameters from k ”. The model k should be small relative to the monster model, but nevertheless somewhat saturated.

Our *types* are always complete, unless explicitly stated otherwise. A set is ∞ -

definable is it can be defined by a partial type. [Hod97] is a good introduction to model theory.

We constantly abuse notation by identifying a formula or a type with the set of its realizations. We consistently use “subset” semantically, to refer to sets of realizations. For example, if p and q are partial types, “ p is a subset of q ” means that p implies q .

We use many notions from stability and simplicity theory without defining them properly. We only include the definitions that are necessary for our purposes, and we often replace the common, general definition with one that is equivalent in our setting and convenient for our purposes. A thorough explanation of forking independence, ranks, orthogonality, and the Zilber trichotomy can be found in [Pil96] for stable theories and in [Wag00] for simple theories.

We say that two sets are *definably isomorphic* if there is a definable bijection between them. We say that two sets A and B are *definably interalgebraic* if there is a definable, finite-to-finite correspondence between them, that is a definable subset of $A \times B$ that projects onto both A and B , with both projections having finite fibers. Note that the sets A and B need not be definable. On the other hand, any set that is definably interalgebraic with a definable set is itself definable.

This thesis is concerned with two first-order theories, in two different languages. One is ACFA described above; the other is the theory of algebraically closed fields, in the language of fields. The second is much easier to work with, and we often try to reduce questions about the first to questions about the second.

1.1.2 Difference algebra: the quantifier-free theory

This purely algebraic treatment of difference fields can be found in Cohn’s “Difference Algebra” [Coh65].

The quantifier-free type of a tuple a over a set of parameters B is completely determined by the difference ideal of difference polynomials with coefficients in B satisfied by a , and these ideals are well-behaved. Such an ideal is always *perfect*. Perfect difference ideals are the appropriate generalization of radical ideals. Every perfect ideal in the ring of difference polynomials over a difference field is finitely generated as perfect ideal. Perfect difference ideals satisfy the ascending chain condition and factor into prime ideals, corresponding to *difference-irreducible* components of their solution sets. These properties give

rise to a *difference-Zariski topology*, similar to the Zariski topology in algebraic geometry.

Definition 1. *The set D of solutions of a prime difference ideal will be called difference-irreducible. More generally, a quantifier-free definable set S will be called difference-irreducible if the smallest difference variety (i.e. solution set of a system of polynomial difference equations) containing S is difference-irreducible. (Note that irreducibility is a well-defined notion, unaffected by adding further parameters, since we are already working over a submodel.)*

One construction from difference algebra is often useful to us, so we define it here:

Definition 2. *The N th prolongation of a set S is*

$$S(N) := \{(s, \sigma(s), \sigma^2(s), \dots, \sigma^N(s)) \mid s \in S\}$$

It follows from the above description of quantifier-free types that the quantifier-free part of ACFA is ω -stable, of rank ω .

Unlike the theory of algebraically closed fields, ACFA does not eliminate quantifiers: in contrast to usual algebraic geometry, not all constructible sets are (quantifier-free) definable. For example, let $F^2 := \{x : \exists y y^2 = x \wedge \sigma(y) = y\}$; the smallest quantifier-free definable set containing it is $F := \{x : \sigma(x) = x\}$, but the two are not equal (except in characteristic 2), since the fixed field of σ is pseudo-finite rather than algebraically closed.

1.1.3 Ranks in ACFA

There are several notions of dimension for definable sets and types in ACFA. Each comes from a kind of niceness of the theory. The values of these ranks are ordinals or the symbol ∞ ; ∞ is defined to be greater than all ordinals.

σ -degree

The σ -degree is a purely algebraic notion, naturally defined for quantifier-free types.

Definition 3. *The σ -degree of a tuple a over a set of parameters B is the transcendence degree of the difference field generated by $B \cup \{a, \sigma(a), \sigma^2(a), \dots\}$ over the difference field generated by B , if this is finite. Otherwise, we say that the σ -degree of a over B is ∞ .*

For example, if any of the a_i are *transformally transcendental over B* , meaning that $\{a, \sigma(a), \sigma^2(a), \dots\}$ satisfy no non-trivial polynomial equations with coefficients in B , then the σ -degree of a over B is ∞ .

The notion of σ -degree can be extended to types, though it will only depend on the quantifier-free part. It can be extended to definable (or ∞ -definable) sets in the usual manner:

Definition 4. *the σ -degree of a definable (or ∞ -definable) set S is defined to be the maximum of the σ -degrees of tuples in S .*

This notion is well-behaved because it is closely tied to difference algebra and to well-behaved difference ideals. This notion is problematic because there are unhappy definable sets of σ -degree 2 without definable subsets of σ -degree 1. Nevertheless, it is a useful notion because of several properties, proved in [CH99]:

- Fact 1.1.**
1. *The σ -degree of types is upper-semi-continuous: if the σ -degree of p is n , then p contains a formula of σ -degree n , i.e. one that forces p to have σ -degree n .*
 2. *The σ -degree of definable sets is definable: for each formula $\phi(x, y)$ and for each $n \in \mathbb{N}$ the set $\{b \mid \deg_\sigma(\phi(x, b)) \geq n\}$ is definable.*

Stability ranks

Since the quantifier-free part of ACFA is ω -stable, all ordinal-valued ranks from stability theory, such as Morley rank, are applicable to quantifier-free formulas and types in ACFA. This thesis does not use these, so we only mention them briefly.

Lascar Rank

ACFA does not eliminate quantifiers, and the full theory is not stable. For example, the partial type $p(x) := x \notin \text{acl}(B) \wedge \sigma(x) = x$ extends to a unique complete quantifier-free type over B , but does not decide whether σ fixes or swaps the square roots of x ; in fact, the formula $\sigma(\sqrt{x+y}) = \sqrt{x+y}$ has Shelah's independence property on the set of realizations of p , witnessing instability. However, ACFA is supersimple ([CH99], Remark 3 on p.3011), so the ranks from simplicity theory work well in ACFA. In particular, Lascar's U-rank, the foundation rank of non-forking independence, makes sense for complete types and is always ordinal-valued (as opposed to $= \infty$).

ACFA satisfies a strong form of the Stable Forking Conjecture: forking in ACFA is always witnessed by a quantifier-free formula. Since the quantifier-free part of ACFA is stable, a type over an algebraically closed set of parameters has a unique quantifier-free non-forking extension. This non-forking extension can be characterized as follows in terms of difference ideals: the type of a tuple a over a bigger field L does not fork over a smaller, algebraically closed field k if and only if the ideal of difference polynomials over L satisfied by a is the tensor product (over k) of L with the ideal of difference polynomials over k satisfied by a . This relation between forking and ideals gives rise to a relation between the Lascar rank and the σ -degree: Lascar rank goes up when a “new” difference polynomial is added to the ideal, reducing the σ -degree.

Here are several useful properties of ranks in ACFA:

- Fact 1.2.**
1. *Both the Lascar rank and the σ -degree witness forking: if q is an extension of p , then $R(q) \leq R(p) < \infty$, with equality iff the extension is non-forking.*
 2. *The Lascar rank of a type p is bounded by the σ -degree of p . This is only exciting when both are finite.*
 3. *$U(p) = 0$ iff $\deg_\sigma(p) = 0$ iff p is an algebraic type.*
 4. *In particular, any type of σ -degree 1 must also have U -rank 1. The converse is false.*
 5. *The Lascar rank of a type p is finite iff the σ -degree of p is finite.*

Even when both are finite, it is possible for Lascar rank to be strictly less than the σ -degree. It is possible for U -rank to not be definable, unlike σ -degree. It may happen that all formulas in a type p of finite Lascar rank have strictly higher Lascar-rank than p . One example of such a type can be constructed from a simple abelian variety of dimension 2 that is as a limit of a family of products of elliptic curves, as described in [HS99].

Definition 5. *Types of U -rank 1 are called minimal; a formula ϕ , and the set defined by it, is called minimal if all types containing ϕ are minimal or algebraic.*

Notation abuse alert: what we call “minimal” definable sets are usually called “weakly minimal”, to distinguish from an older notion that does not appear in this thesis.

This thesis is about certain minimal sets of σ -degree 1.

Fact 1.3. *Minimal types of σ -degree 1 are pleasant because they always contain a (quantifier-free) minimal formula. (This formula is the conjunction of the finitely many difference polynomials forcing the σ -degree to be 1.)*

There exist minimal types of higher σ -degree, but they are significantly harder to work with. For example, they might not contain a minimal formula, i.e. they might occur as limits of families of non-minimal sets. Even the minimal formulas of higher σ -degrees are problematic, because the algebraic geometry involved is much more involved.

1.1.4 Genericity

There are several notions of a “big” subset of a definable set. If the set of realizations of a type is a big subset of T , then the type is called *generic* in T . More generally, we say that a statement holds *generically* if it holds on a big subset. The precise meaning of “generic” and “generically” depend on the precise meaning of “big”. We make use of the last two of the following notions of “big”:

Definition 6. • $S \subset T$ is difference-Zariski dense if the closure of S in the difference-Zariski topology contains T . This is the strongest notion. Note that a complete type cannot be difference-Zariski dense in a difference-reducible set T .

- For sets T of a very special form, we will define a notion of very dense in Definition 20 in Chapter 2. This notion is weaker than “difference-Zariski dense” and stronger than model-theoretic genericity in general, but equivalent to model-theoretic genericity in the special case studied by this thesis.
- An ∞ -definable subset S of an ∞ -definable set T is model-theoretically generic in T iff the rank of S is the same as the rank of T . Remember, for sets of finite rank, this definition is the same regardless of which rank you use.

Sometimes, especially in algebraic geometry, especially for subsets of curves in algebraic geometry where the last two notions of genericity coincide with co-finiteness, we say “almost everywhere” instead of “generically”. This thesis is mostly about curves, so the distinction between these two notions is only relevant in Chapter 2 where everything is done in unnecessary generality.

1.1.5 Non-orthogonality

Another model-theoretic concept central to this thesis is *non-orthogonality*. Throughout this thesis, we work over a small but sufficiently saturated model, so we do not need to distinguish between orthogonality and almost-orthogonality; the definitions we use properly belong to almost-orthogonality.

Definition 7. *Complete types p and q are almost-orthogonal if for all realizations $a \models p$ and $b \models q$ the rank of the type of a over M is the same as the rank of the type of a over Mb .*

Fact 1.4. *Complete types p and q over a sufficiently saturated model M are orthogonal if and only if they are almost-orthogonal.*

It is true, though by no means obvious, that this notion is symmetric; and that any of the ranks mentioned above give the same notion of orthogonality, except that σ -degree only works for finite-rank types. (The first assertion follows from the symmetry of forking in simple theories; the second from the fact that each of those ranks witnesses forking.)

This thesis is all about non-orthogonality between minimal types. The interesting potential conclusions of this thesis rely on Hrushovski's analysis of finite-rank types in terms of minimal types, which in turn relies on non-orthogonality between a finite-rank type and a minimal one. Note that

Fact 1.5. *If p and q are complete types over M , and p is minimal and non-orthogonal to q , then for every realization $a \models p$ there is some realization $b \models q$ such that $R(a/Mb) \leq R(a/M) = 1$. I.e. $a \in \text{acl}(Mb)$.*

By compactness, this algebraicity is witnessed by a formula that depends only on the types p and q and not on particular realizations. If both types are minimal, algebraicity becomes *inter-algebraicity*:

Fact 1.6. *If p and q are complete, minimal, non-orthogonal types over M , then there is a formula $\theta(x, y)$, over M , such that*

- *for all $(a, b) \models \theta$, $a \in \text{acl}(Mb)$ and $b \in \text{acl}(Ma)$; and*
- *for all $a \models p$ there exists $b \models q$ such that $\models \theta(a, b)$, and vice versa.*

Fact 1.7. *Non-orthogonality is an equivalence relation on the set of minimal types over a model. In particular, it is transitive.*

Definition 8. *A type p is non-orthogonal to a definable set defined by a formula ψ if some type q containing ψ is non-orthogonal to p .*

One more application of compactness shows

Lemma 1.8. *If a formula θ witnesses non-orthogonality between a minimal type p and a type q , then there are formulas $\phi \in p$ and $\psi \in q$ such that θ defines a many-to-finite multi-valued function from ψ to ϕ .*

If a formula θ witnesses non-orthogonality between two minimal types p and q , then there are formulas $\phi \in p$ and $\psi \in q$ such that θ defines a finite-to-finite correspondence between them.

This motivates the following definition:

Definition 9. *We call two definable sets interalgebraic if there is a definable finite-to-finite correspondence between them; equivalently, if every type containing the formula defining one set is non-orthogonal to some type containing the defining formula of the other set.*

We call a definable set S algebraic over a definable set T if there is a definable, many-to-finite multi-valued function from T to S .

1.2 Previous foundational work

Difference equations first arose as functional equations in analysis. The abstract algebraic setting of difference algebra is described in Cohn's book [Coh65]. The model theory of difference fields in general and ACFA in particular was first studied in [Mac97], where a first-order axiomatization of ACFA is first given. In [CH99] and [CHP02], Chatzidakis, Hrushovski, and Peterzil develop the modern simplicity-theoretic approach. They show that ACFA is supersimple, that its quantifier-free part is ω -stable, and that forking is witnessed by quantifier-free formulas. They show that types of finite U-rank can be analyzed in terms of minimal types, and that the minimal types satisfy a version of Zilber's trichotomy:

Fact 1.9. *Let p be a minimal type over an algebraically-closed set in ACFA. Then exactly one of the following holds:*

- (field-like) p is non-orthogonal to the fixed field of a definable field automorphism.
- (group-like) p is locally-modular and non-orthogonal to a minimal definable group.
- (trivial) p is locally-modular, and the algebraic closure operator on the set of realizations of p is trivial: $\text{acl}(A) = \cup_{a \in A} \text{acl}(a)$.

The first case is separated from local modularity in the theorem in section 6 of [CHP02]. Fixed fields of definable automorphisms are stably-embedded, and all induced definable structure on them is already definable in the language of fields; they are pseudo-finite. Field-like minimal types have been completely characterized in [CH99] (characteristic zero) and in [CHP02] (positive characteristic).

The separation between the last two cases is also proved in that paper, and fits in with general group configuration theorems for stable and simple theories. This is stronger than the usual theorems in that the group is honestly definable rather than type-definable imaginary, or hyper-definable hyper-imaginary.

The Zilber trichotomy sometimes also make sense for minimal definable sets:

Definition 10. *We will call a minimal definable set (or the formula defining it) field-like (respectively, group-like, trivial) if all sufficiently generic types in the definable set (containing the defining formula) are field-like (resp., group-like, trivial). Otherwise we will call the definable set mixed.*

Unfortunately, we must postpone the explanation of what “sufficiently generic” above means until Definition 20.

Definable minimal modular groups are characterized in [CH99] as certain subgroups of algebraic groups:

Fact 1.10. *If G is a minimal group definable over a small model of ACFA, then there exists an algebraic group A and a Zariski-dense, definable, minimal subgroup $H \leq A$, which is interalgebraic with (indeed, finitely covered by) a finite-index subgroup of G .*

In characteristic zero, minimal modular groups are stable and stably embedded, so they can be analyzed with the tools of stable group theory. In positive characteristic, at least the tools of supersimple groups are still available.

Very little is known at present about trivial minimal types and sets in ACFA, except that in characteristic zero they are stable, stably embedded.

This thesis gives a surprisingly nice characterization of certain group-like σ -degree 1 minimal sets, thereby providing many new explicit examples of trivial minimal sets.

1.3 Questions and answers

This thesis concerns the definability of the cases of the Zilber trichotomy for families of minimal sets. Every definable set in ACFA admits a finite cover by a quantifier-free definable set, so we may restrict our attention to quantifier-free definable sets without loss of generality.

On the other hand, we do lose generality by describing only minimal sets of σ -degree 1. We expect to find a similar characterization of group-like minimal sets of higher σ -degree, but the reduction to algebraic geometry is more delicate in that case, and the algebraic geometry itself is more involved.

It is easy to see (Proposition 2.13) that σ -degree is invariant under non-orthogonality, so we only need to consider σ -degree 1 groups and fields in deciding to which case of the trichotomy a given minimal set of σ -degree 1 belongs.

We make one more technical assumption on the minimal sets we work with. In Chapter 2, we will show that any quantifier-free definable, σ -degree 1 minimal set can be broken into pieces each of which is “encoded” by a pair of curves V and W where W is a finite-to-finite correspondence between V and V^σ . In this thesis, we further assume that W is finite-to-*one*: i.e. that W is a graph of a function from V to V^σ . This definitely loses generality; for example, the reasonable analog of Fact 1.11 is false without this assumption. However, this special case is sufficient for many applications of ACFA and sufficiently rich to be interesting in its own right. Unlike the σ -degree, this assumption is not preserved under non-orthogonality, so we cannot assume that the minimal sets witnessing, for example, group-likeness of one of our minimal sets, satisfies this assumption. Once we have made this assumption, we may assume without loss of generality that the curve V is \mathbb{P}^1 : if the genus of V is 1, it is not hard to see that the minimal set must be group-like or field-like, and it is easy to distinguish the two (see Lemma 4.6); if the genus of V is greater than 1, the function from V to V^σ must be purely inseparable, making the minimal set field-like.

Having made these assumptions and reductions, we are ready to state the result:

Definition 11. *Let $P(x; y)$ and $Q(x; y)$ be polynomials over \mathbb{Z} in $n+1$ variables x, y_1, \dots, y_n . Let S_b be the minimal set of σ -degree 1 defined by $\sigma(x) = P(x, b)/Q(x, b)$.*

Theorem 1. S_b is never mixed.

Further, $\{b \mid S_b \text{ is field-like}\}$ is definable, and similarly for the other two cases of Zilber’s trichotomy.

The proof of this theorem (the content of this thesis) relies on a characterization of such minimal sets that are non-orthogonal to a group-like minimal type, because the field-like case is already taken care of in [CH99] and [CHP02]:

Fact 1.11. *The following are equivalent for S_b as above:*

- S_b is non-orthogonal to a field-like type;
- $P(x, b)/Q(x, b)$ is purely inseparable;
- S_b is field-like.

We obtain a similar characterization for the group-like case:

Theorem 2. *For each n and for each pair (P, Q) of polynomials over \mathbb{Z} there is a formula $\delta_{n,P,Q}(y)$ such that the following are equivalent for S_b :*

- S_b is non-orthogonal to a group-like type;
- $\models \delta_{n,P,Q}(b)$;
- S_b is group-like.

We should point out that in both characterizations, the formula should preclude the possibility that $P(x, b)/Q(x, b)$ is constant or undefined, because then S_b is not minimal.

The two above results show that such minimal sets cannot be mixed, and that the field-like and the group-like case are each definable. Then the trivial case is definable as “not field-like and not group-like”.

Sketch of proof of Theorem 2.

In chapter 2, we remind the reader how to translate model-theoretic notions into algebraic geometry. In Theorem 4 we obtain a characterization of group-likeness in terms of algebraic geometry; unfortunately, this characterization involves quantification over curves and morphisms, which is distinctly not first-order.

In chapter 3, we remind the reader about jet spaces, which are needed in chapter 4 to make the characterization from chapter 2 first-order. We describe this construction in the language of schemes, for the sake of simplifying the exposition, but we explain that this construction can be brought back into the classical setting of affinely-embedded varieties and first-order logic.

The new mathematics occurs in chapter 4, where the large commutative diagram obtained in chapter 2 is examined and made simpler. After two sections of ugly technicalities (4.1 and 4.2), we find out that the minimal set defined by $\sigma(x) = f(x)$ is group-like if and only if it is covered by a minimal *group* defined by another equation of the same form. In the next section, we use jet spaces to show that the cover must be a quotient by a finite group of automorphisms. Then it is an easy exercise to show that the degree of the cover is bounded, as we do in section 4.4. After all these metamorphoses, we obtain a characterization of group-like minimal sets in terms of algebraic geometry that only involves quantification over curves of a given genus, and over morphisms of a bounded degree: this is a first-order characterization.

1.3.1 Motivation

This work is inspired by the work of Hrushovski and Itai on differentially closed fields (DCF). In [HI03], they find construct new model-complete theories of differential fields; in contrast to differentially closed fields, in these new theories some systems of differential equations do not have solutions. More precisely, they choose a particular class C of minimal sets in DCF and say that a system of equations has solutions only if in a model of DCF its set of solutions would be orthogonal to every minimal set in C . In order for this new theory to be first-order axiomatizable, being *orthogonal to every minimal set in C* must be definable in families. This thesis is a first step in a similar approach to the search for model-complete theories of difference fields other than ACFA.

So far, we have only considered definability of orthogonality between two minimal sets. To obtain a new model-complete theory, it is also necessary to investigate non-orthogonality of a finite-rank set to a minimal set. This is a rather delicate question, because the Hrushovski analysis of finite-rank types in terms of minimal types occurs on the level of *types*, not on the level of definable sets. This thesis does not completely settle the questions about orthogonality between two minimal sets, either, because we only consider minimal

sets of a special form. Nevertheless, this work suggests where one might hope to find a class C of minimal sets such that orthogonality to it would be definable.

Previous work [CH99] by Chatzidakis and Hrushovski immediately implies that the class of minimal sets non-orthogonal to a definable fixed field is not definable. Although non-orthogonality to *some* minimal group is shown here to be definable (at least for minimal sets of our special form), the results of this thesis suggest that non-orthogonality to a *particular* minimal group will not be definable. Omitting all group-like minimal sets does not appear possible. Therefore, we expect to find our class C among trivial minimal sets. Indeed, it is among trivial minimal sets in DCF that Hrushovski and Itai find their class C .

This thesis gives an easily computable characterization of group-like minimal sets of a certain form. This characterization can be used to easily verify that many minimal sets of that form are trivial. We hope to find our class C among these trivial sets.

1.4 Algebraic geometry notation, conventions and lemmas

1.4.1 Algebraic geometry: definitions

We assume the reader is familiar with algebraic geometry as described in, for example, Hartshorne’s *Algebraic Geometry* [Har77]; chapters 1, 2, and 4 is all we really use. In this section we fix precise meanings for some words that mean different things to different people. Some of these choices are somewhat unorthodox.

The word *morphism* is reserved for geometry; speaking of algebraic objects (groups, rings, fields, modules, etc) we will try to stick to *homomorphism*. Thus we distinguish an abstract group *homomorphism* between algebraic groups from an algebraic group *morphism*: a morphism between the underlying varieties respecting the group law.

For the duration of this section, “definable” means definable in the language of fields. The model-theoretic assumption that we work over a fixed small model k of parameters implies that we are doing algebraic geometry over a fixed algebraically-closed base field k .

We take the most general sheafless meaning of *variety* because we mostly use varieties while describing general and well-known constructions. The actual new work in this thesis, namely chapter 4, is only concerned with varieties of dimension 1, and there we make all possible reductions before beginning the work. Hence we adopt the most restrictive

possible definition of a *curve*.

Varieties

Our varieties are quasi-projective. In particular, varieties do not have to be smooth, irreducible, or closed. In this way, if a subset of a variety is definable in the language of fields, then it is itself a finite union of varieties.

Most of our varieties are affine and affinely embedded: such a variety is a Zariski-open subset of a solution set of a system of polynomial equations in several variables. Again, a definable (in the language of fields) subset of an affine, affinely embedded variety is itself a finite union of affine, affinely embedded varieties.

When R is a k -algebra of the right kind, we abuse notation a little by saying that the scheme $\text{Spec}(R)$ is an affine variety; in that case, choosing generators of R as a k -algebra corresponds to choosing an affine embedding of $\text{Spec}(R)$.

An *algebraic group* is a variety G with a morphism of varieties $G \times G \rightarrow G$ giving the group law.

An *algebraic finite-to-finite correspondence* between varieties A and B is a (pure field language!) definable finite-to-finite correspondence. This correspondence is a (pure field language!) definable subset, hence a finite union of subvarieties, of $A \times B$, which projects dominantly onto both A and B , with both projections having finite fibers.

Notation abuse alert!: morphisms of varieties are *rational*, not necessarily regular.

Curves

In sharp contrast to varieties, our curves are always projective, closed, irreducible, and non-singular. The first three properties can be obtained relatively painlessly for any variety. There is no known canonical way to remove singularities from higher-dimensional varieties.

For dimension 1, there is a functor called “normalization” that takes an irreducible variety and returns a curve birational to it. This functor takes *rational* morphisms between varieties of dimension 1 to *regular* morphisms between curves. The rectangle made of the old morphism between the old varieties, the new morphism between the new varieties, and the two birational isomorphisms commutes almost everywhere.

Definition 12. *A curve is a projective, closed, irreducible, non-singular variety of dimension 1. Morphisms between curves are assumed to be regular.*

Note that “non-constant”, “dominant”, and “surjective” are all the same for morphisms of curves.

Algebraic groups play a prominent role in this thesis. The following definition, rather strange at first glance, is introduced in order to treat algebraic groups of dimension 1 as curves.

Definition 13. *An algebraic group curve is a projective curve equipped with a group law on a Zariski-dense subvariety called the group of G . The group of G is an algebraic group.*

A morphism between algebraic group curves which respects the group law is called an algebraic group morphism, or sometimes simply a group morphism.

Curves and morphisms of curves possess a number of integer invariants. Every curve has a *genus*. A non-constant morphism of curves has a *degree*, a *separable degree*, and an *inseparable degree*. A morphism may ramify at a point, in which case it has a *ramification index* greater than one. These invariants are interconnected by the Hurwitz formula. A reader who is not familiar with these should look at Silverman’s account in [Sil86] before trying to read Chapter 4.

Schemes.

In chapter 3, the construction of jet bundles is carried out in the formalism of sheaves and schemes. All notation is standard, taken from Grothendick’s [Gro60].

By *classical algebraic geometry*, we mean the study of varieties that does not involve schemes and, therefore, lives in the realm of first-order logic..

1.4.2 Definability in algebraic geometry

For the duration of this section, “definable” is in the pure field language.

Automorphisms

Every field of positive characteristic p has a definable automorphism $F : x \mapsto x^p$ called the Frobenius automorphism.

Definition 14. *If X is an algebraic object (curve, variety, morphism, algebraic group, etc) defined over a field k and τ is an automorphism of k , let X^τ denote the object defined by applying τ to all parameters in the definition of X . For a variety X , τ gives a bijection between points of X and those of X^τ .*

Note that even if τ is not definable, X^τ is as definable as X but with different parameters; however, the bijection given by τ is only as definable as τ . If τ is the Frobenius automorphism, this bijection is also a morphism of varieties.

Definable sets and functions

Since the theory of algebraically closed fields eliminates quantifiers, any definable set is quantifier-free definable. Any quantifier-free definable set is a finite union of varieties: this can be seen by putting the defining formula into disjunctive normal form.

Definable functions are nearly morphisms: the graph of a definable function $f : A \rightarrow B$ is a definable subset S , hence a union $\cup_i V_i$ of subvarieties, of the product $A \times B$ of the domain and the range. The projections to the domain A and to the range B are morphisms of varieties, and the projection onto the domain A is a bijection.

Each V_i is a graph of a definable function f_i whose domain is a subvariety A_i of A , and the projection from V_i to A_i is a bijective morphism of varieties. The only way that this bijective morphism can fail to be an isomorphism of varieties is by being inseparable. In that case, f_i can be written as a composition of a morphism of varieties with a negative power of the Frobenius isomorphism. The negative power of the Frobenius can be tacked onto the end or the beginning of the morphism.

This separation into pieces is necessary: a definable function could be defined by a case-out; for example, the function swapping two points and leaving the rest of the variety fixed is definable but is not a morphism of varieties. However, the cases themselves must be definable sets, hence can be taken to be subvarieties at the expense of increasing the number of cases. In particular, each irreducible component of the domain will always have precisely one case of the definition Zariski-dense in it:

Lemma 1.12. *If f is a definable function from an irreducible variety A , there is a morphism of varieties f_0 , an integer m , and a Zariski-dense subvariety A_0 of A such that $f = F^{-m} \circ f_0$ on A_0 .*

Several times we encounter a definable function from a curve to itself that has finite order under composition; such a function is automatically equal to a birational morphism on a Zariski-dense subvariety: a negative power of the Frobenius would contradict the finite order, since for a curve we may assume that f_0 is separable in the lemma above.

Since the theory of algebraically closed fields also eliminates imaginaries, a definable equivalence relation on the points of a variety gives rise to a definable function to a definable set. More precisely,

Lemma 1.13. *If V is a variety and $E \subset V \times V$ is a definable equivalence relation, then there exists a definable set S and a definable function $f : V \rightarrow S$ such that $f(v) = f(v')$ iff $(v, v') \in E$.*

As any definable set, S can be written as a union of varieties, and if V is irreducible and we are content with the domain of f being a Zariski-dense subvariety of V , and the image of f being Zariski-dense in S , then we can assume that S is itself an irreducible variety and f is equal to a morphism of algebraic varieties on a Zariski-dense subvariety, up to a negative power of the Frobenius automorphism. If we move the negative power of the Frobenius to the end, so that $f = F^{-n} \circ \hat{f}$ for some morphism of varieties \hat{f} , then the image of \hat{f} is S^{F^m} , another irreducible variety.

Any finite morphism of varieties can be written as a composition of a separable one and a purely-inseparable one; the purely inseparable one will not affect the fibers, so we may replace \hat{f} with its separable part. Since f and \hat{f} have the same fibers, we can refine the above lemma to:

Lemma 1.14. *If V is an irreducible variety and $E \subset V \times V$ is a definable equivalence relation with generically finite equivalence classes, then there exists a Zariski-dense subvariety V_0 of V , a variety W , and a finite separable morphism $g : V_0 \rightarrow W$ such that $g(v) = g(v')$ iff $(v, v') \in E$ for all v, v' in V_0 .*

Definable functions between curves and definable equivalence relations on curves are even nicer. For every rational morphism f between curves, there is a regular morphism g which is equal to f almost everywhere. Any morphism f between curves can be written as $f_0 \circ F^m$ and as $F^m \circ f_1$ for some power m of the Frobenius automorphism and some separable morphisms of curves f_0 and f_1 . Therefore, the lemma above can be further refined, brushing the Zariski-dense subvariety under the “almost everywhere” rug.

Lemma 1.15. *If V is a curve and $E \subset V \times V$ is a definable equivalence relation, then there exists a curve W , a separable morphism of curves $g : V \rightarrow W$, and a Zariski-dense subvariety U of V such that for all $v, v' \in U$, $g(v) = g(v')$ iff $(v, v') \in E$.*

Further, V and E determine W and g up to isomorphism: if another $g' : V \rightarrow W'$ satisfies the conclusion of this theorem, then there is an isomorphism $f : W \rightarrow W'$ such that $g' = f \circ g$.

Schemes

Schemes live outside first-order logic. It is important for us to stay inside first-order logic. When we make use of the scheme-theoretic construction of jet bundles, we need to know that the construction can be brought back down to first-order definability. How this happens is detailed in section 3.4.2.

1.4.3 Funny notions

Here we collect a few notions and lemmas that will be used repeatedly in the thesis.

We often need to put two morphisms together in one of the following ways:

Definition 15. *When $f_i : C \rightarrow D_i$, “ $f_1 \boxtimes f_2$ ” is the morphism from C to $D_1 \times D_2$ given by $x \mapsto (f_1(x), f_2(x))$.*

When $f_i : C_i \rightarrow D_i$, “ $f_1 \times f_2$ ” is the morphism from $C_1 \times C_2$ to $D_1 \times D_2$ given by $(x, y) \mapsto (f_1(x), f_2(y))$.

Generically injective

Definition 16. *A morphism of varieties $h : C \rightarrow E$ is called an initial compositional factor of another morphism of varieties $f : C \rightarrow A$ if there exists a morphism $f' : E \rightarrow A$ such that $f = f' \circ h$.*

Let $f : C \rightarrow A$ and $g : C \rightarrow B$ be two morphisms of varieties, and let $D \subset A \times B$ be the image of $f \boxtimes g$. Notation abuse alert: the morphism from C to D induced by $f \boxtimes g$ will also be called $f \boxtimes g$; this makes a difference for the notion of compositional factors, defined above. We investigate the connection between two related properties:

Proposition 1.16. *(A) holds if and only if (B) does not:*

- (A) There is a non-trivial morphism $h : C \rightarrow E$ which is an initial compositional factor of both f and g .
- (B) $(f \boxtimes g) : C \rightarrow D$ is birational.

First, a lemma:

Lemma 1.17. *The maximal initial compositional factor shared by f and g is $f \boxtimes g$. I.e.*

1. $f \boxtimes g$ is an initial compositional factor of both f and g .
2. If $h : C \rightarrow E$ is an initial compositional factor of both f and g , then it is also an initial compositional factor of $f \boxtimes g$.

Proof. 1. Let $f' : D \rightarrow A$ be the restriction of the projection $A \times B \rightarrow A$. Clearly $f = f' \circ (f \boxtimes g)$. Same story with g .

2. Suppose there are $f' : E \rightarrow A$ and $g' : E \rightarrow B$ such that $f = f' \circ h$ and $g = g' \circ h$. Then the image of $f' \boxtimes g' : E \rightarrow A \times B$ is D and $f \boxtimes g = (f' \boxtimes g') \circ h$.

□

Now the proposition follows immediately, as a morphism that has no non-trivial initial compositional factors (including itself!) is birational.

Definition 17. *When f and g satisfy (B) (and, therefore, not (A)) above, we say that $f \boxtimes g$ is generically injective.*

Full fiber product

Definition 18. *Given two morphisms of varieties $g : B \rightarrow D$ and $f : C \rightarrow D$, the fiber product P_0 , in the category of varieties, of B and C over D is the subvariety of $C \times B$ defined by $(b, c) \in P_0$ iff $g(b) = f(c)$.*

If f and g are morphisms of curves, and P_0 happens to be irreducible, then its normalization P is the fiber product in the category of curves. It comes with morphisms of curves $f'' : P \rightarrow B$ and $g'' : P \rightarrow C$ such that

1. $f'' \boxtimes g''$ is generically injective.
2. $\deg(f'') = \deg(f)$ and $\deg(g'') = \deg(g)$.

3. The following commutes:

$$\begin{array}{ccc} P & \xrightarrow{f''} & B \\ \downarrow g'' & & \downarrow g \\ C & \xrightarrow{f} & D \end{array}$$

4. For almost all pairs (b, c) such that $g(b) = f(c)$ there exists a unique $a \in P$ such that $f''(a) = b$ and $g''(a) = c$.

Lemma 1.18. *Suppose the following is a commutative diagram of separable morphisms of curves:*

$$\begin{array}{ccc} A & \xrightarrow{f'} & B \\ \downarrow g' & & \downarrow g \\ C & \xrightarrow{f} & D \end{array}$$

Then the following are equivalent:

- *A is the fiber product of B and C over D and this diagram witnesses it.*
- *Conditions (1) and (2) above.*
- *Condition (4) above.*

Proof. The fiber product obviously satisfies all the conditions.

Condition (4) gives the birational isomorphism between A and P_0 , the fiber product in the category of varieties.

Conditions (1) and (2) conspire to make condition (4) true: (1) gives uniqueness, and uniqueness together with (2) gives existence. \square

Factors of algebraic group curve morphisms

Lemma 1.19. *Compositional factors of a group morphism between curves are group morphisms.*

First, let us explain what we mean: Suppose that two curves A and B carry an algebraic group structure, and a morphism $\theta : A \rightarrow B$ is a group morphism. Suppose further that there is a curve C and morphisms

$$A \xrightarrow{\phi} C \xrightarrow{\psi} B$$

such that $\theta = \psi \circ \phi$. We show that C can be given an algebraic group structure that makes ϕ and ψ group morphisms.

Proof. Note that $|\ker(\theta)| = \deg_s(\theta)$ is finite. If $\phi(x) = \phi(y)$, then $\theta(x) = \theta(y)$, so $y - x \in \ker(\theta)$. So for each x in G , for some $u \in \ker(\theta)$ $\phi(x) = \phi(x + u)$. There are only finitely many choices for u , so some must occur infinitely often; any that occur infinitely often must occur almost everywhere. This means that almost all fibers are closed under addition of u , so the values of u that occur almost everywhere form a (finite) subgroup A_0 of $\ker \theta$. Almost all fibers of ϕ are cosets of A_0 , so the image of ϕ is isomorphic to the algebraic group curve A/A_0 . Then A_0 is $\ker(\phi)$, and ψ is the quotient by $\phi(\ker(\theta))$, with Frobenii added to taste if the morphisms aren't separable. \square

Chapter 2

Encoding

The purpose of this chapter is to characterize minimal sets of a special form that have something to do with groups. The special form is not too special: any quantifier-free definable minimal set is a union of finitely many such pieces, and any definable minimal set admits a finite cover by a quantifier-free-definable one. The characterization is useful because it reduces the question of group-likeness to pure algebraic geometry, enabling the rest of this thesis to use power tools from that area. The purpose of this chapter is

Theorem 3. *If*

A and B are irreducible varieties,

B is a finite-to-finite correspondence between A and A^σ ,

$(A, B)^\sharp = \{x \in A \mid (x, \sigma(x)) \in B\}$ is a minimal set,

and one of its very generic types is group-like,

then *there are algebraic groups G and $\Gamma \leq G \times G^\sigma$, varieties C and D , and dominant finite morphisms of varieties making the following diagram commute:*

$$\begin{array}{ccccc}
 G & \longleftarrow & \Gamma & \longrightarrow & G^\sigma \\
 \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
 C & \longleftarrow & D & \longrightarrow & C^\sigma \\
 \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
 A & \longleftarrow & B & \longrightarrow & A^\sigma
 \end{array}$$

In this diagram, the horizontal arrows are projections: $B \subset A \times A^\sigma$, $D \subset C \times C^\sigma$, and $\Gamma \subset G \times G^\sigma$, and the two middle vertical arrows are restrictions of $\pi \times \pi^\sigma$ and $\rho \times \rho^\sigma$ to D .

Some words in this theorem will be defined later in the chapter, such as the notion of a “very generic” type (Definition 20).

Most of the results in this chapter are rather technical, and easier to prove than to state, so there is no outline. None of the results in this chapter are new, but as usual the details are scattered or missing in the literature.

An immediate consequence of this theorem is that a minimal set of the form $(A, B)^\sharp$ for irreducible A and B , cannot be mixed: if one of its very generic types is group-like, then all of its very generic types are group-like.

This thesis only deals with the case where A is a curve and B is the graph of a function. We will make use of these restrictions only in the last section of this chapter.

2.1 A collection of well-known facts about sharps.

This section provides details of a construction mentioned in [CH99] and described in more detail in [Hru01].

2.1.1 Sharps

Definition 19. *Suppose V, W are affine varieties with $W \subset V \times V^\sigma$.*

$$(V, W)^{sh} := \{x \in V \mid (x, \sigma(x)) \in W\}$$

If in addition the projection of W to V is Zariski-dense in V , and the projection of W to V^σ is Zariski-dense in V^σ , we write $(V, W)^\sharp$ instead of $(V, W)^{sh}$ and say that the pair (V, W) encodes the quantifier-free definable set $(V, W)^\sharp$.

If $f : V \rightarrow V^\sigma$ is a morphism, we will write $(V, f)^\sharp$ instead of $(V, (\text{graph of } f))^\sharp$.

We will also use this notation when W is not a subset of $V \times V^\sigma$, but comes with morphisms to V and V^σ .

It is for sets of this form that we have a notion of a very generic subset:

Definition 20. $S \subset (V, W)^\sharp$ is called very dense in $(V, W)^\sharp$, if S is dense in V and the first prolongation of S (defined in the introduction, in the section about difference algebra) is dense in W . A very dense type is called very generic.

Definable sets encoded by a pair of varieties are particularly easy to handle, especially if the varieties are irreducible. For example, the following is an axiom of ACFA (or an immediate consequence of the axioms, depending on the axiomatization chosen):

Lemma 2.1. *If W is irreducible, then $(V, W)^\sharp \subset V$ is Zariski-dense in V and its first prolongation is Zariski-dense in W : $(V, W)^\sharp$ is very dense in itself.*

We will show that any quantifier-free definable set is definably isomorphic to a finite union of sets encoded by irreducible varieties. We first work towards showing that any quantifier-free definable set is definably isomorphic to a finite union of sets of the form $(V, W)^\sharp$ for some V, W .

Lemma 2.2. *Any set defined by a quantifier-free formula $\phi_0(y)$ is definably isomorphic to one defined by $\phi(x) := \tilde{\phi}(x, \sigma(x))$, where $\tilde{\phi}$ is (a quantifier-free formula) in the language of fields, i.e. doesn't mention σ .*

Proof. $\phi_0(y) = \tilde{\phi}_0(\sigma^m(y), \sigma^{m+1}(y), \dots, \sigma^{m+n}(y))$ for some integers m and n and some formula $\tilde{\phi}_0$ that doesn't mention σ .

Let $x = f(y) := (\sigma^m(y), \sigma^{m+1}(y), \dots, \sigma^{m+n-1}(y))$;

then $y = g(x) := \sigma^{-m}(x_0)$, where x_0 is the first component of x .

f and g are definable bijections between the set defined by $\phi_0(y)$ and the set defined by $\phi(x) = \tilde{\phi}(x, \sigma(x))$, where $\tilde{\phi}(x, x')$ is the formula

$$\tilde{\phi}_0(x_0, x_1, \dots, x_{n-1}, x'_{n-1}) \wedge (\wedge_{i=0}^{i=n-2} \sigma(x_i) = x_{i+1}) \wedge (\wedge_{i=0}^{i=n-2} x'_i = x_{i+1})$$

□

Lemma 2.3. *Any quantifier-free definable set is definably isomorphic to a finite union of set of the form $(V_0, W_0)^{sh}$.*

Proof. As in algebraic geometry, every quantifier-free formula is equivalent to a disjunction of formulas of the form

$$\phi(x) := (\wedge_i p_i(x) = 0) \wedge (\wedge_j q_j(x) \neq 0)$$

for some difference polynomials p_i, q_j . We show that each of these is definably isomorphic to $(V_0, W_0)^{sh}$ for some V_0, W_0 . Using the previous lemma, we may assume without loss of generality that p_i and q_j only mention x and $\sigma(x)$, i.e. that there are (pure field language!) polynomials P_i, Q_j such that

$$p_i(x) = P_i(x, \sigma(x)) , q_j(x) = Q_j(x, \sigma(x))$$

If x is an n -tuple, let U be the subvariety of affine $(2n)$ -space defined by

$$(\wedge_i P_i(u_1, u_1, \dots, u_{2n}) = 0) \wedge (\wedge_j Q_j(u_1, \dots, u_{2n}) \neq 0)$$

Let V_0 be the projection of U onto the the first n coordinates, and let

$$W_0 := (V_0 \times V_0^\sigma) \cap U$$

Clearly, the set defined by ϕ is $(V_0, W_0)^{sh}$:

By definition of U , $\phi(x) \Leftrightarrow (x, \sigma(x)) \in U$.

Note that $(x, \sigma(x)) \in U \Rightarrow x \in V_0 \Rightarrow \sigma(x) \in V_0^\sigma$.

Therefore, $\phi(x) \Leftrightarrow ((x, \sigma(x)) \in U \wedge x \in V_0 \wedge \sigma(x) \in V_0^\sigma)$.

By definition of W_0 , $\phi(x) \Leftrightarrow ((x, \sigma(x)) \in W_0 \wedge x \in V_0)$.

So $\phi(x) \Leftrightarrow x \in (V_0, W_0)^{sh}$.

□

Let V_0^s be the projection of W_0 onto the the last n coordinates. If $V_0^s = V_0^\sigma$, ϕ defines $(V_0, W_0)^\sharp$, and we are done. But it may be that we lost information as we passed from difference equations to algebraic equations, forgetting that u_{n+i} is supposed to be $\sigma(u_i)$. For example, if $n = 2$ and $\phi(x)$ is $x_1 = 17 \wedge x_1 + \sigma(x_1) + \sigma(x_2) = 51$, then V_0 is given by $u_1 = 17$, while V_0^s is given by $u_3 + u_4 = 34$.

Indeed, we are done if $V_0^s \cap V_0^\sigma$ is Zariski-dense in each of V_0^s and V_0^σ : then we can take $V^\sigma = V_0^s \cup V_0^\sigma$, and $W = W_0$. Otherwise, we need more work.

Lemma 2.4. *Suppose $V_0 \subset \mathbb{A}^n$, $W_0 \subset V_0 \times V_0^\sigma$, and V_0^s is the projection of W_0 onto the last n coordinates. If $V_0^s \cap V_0^\sigma$ is not Zariski-dense in V_0^σ , or not Zariski-dense in V_0^s , then there exist $V_1 \subset V_0$, $W_1 \subset W_0$ such that $(V_0, W_0)^{sh} = (V_1, W_1)^{sh}$, and V_1 is not Zariski-dense in V_0 .*

Proof. Suppose that $V_0^s \cap V_0^\sigma$ is not Zariski-dense in V_0^σ . The other half of the proof, when it is not Zariski-dense in V_0^s , is essentially the same.

Let $U_0 := V_0 \cap \sigma^{-1}(V_0^s)$, not Zariski-dense in V_0 (apply σ to first sentence of the proof of this lemma).

Let $W_1 := W_0 \cap (U_0 \times U_0^\sigma)$.

Let $V_1 \subset U_0$ be the projection of W_1 onto the the first n coordinates.

Note that $(V_0, W_0)^{sh} = (V_1, W_1)^{sh}$.

□

We now have all the ingredients:

Proposition 2.5. *Any quantifier-free definable set is definably isomorphic to a finite union of sets of the form $(V, W)^\sharp$.*

Proof. Lemma 2.2 provides the isomorphism. Lemma 2.3 reduces to sets of the form $(A, B)^{sh}$. Lemma 2.4 provides the induction step: there are no infinite descending chains of varieties $V_0 \supset V_1 \supset V_2 \dots$ such that each V_{i+1} is not Zariski-dense in V_i . The discussion preceding lemma 2.4 furnishes the base case for the induction. □

We now turn to the question of irreducibility.

Proposition 2.6. *Any $(V, W)^\sharp$ is a union of $(V_k, W_k)^\sharp$, where all V_k, W_k are irreducible.*

Proof. We proceed again by induction on the length of a (necessarily finite!) chain of non-Zariski-dense subvarieties of V .

For the base case of the induction, note that proposition is stupidly true when the dimension of V is zero.

Of course, $(V, W)^\sharp = \cup_i (V, W_i)^{sh}$ where W_i are irreducible components of W , but the projections of W_i may fail to be Zariski-dense in V or V^σ . In that case, Proposition 2.5 says that $(V, W_i)^{sh} = (V'_i, W'_i)^\sharp$, where V'_i is not Zariski-dense in V .

So,

$$(V, W)^\sharp = (\cup_{i \in I} (V, W_i)^\sharp) \cup (\cup_{i \in R} (V'_i, W'_i)^\sharp)$$

where

- Each W_i is irreducible. Since W_i projects dominantly onto V , this means V is also irreducible, so the sets in the first part of the union are already encoded by irreducible varieties.

- Each V'_i is not Zariski-dense in V . By induction, each $(V'_i, W'_i)^\sharp$ can be written as a finite union $\cup_j (V_{ij}, W_{ij})^\sharp$, where all V_{ij}, W_{ij} are irreducible.

Then the desired decomposition is

$$(V, W)^\sharp = (\cup_{i \in I} (V, W_i)^\sharp) \cup (\cup_{i \in R} \cup_j (V_{ij}, W_{ij})^\sharp)$$

□

Note that even for irreducible V, W , the set $(V, W)^\sharp$ may be σ -reducible: it is possible that the non-primeness of the difference ideal generated by the equations of $(V, W)^\sharp$ is only witnessed by difference polynomials involving further transforms $\sigma^N(x)$.

2.1.2 Definability discussion

We have just shown that any quantifier-free definable set is isomorphic to a finite union $\cup_k (V_k, W_k)^\sharp$, where all V_k, W_k are irreducible. This construction shows that (the parameters in the equations defining) V_k and W_k can be recovered definably from (the parameters in the equations defining) the original quantifier-free definable set, but this brushes three hard algebraic-geometry assertions under the rug.

Firstly, for each induction step in obtaining the varieties (V, W) encoding the minimal set, we need to get, uniformly definably, the equations defining the projection of a variety from the parameters in equations defining the variety itself. The existence of this procedure follows (by compactness) from quantifier elimination in the theory of algebraically closed fields. The procedures actually used to compute such things are studied in elimination theory. Secondly, to carry out the whole induction obtaining the varieties (V, W) encoding the minimal set, we need to predict how many steps this induction will take; in particular, we need a bound on the number of irreducible components of various dimensions that appear when dimension decreases. The existence of the bound again follows from compactness, and the actual bounds have been computed by Hrushovski in his [Hru01]. Thirdly, to recover the (equations of the) finitely many pieces encoded by irreducible varieties, we also need to know that irreducible components of a variety can be definably recovered from the variety itself; this is a hard theorem of van den Dries ([vdDS84]).

This thesis does not need this and will not prove it carefully. The algorithm to be constructed in this thesis takes in data describing a very special kind of quantifier-free definable set, which is already encoded by a pair of varieties. In that special case,

(uniformly definably) extracting the finitely many pieces of the set that are encoded by irreducible varieties is also easier, and only needs van den Dries' result on definability of irreducible components.

2.1.3 Finite rank

Proposition 2.7. *Any finite-rank, definable or type-definable set S is definably isomorphic to a very dense subset $S(N)$ of some $(V, W)^\sharp$, where W is a finite-to-finite correspondence between V and V^σ*

If the σ -degree of S and the quantifier-free formula witnessing it are known, then V and W can be produced uniformly definably from those polynomials, much as in the previous section. Otherwise, we have to settle for this proposition, which makes no claims of uniformity.

Proof. Let N be the σ -degree of S , so that $\sigma^{N+1}(s)$ is in the *field-theoretic* algebraic closure of $(s, \sigma(s), \dots, \sigma^N(s))$, for all $s \in S$, and then s is in the field-theoretic algebraic closure of $(\sigma(s), \sigma^2(s), \dots, \sigma^{N+1}(s))$. Let $S(N)$ be the N th prolongation of S , i.e. $\{(s, \sigma(s), \dots, \sigma^N(s)) \mid s \in S\}$, which is clearly definably isomorphic to S . Let V be the Zariski closure of $S(N)$. Let W be the Zariski closure of

$$X := \{(s, \sigma(s), \dots, \sigma^N(s); \sigma(s), \sigma^2(s), \dots, \sigma^{N+1}(s)) \mid s \in S\} \subset V \times V^\sigma$$

It is clear that W is a finite-to-finite correspondence, since the polynomial equations witnessing that

$$\sigma^{N+1}(s) \in \text{acl}(s, \sigma(s), \dots, \sigma^N(s)) \text{ and } s \in \text{acl}(\sigma(s), \dots, \sigma^{N+1}(s))$$

hold on X and hence on its Zariski closure W . □

In particular, every minimal set or type is isomorphic to a very dense subset of a set encoded by finite-to-finite correspondence.

Lemma 2.8. *If W is a finite-to-finite correspondence between V and V^σ , and both V and W are irreducible, then the dimension of the variety V is the σ -degree of $(V, W)^\sharp$.*

Proof. It is obvious that the dimension of V is an upper bound for the σ -degree of $(V, W)^\sharp$. To show that it is also a lower bound, note that if U is a not-Zariski-dense subvariety of V , then $W' := W \cap (U \times U^\sigma)$ is a non-Zariski-dense subvariety of W . So $\exists x \in (V, W)^\sharp - W'$ is an axiom of ACFA; note that this $x \notin U$. By compactness, for any given (small) base field L , there exists some $a \in (V, W)^\sharp$ which does not lie on *any* not-Zariski-dense subvariety of V defined over L , which implies that the transcendence degree of a over L (i.e. the σ -degree of a over L) is already equal to the dimension of V . \square

Another pleasant consequence of finite rank is that definable interalgebraicity between finite-rank definable sets extends to algebraic interalgebraicity between underlying varieties.

Proposition 2.9. *Suppose that A, B, C , and D are varieties, S and T are sets, and θ is a formula. Suppose further that*

- B is a finite-to-finite correspondence between A and A^σ ;
- D is a finite-to-finite correspondence between C and C^σ ;
- S is a very dense subset of $(A, B)^\sharp$;
- T is a very dense subset of $(C, D)^\sharp$;
- $\theta(x; y)$ witnesses that S and T are uniformly interalgebraic: there is a uniform bound N such that

$$\forall s \in S \exists_{\leq N} t \in T \theta(s, t)$$

and vice versa.

Then there is a (quantifier-free) formula in the language of fields $\zeta(x; y)$, and Zariski-dense subvarieties $A_1 \subset A$ and $C_1 \subset C$ such that

1. $\theta(x, y) \wedge x \in S \wedge y \in T$ implies $\zeta(x, y)$;
2. $\zeta(x; y)$ defines an algebraic finite-to-finite correspondence between $A_1 \subset A$ and $C_1 \subset C$;
3. ζ takes B to D

The proof of this proposition is broken into three lemmas.

Lemma 2.10. *Under the assumptions in the proposition, θ can be replaced by a quantifier-free, pure field language ζ such that θ implies ζ and ζ still witnesses the interalgebraicity between S and T .*

Proof. Take an arbitrary $a \in S$. Then there is some $c \in T$ such that $\theta(a, c)$ holds, so that $a \in \text{acl}(c)$. The characterization of model-theoretic closure from Proposition 1.7 and paragraph 1.8 in [CH99] implies that a is in the *field-theoretic* algebraic closure of $\{c, \sigma(c), \sigma^2(c), \dots\}$. But, since $(c, \sigma(c)) \in D$, and D is a finite-to-finite correspondence, $\sigma(c)$ is already in the field-theoretic algebraic closure of c , and so are all the $\sigma^n(c)$. So a is in the *field-theoretic* algebraic closure of c , i.e. there is a (quantifier-free) formula $\zeta_0(x, y)$ in the pure field language witnessing that $a \in \text{acl}(c)$. Exactly the same reasoning produces a (quantifier-free) formula $\zeta_1(x, y)$ in the pure field language witnessing that $c \in \text{acl}(a)$. Let ζ_a be $\zeta_0 \wedge \zeta_1$. This construction may produce different formulas for different a , but by compactness, some ζ will work for all a . \square

We now show that ζ defines a finite-to-finite correspondence, not only between S and T but also between A and all of C . We may and will assume without loss of generality that $\zeta(x, y)$ implies $x \in A \wedge y \in C$.

Lemma 2.11. *$\zeta(x; y)$ defines an algebraic finite-to-finite correspondence between $A_1 \subset A$ and $C_1 \subset C$;*

Proof. The (pure field language!) formula $\exists^{\leq N} y \zeta(x, y)$ holds for all $s \in S$, so it holds on a Zariski-dense subset of A . Therefore, it holds on a Zariski-dense subvariety A_0 of A . The same argument produces a Zariski-dense subvariety $C_0 \subset C$ with the same property. Note that $S \subset A_0$ and $T \subset C_0$.

We are almost done, but not quite: after we add the inequations defining A_0 and C_0 to ζ , it will be true that for every $c \in C_0$ there are at most N $a \in A_0$ such that $\zeta(a, c)$ holds, but there may not be any such a ; and vice versa. To lighten notation in the rest of this paragraph, we will say that $a \in A$ and $c \in C$ are *friends* if $\zeta(a, c)$ holds. If a has no friends left in C_0 , we say that a is *friendless*, and similarly for c who has no friends left in A_0 . The collection of people in A_0 who have no friends left in C_0 is a Zariski-codense subvariety of A_0 , because none of the people in S are friendless: they have friends in $T \subset C_0$. So we can throw out the Zariski-codense collection of friendless people from A_0 and let $A_1 = A_0 - \{\text{the friendless}\}$. Similarly let $C_1 = C_0 - \{\text{the friendless}\}$, and note

that people in A_1 still have friends in C_1 , because they had friends in C_0 and their friends weren't friendless, hence weren't thrown out. \square

We can now state assertion 3 in Proposition 2.9 precisely and prove it:

Lemma 2.12. *Let $E_0 \subset A_0 \times C_0$ be the Zariski closure of the set defined by ζ , a closed subvariety of $A_0 \times C_0$. Then there is an algebraic finite-to-finite correspondence $F_0 \subset E_0 \times E_0^\sigma$ which projects dominantly onto B and onto D .*

Proof. Let F_0 be the collection of points $(a, c; a', c')$ in $(A_0 \times C_0) \times (A_0^\sigma \times C_0^\sigma)$ such that:

- $(a, c) \in E_0$,
- $(a', c') \in E_0^\sigma$,
- $(a, a') \in B$, and
- $(c, c') \in D$.

Note that the last two requirements imply that a and a' are field-theoretically interalgebraic, and so are c and c' ; this shows that $F - 0$ is a finite-to-finite correspondence between E_0 and E_0^σ . We show that F_0 projects dominantly onto B , i.e. that

$$\phi(x, x') := \exists(y, y') \in D (x, y, x', y') \in F_0$$

holds on a Zariski-dense subvariety of B . This is a formula in the pure-field language, so it is enough to show that it holds on a Zariski-dense subset of B . We now show that it does hold on $S(1)$, the first prolongation of S . For $(a, \sigma(a)) \in S(1)$, there exists $c \in T$ such that $\models \theta(a, c)$, and hence also $\models \zeta(a, c)$, so $(a, c) \in E_0$. Since $c \in T$, $(c, \sigma(c)) \in D$. The last piece follows by applying the field automorphism σ to the formula $(a, c) \in E_0$. \square

The proposition is now proved.

If B and D are irreducible, $S = (A, B)^\sharp$ and $T = (C, D)^\sharp$ is a special case of this proposition. The result is useful much more generally, though, because such A , B , C , and D can be found for any finite-rank definable or type-definable sets S and T .

Note that, even in the special case, this is as good as it gets: we cannot get $\theta = \zeta \cap \left((A, B)^\sharp \times (C, D)^\sharp \right)$. For example, σ is a definable map with finite fibers between any $(A, B)^\sharp$ and $(A^\sigma, B^\sigma)^\sharp$, but not every A admits an algebraic-geometry morphism to A^σ .

For example, A could be a sufficiently general high-genus curve with A^σ not isomorphic to A .

The purpose of all this is to relate two distinct concepts:

1. A chunk of $(A, B)^\sharp$ is definably interalgebraic with a chunk of $(C, D)^\sharp$.
2. There is an algebraic finite-to-finite correspondence E between A and C that takes B to D .

It is useful to characterize (1) in terms of (2), because (2) lives entirely inside algebraic geometry, where power tools from algebraic geometry can be used.

One immediate useful consequence of this characterization is

Proposition 2.13. • *If $(A, B)^\sharp$ is definably interalgebraic with $(C, D)^\sharp$, then they have the same σ -degree.*

- *Two non-orthogonal types have the same σ -degree*

Proof. The σ -degrees of $(A, B)^\sharp$ and $(C, D)^\sharp$ are the dimensions of the varieties A and C , respectively. Since A and C are interalgebraic as varieties, these dimensions must be equal.

The second assertion follows immediately from the first and Proposition 2.7. \square

2.2 Group-like minimal types

This section is devoted to proving

Proposition 2.14. *For any group-like minimal type p there exist algebraic groups E and Γ and a type \hat{p} such that*

- *\hat{p} is a minimal type non-orthogonal to p ;*
- *$\hat{p} \subset (E, \Gamma)^\sharp$;*
- *\hat{p} is Zariski-dense in E , and its first prolongation is Zariski-dense in Γ ;*
- *Γ is a finite-to-finite correspondence between E and E^σ .*

Proof. By definition of “group-like”,

Step 2.2.1. *Let p_1 be a (model-theoretically) generic type of a definable minimal modular group G , with p non-orthogonal to p_1 .*

[CH99] gives the following characterization of definable minimal modular groups in ACFA:

Fact 2.15. *If G is a minimal modular group definable over a small model of ACFA, then there exists an algebraic group A and a Zariski-dense, definable, minimal subgroup $H \leq A$, which is interalgebraic with (indeed, finitely covered by) a finite-index subgroup G_0 of G .*

Any type in G is a type in one of the finitely many cosets of G_0 : since we're working over a small model, we have names for the cosets. Any type p_1 in a coset $G_0 + g$ is non-orthogonal to $p_2 := p_1 - g$, a type in G_0 .

Step 2.2.2. • p_2 is a minimal type non-orthogonal to p ;

- p_2 is a generic type of a definable minimal group G_0 ;
- G_0 is interalgebraic with a definable minimal group H ;
- H is a Zariski-dense subgroup of an algebraic group A .

Let p_3 be the type in H which is interalgebraic with p_2 . More formally, let $\theta(x, y)$ witness the interalgebraicity between H and G_0 and let

$$p_3(y) := \{\exists x \theta(x, y) \wedge \phi(x) \mid \phi(x) \in p_2(x)\}$$

Now p_3 is a generic type of H , since interalgebraicity preserves rank.

Step 2.2.3. • p_3 is a minimal type non-orthogonal to p ;

- p_3 is a generic type of a definable minimal group H ;
- H is a Zariski-dense subgroup of an algebraic group A .

We will repeatedly use

Fact 2.16. *If $H \leq G$ is an abstract subgroup of a topological group G , then the (topological) closure of H in G is also a subgroup of G .*

Let us carry out the proof of Proposition 2.7 inside A and its prolongations. Since p_3 is a minimal type, it has finite σ -degree, so for some N ,

$$\forall a \models p_3 \quad \sigma^{N+1}(a) \in \text{acl}(a, \sigma(a), \dots, \sigma^N(a))$$

Let $E \subset A \times A^\sigma \times \dots \times A^{\sigma^N}$ be the Zariski-closure of the N th prolongation of p_3 , i.e. of $\{(a, \sigma(a), \dots, \sigma^N(a)) \mid a \models p_3\}$. Clearly, this prolongation is a subgroup, so its Zariski closure is a subgroup also. Similarly, let Γ be the Zariski closure of

$$\{(a, \sigma(a), \dots, \sigma^N(a); \sigma(a), \sigma^2(a), \dots, \sigma^{N+1}(a)) \mid a \models p_3\} \subset E \times E^\sigma$$

Again, this is the Zariski closure of a subgroup, hence itself a subgroup. As in the proof of Proposition 2.7, Γ is a finite-to-finite correspondence, as wanted.

Let

$$\hat{p}(x_0, \dots, x_N) := p_3(x_0) \wedge (\bigwedge_{i=0}^{N-1} x_{i+1} = \sigma(x_i))$$

By construction, \hat{p} is very dense in $(E, \Gamma)^\sharp$. It is a prolongation of, hence interdefinable with, p_3 , which is interalgebraic with p , and the proof of the proposition is finished. \square

2.3 The diagram appears

We finally have the technical tools to produce the diagram we desire.

Suppose that B a finite-to-finite correspondence between A and A^σ , and p is a group-like minimal type, very generic in $(A, B)^\sharp$. Since p is group-like, it is (uniformly) definably interalgebraic with some \hat{p} , which is very generic in $(G, \Gamma)^\sharp$, for some algebraic groups G and Γ with Γ a finite-to-finite correspondence between G and G^σ (Proposition 2.14). Now we can apply Proposition 2.9 to $S = p$, $T = \hat{p}$ to conclude that there are varieties C and D and finite morphisms of varieties π and ρ making diagram 3 commute:

$$\begin{array}{ccccc} G & \longleftarrow & \Gamma & \longrightarrow & G^\sigma \\ \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\ C & \longleftarrow & D & \longrightarrow & C^\sigma \\ \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\ A & \longleftarrow & B & \longrightarrow & A^\sigma \end{array}$$

In this diagram, the horizontal arrows are projections, since $B \subset A \times A^\sigma$, $D \subset C \times C^\sigma$, and $\Gamma \subset G \times G^\sigma$, and the two middle vertical arrows are restrictions of $\pi \times \pi^\sigma$ and $\rho \times \rho^\sigma$ to D .

We have now proved Theorem 3.

2.4 Stronger conclusions under extra assumptions

So far in this chapter we have dealt with all possible quantifier-free definable minimal sets, while this thesis only concerns minimal sets of σ -degree 1. In that case, we can make further reductions in the algebraic geometry. For the duration of this section, A and B are irreducible varieties of dimension 1. In this case, the distinction between “model-theoretically generic” and “very generic” becomes moot:

Lemma 2.17. *If A and B are irreducible varieties of dimension 1, then any non-algebraic type p in $(A, B)^\sharp$ is dense in A and its first prolongation is dense in B .*

Proof. Any infinite subset of an irreducible variety of dimension 1 is Zariski-dense. \square

Suppose further that some generic type in $(A, B)^\sharp$ is group-like. Then all the varieties in diagram 3 have dimension 1, because all arrows are finite and dominant. Since the (complete) type \hat{p} is very generic in $(G, \Gamma)^\sharp$, G and Γ in that diagram must also be irreducible. The proof of Lemma 2.12 then shows that some irreducible component D' of D must project dominantly onto both B and Γ . Letting C' be the projection of D' onto C , and noting that the commutativity of the diagram forces D' to project dominantly onto C'^σ , we obtain a new diagram of the same shape, made up of irreducible varieties of dimension 1.

Normalizing an irreducible variety of dimension 1, i.e. finding a curve (a smooth, projective, irreducible closed variety of dimension 1) birational to it, is a functor. Since curves are much easier to work with, we will apply this normalization functor to the entire diagram 3 before continuing.

This functor does not respect the notion “subset”. In a row

$$A \xleftarrow{f} B \xrightarrow{g} A^\sigma$$

of the original diagram, B was a subset of $A \times A^\sigma$, but the normalization $N(B)$ of B is merely birational to its image in $N(A) \times N(A^\sigma)$ under $N(f) \boxtimes N(g)$. What we can salvage is the generic injectivity property described in 1.4.3: since $f \boxtimes g$ was generically injective, $N(f) \boxtimes N(g)$ still is.

A version of Theorem 3 summarizes this discussion:

Theorem 4. *Suppose that*

A_0 and B_0 are irreducible varieties **of dimension 1**; and

B_0 is a finite-to-finite correspondence between A_0 and A_0^σ ; and

some (non-algebraic) type in (the minimal set) $(A_0, B_0)^\#$ is group-like.

Let A be the normalization of A_0 and B be the normalization of B_0 .

Then there are algebraic group curves G and $\Gamma \leq G \times G^\sigma$, curves C and D , and surjective finite morphisms making the following diagram commute:

$$\begin{array}{ccccc}
 G & \longleftarrow & \Gamma & \longrightarrow & G^\sigma \\
 \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
 C & \longleftarrow & D & \longrightarrow & C^\sigma \\
 \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
 A & \longleftarrow & B & \longrightarrow & A^\sigma
 \end{array}$$

Furthermore, the funny product \boxtimes of the two arrows in each row is generically injective, and so is $\pi \boxtimes \rho$.

2.4.1 A small reduction

This proposition belongs here more than anywhere else. It is too technical for the introduction, and too model-theoretic for chapter 4.

Proposition 2.18. *It is sufficient to prove Theorem 2 under an additional assumption that the function $x \mapsto P(x, b)/Q(x, b)$ is separable.*

Proof. The inseparable degree m and the separable component f' of a rational function f in one variable (i.e. integer m and rational function f' such that $f = f' \circ F^m$, where F is the Frobenius automorphism) can be recovered definably from f .

Let $\tau = F^{-m} \circ \sigma$; it is another automorphism of the field K , and the field K with the automorphism τ is also a model M of ACFA ([CH99]).

The minimal set defined by $\sigma(x) = f(x)$ is group-like in the original model if and only if the set defined by $\tau(x) = f'(x)$ is group-like in M .

Since f' is separable, we can evaluate the characterizing formula from the separable case of Theorem 2 in M to determine whether $\tau(x) = f'(x)$ is group-like in M . \square

We pause for a chapter to construct jet spaces and then in chapter 4 we examine the diagram created in Theorem 4, with the additional separability assumption (which does not lose generality) and the additional assumption that B is the graph of a function from A to A^σ (which does lose generality, but that is as much generality as we promised in chapter 1).

Chapter 3

Construction of jet spaces

Here we describe in great detail the construction of jet bundles. Jet spaces are one of the higher-order analogs of tangent space; they are different from arc spaces, even if some people use the word “jet spaces” to refer to arc spaces. Like tangent bundles, these vector “bundles” may fail to be locally trivial at singularities; we see no easy way around this bit of notation abuse, and we have the excuse that, in this thesis, we will only look at jet bundles of smooth varieties.

This construction has been known and used for a long time, but we could not find a detailed reference, or indeed any reference that asserted all the properties we need in the generality we need. Our construction borrows heavily from Moosa’s “Jet spaces in complex analytic geometry: an exposition” [Moo04]; Grothendieck’s “Techniques de construction en géométrie analytique VII. Étude locale des morphismes: éléments de calcul infinitésimal.” [Gro62] and EGA.II [Gro61]; and Moosa and Scanlon’s unpublished notes on differential jets and arcs [MS].

Although we will only use jet spaces for classical algebraic varieties, i.e. zero-sets of polynomials, it is easier to describe the construction of jet spaces in the language of schemes. Nevertheless, we try to work as naively as possible, in order to make this account as accessible as possible.

We construct covariant functors J^m , one for each for each $m \in \mathbb{N}$, that takes a scheme X over an algebraically closed base field k and returns a finite-dimensional, possibly not locally-trivial, k -vector bundle J_X^m over X , called the *m th jet bundle of X* . The fiber of J_X^m above a point $p \in X$ is a finite-dimensional k -vector-space, denoted $J_{X,p}^m$ and called the *m th jet space of X at p* . We show that these functors have the following

properties:

- If $f : X \rightarrow Y$ is a morphism of varieties, the morphism $J^m(f) : J_X^m \rightarrow J_Y^m$ is k -linear on each fiber of $J^m(X)$. (section 3.4.4)
- If $f : X \rightarrow Y$ is a constant morphism, then the induced linear maps on fibers of J_X^m are all zero. (section 3.5.1)
- If for two morphisms $f, g : X \rightarrow Y$, $f(p) = g(p)$ and for all m , the induced linear maps $J^m(f)_p$ and $J^m(g)_p$ on the fiber of J_X^m above p agree, then $f = g$. (section 3.5.3)
- If f is étale at p , then the linear map induced by f on the fiber of J_X^m above p is an isomorphism. (section 3.5.4)
- Higher jet spaces have more information: Suppose $f, g : X \rightarrow Y$ are morphisms of schemes, $f(p) = g(p) = q$ for some points $p \in X$ and $q \in Y$, and $J^n(f)_p = J^n(g)_p$. Then $J^m(f)_p = J^m(g)_p$ for all $m \leq n$. (section 3.5.2)

Sheaves and definability

It will also be important for us to know that the construction of jet bundles of classical varieties occurs inside the first-order theory of algebraically closed fields.

In classical algebraic geometry over a fixed algebraically-closed field k , a variety is a certain boolean combination of solution-sets of polynomials in a cartesian power k^n . Thus, classical varieties are first-order definable objects. Sheaves and schemes, on the other hand, live outside the first-order setting. In particular, local rings $\mathcal{O}_{X,p}$ are infinitary objects (direct limits) outside the first-order setting. Jet spaces are a remedy for this.

The sheaf-theoretic construction of jet spaces naturally occurs, and is easiest to describe, outside the first-order, classical algebraic geometry setting. However, we will see that when this construction is applied to a classical variety, everything remains first-order definable; which is the point of the exercise. Details and proofs are in section 3.4.2.

3.1 k -algebra functors

This section describes a functor that takes a k -algebra R and returns a complicated commutative diagram of k -algebras and their ideals. We are really only interested in some

of the objects in the diagram, namely the R -modules $J^m(R)$, but we need the rest of the diagram to construct them and understand them. This section is devoted to proving that each of the object in the diagram (both the algebras and the ideals) is a functor returning an R -module, some with extra structure.

Notation/Assumption 3.1.1. *For the rest of this chapter, k is an algebraically closed field, all rings are k -algebras, and all ring homomorphisms and tensor products are over k , unless otherwise stated.*

3.1.1 One ring

- Let R be ring. For example, R could be a function field over k .
- Let $T(R) := R \otimes R$, another ring.
- Let $\alpha_R : R \rightarrow T(R) : x \mapsto x \otimes 1$, a ring homomorphism making $T(R)$ an R -algebra.
- Let $\delta_R : T(R) \rightarrow R : x \otimes y \mapsto xy$, another ring homomorphism.

Note that $\delta_R \circ \alpha_R = \text{id}_R$, so δ_R is an R -algebra morphism.

- Let $I(R) := \ker(\delta_R)$, an ideal of $T(R)$.
- For each $m \in \mathbb{N}$, let $I^m(R) := (I(R))^m$, also an ideal of $T(R)$.
- For each $m \in \mathbb{N}$, let $A^m(R) := T(R)/I^{m+1}(R)$, a quotient ring.

E.g. $A^0(R) = R$.

- For each $m \in \mathbb{N}$, let $\delta_R^m : T(R) \rightarrow A^m(R)$, a quotient homomorphism. E.g. $\delta^0 = \delta$.

Note that $\delta_R^m \circ \alpha_R$ makes $A^m(R)$ into an R -algebra, and then δ_R^m is an R -algebra homomorphism.


- For each $m \in \mathbb{N}$, let $J^m(R) = I(R)/I^{m+1}(R) = \delta_R^m(I(R)) \trianglelefteq A^m(R)$, an ideal in $A^m(R)$, the image of $I(R)$ under the projection $T(R) \rightarrow A^m(R)$.

E.g. $J^0 = \{0\}$, and in general J^m is nilpotent of power $(m + 1)$.

- Since $I^n(R) \subset I^m(R)$ for every $n \geq m$, there are canonical morphisms $\pi^{(n,m)} : A^n(R) \rightarrow A^m(R)$.

Notation abuse alert: in these definitions, the superscript is only a label, not a power, except for I^m .

3.1.2 Commutative diagram for jets



(R) has been suppressed to reduce clutter. The homomorphisms $\pi^{(n,m)}$ are vertical arrows on the right. Dashed arrows are α and $\delta^k \circ \alpha$ for various k , making $T(R)$ and $A^k(R)$ into R -algebras, and their ideals $I(R)$ and $J^k(R)$ into R -modules. Left-to-right straight lines are exact: for example, I^2 is the kernel of both $T \rightarrow A^1$ and $I \rightarrow J^1$.

3.1.3 Two rings

We would like the construction in the previous section to be functorial. I.e. for each k -algebra morphism $f : R \rightarrow S$ and for each object Q in the commutative diagram 3.1.2, we would like to define a morphism $Q(f) : Q(R) \rightarrow Q(S)$, in the appropriate category for each Q , compatible with the arrows of 3.1.2. For example, $(T(f), I^m(f))$ should be a bimorphism, i.e. $I^m(f)$ should be a $T(f)$ -morphism.

- Let $T(f) := f \otimes f$, the homomorphism $T(R) \rightarrow T(S)$ that takes $r \otimes r'$ to $f(r) \otimes f(r')$.

This is a well-defined ring homomorphism, compatible with α and δ .

- I is a subset of T , so $I(f)$ should be the restriction of $T(f)$ to $I(R)$. We need to verify that the image of this restriction is a subset of $I(S)$. This follows by simple commutative algebra (Lemma 3.22) from the commutativity of

$$\begin{array}{ccccc} I(R) & \longrightarrow & T(R) & \xrightarrow{\delta_R} & R \\ & & T(f) \downarrow & & f \downarrow \\ I(S) & \longrightarrow & T(S) & \xrightarrow{\delta_S} & S \end{array}$$

- For each m , I^m is a subset of T , so $I^m(f)$ should be the restriction of $T(f)$ to $I^m(R)$. We already know that the image of $I(R)$ under $T(f)$ is a subset of $I(S)$, and we need to verify that the image of $I^m(R)$ under $T(f)$ is a subset of $I^m(S)$. This is again an easy ring theory exercise that does not rely on the particulars of our construction.
- Now for each m , $A^m(f)$ is given by the other direction of Lemma 3.22, since

$$\begin{array}{ccccc} I^{m+1}(R) & \longrightarrow & T(R) & \xrightarrow{\delta_R^m} & A^m(R) \\ I^{m+1}(f) \downarrow & & T(f) \downarrow & & \\ I^{m+1}(S) & \longrightarrow & T(S) & \xrightarrow{\delta_S^m} & A^m(S) \end{array}$$

already commutes. Note that the diagram above witnesses that $T(f)$ and $A^m(f)$ are compatible with δ^m s.

- $T(f)$ takes $I(R)$ to $I(S)$ and $I^{m+1}(R)$ to $I^{m+1}(S)$, so it induces a morphism $J^m(f)$ from $J^m(R)$ to $J^m(S)$.

3.1.4 Three rings

To finish functoriality, we need to know that given three rings and two morphisms $P \xrightarrow{f} R \xrightarrow{g} S$, for each object Q in diagram 3.1.2, $Q(g \circ f) = Q(g) \circ Q(f)$. This is clear.

Theorem 5. T, I, I^m, A^m , and J^m are functors that take a k -algebra R and return an R -module (I, I^m , and J^m) or an R -algebra (T and A^m).

Some of the resulting R -modules have more structure (R -algebras, $T(R)$ -modules), and this extra structure is respected by the functors.

3.2 Generators

This section is a collection of fairly obvious facts. They provide the finite-generatedness results which are necessary for parts of the sheaf-theoretic construction and crucial for first-order definability. First, some notation.

- N is a fixed positive integer.
- r will always be a single element of R .
- i will always be a single integer index.
- \mathbf{r} will denote an N -tuple from R .
- \mathbf{j} will be a N -tuple of positive integers.
- $\mathbf{r}^{\mathbf{j}} := \prod_i r_i^{j_i}$.
- For $x \in R$, $\widehat{x} := 1 \otimes x - x \otimes 1 \in T(R)$

For example, $\delta^1(\widehat{x})$ lives in $J^1(R)$, the usual module of differential 1-forms, where it is usually called “ dx ”.

- Note that $\widehat{\mathbf{r}^{\mathbf{j}}} = 1 \otimes \mathbf{r}^{\mathbf{j}} - \mathbf{r}^{\mathbf{j}} \otimes 1$ and $\widehat{\mathbf{r}^{\mathbf{j}}} = \prod_i (1 \otimes r_i - r_i \otimes 1)^{j_i}$ are different.

Proposition 3.1. Let $\{r_i \mid i \in N\}$ generate R as a k -algebra. Then

1. $\{\mathbf{r}^{\mathbf{j}}\}$ (including the empty product) generates R as a k -module.
2. $\{1, \widehat{\mathbf{r}^{\mathbf{j}}}\}$ generates $T(R)$ as an R -module.

Hence, $\{\mathbf{r}^{\mathbf{j}'}, \mathbf{r}^{\mathbf{j}'} \widehat{\mathbf{r}^{\mathbf{j}}}\}$ generates $T(R)$ as an k -module.

3. $\{1, \widehat{\mathbf{r}}^{\mathbf{j}}\}$ also generates $T(R)$ as an R -module.

Hence, $\{\mathbf{r}^{\mathbf{j}'}, \mathbf{r}^{\mathbf{j}'} \widehat{\mathbf{r}}^{\mathbf{j}}\}$ generates $T(R)$ as a k -module.

4. $\{\widehat{\mathbf{r}}^{\mathbf{j}} \mid \mathbf{j} \neq \emptyset\}$ generates $I(R)$ as an R -module.

Hence, $\{\mathbf{r}^{\mathbf{j}'} \widehat{\mathbf{r}}^{\mathbf{j}} \mid \mathbf{j} \neq \emptyset\}$ generates $I(R)$ as a k -module.

5. $\{\widehat{\mathbf{r}}^{\mathbf{j}} \mid \sum \mathbf{j} > m\}$ generates $I^{m+1}(R)$ as an R -module.

6. $\{\delta^m(\widehat{\mathbf{r}}^{\mathbf{j}}) \mid 1 \leq \sum \mathbf{j} \leq m\}$ generates $J^m(R)$ as an R -module.

Furthermore, if $\{r_i \mid i \in N\}$ is algebraically independent over k , each of those sets of generators is in fact a basis; and $T(R)$ is graded by $M^l(R)$, the R -module generated by $\{\widehat{\mathbf{r}}^{\mathbf{j}} \mid \sum \mathbf{j} = l\}$, the set of l -fold products of \widehat{r}_i .

Proof. 1. is clear.

2. By definition, $\{a \otimes b \mid a, b \in R\}$ generate $T(R)$ as a k -module.

$$\begin{aligned} a \otimes b &= (a \otimes 1)(1 \otimes b - b \otimes 1 + b \otimes 1) = \\ &= ab \otimes 1 + (a \otimes 1) \left(\sum_{\mathbf{j}} c_{\mathbf{j}} (1 \otimes \mathbf{r}^{\mathbf{j}} - \mathbf{r}^{\mathbf{j}} \otimes 1) \right) = \\ &= ab \otimes 1 + (a \otimes 1) \left(\sum_{\mathbf{j}} c_{\mathbf{j}} \widehat{\mathbf{r}}^{\mathbf{j}} \right) \end{aligned}$$

3. Easy computation verifies that

$$\widehat{r}s = \widehat{r}\widehat{s} + (r \otimes 1)(\widehat{s}) + (s \otimes 1)(\widehat{r})$$

Iterate this to write $\widehat{\mathbf{r}}^{\mathbf{j}}$ as an R -linear combination of 1 and $\widehat{\mathbf{r}}^{\mathbf{j}}$ s, and use (2).

4. By definition, I is the kernel of the R -linear

$$\delta : T(R) \rightarrow R : \sum_i a_i \otimes b_i \mapsto \sum_i a_i b_i$$

Note that $\delta(\widehat{x}) = 0$ for all x , and use the expansion in (3) to get

$$\delta \left(a \otimes 1 + \sum_{\mathbf{j}} (b_{\mathbf{j}} \otimes 1) \widehat{\mathbf{r}}^{\mathbf{j}} \right) = \delta(a \otimes 1) = a$$

So $w := \left(a \otimes 1 + \sum_{\mathbf{j}} (b_{\mathbf{j}} \otimes 1) \widehat{\mathbf{r}}^{\mathbf{j}} \right) \in I$ iff $\delta(w) = 0$ iff $a = 0$, as wanted. This also shows that the elements \widehat{r}_i generate I as a $T(R)$ -ideal.

5. I^{m+1} is generated as a $T(R)$ -ideal by $(m+1)$ -fold products of generators of I as a $T(R)$ -ideal, namely by $\{\widehat{\mathbf{r}}^{\mathbf{j}} \mid \sum \mathbf{j} = m+1\}$. So anything in I^{m+1} can be written as a $T(R)$ -linear combination of these $(m+1)$ -fold products $\sum_{\mathbf{j}} t_{\mathbf{j}} \widehat{\mathbf{r}}^{\mathbf{j}}$, where $t_{\mathbf{j}}$ are the coefficients in $T(R)$ and the sum is over multi-indices \mathbf{j} such that $\sum \mathbf{j} = m+1$. Expanding each $t_{\mathbf{j}}$ as in (3), we get the result.
6. δ^{m+1} is R -linear; $\delta^{m+1} : I \rightarrow J^m$ is a surjective map of R -modules.

The non-zero δ^{m+1} -images of the generators of the R -module I generate the R -module J^m .

$\{\widehat{\mathbf{r}}^{\mathbf{j}} \mid \mathbf{j} \neq \emptyset\}$ generates $I(R)$ as an R -module (4).

$\delta^{m+1}(\widehat{\mathbf{r}}^{\mathbf{j}}) = 0$ if $\sum \mathbf{j} \geq m+1$ (5).

“Furthermore” for $T(R)$ follows from the universality of the tensor product; direct computation then verifies the rest of “furthermore”.

□

3.3 Functors and (pre)sheaves

Proposition 3.2. *A functor F from a category C to a category D induces an associated functor \mathcal{F} that takes presheaves with values in C and returns presheaves with values in D .*

Proof. A functor F from a category C to a category D can be applied to a commutative diagram of objects and arrows of C to produce an identically-shaped commutative diagram in D . For example, applying F to a presheaf with values in C , one obtains a presheaf with values in D . For another example, we can apply F to the giant commutative diagram witnessing a morphism of two presheaves with values in C (i.e. the diagram made up of the two presheaves and the C -morphisms that make up the morphism between the two presheaves) to obtain the giant commutative diagram witnessing the morphism between the two resulting presheaves with values in D . □

We will apply the (presheaf versions of) functors defined in section 3.1 to the structure sheaf \mathcal{O}_X of a scheme X over k . The resulting presheaves on $\mathrm{Sp}(X)$, the under-

lying topological space of X , will be denoted \mathcal{T}_X , \mathcal{I}_X^m , \mathcal{A}_X^m and \mathcal{J}_X^m . \mathcal{T}_X is a presheaf of \mathcal{O}_X -algebras, via α (notation abuse alert: α denotes both the morphism of rings and the morphism of sheaves of rings on $\mathrm{Sp}(X)$). The \mathcal{I}_X^m are presheaves of ideals of \mathcal{T}_X . The \mathcal{A}_X^m are presheaves of \mathcal{O}_X -algebras, via $\delta^m \circ \alpha$. For each m , \mathcal{J}_X^m is a presheaf of nilpotent (of power $m + 1$) ideals of \mathcal{A}_X^m , and also a presheaf of \mathcal{O}_X -modules.

We are interested in the local properties of the sheafifications of these presheaves, so it will be sufficient to understand the behavior of these presheaves on affine pieces of a scheme.

Notation/Assumption 3.3.1. *For the duration of this section, R is a k -algebra, $X = \mathrm{Spec}R$ is an affine scheme; $S \subset R$ is a multiplicative subset, and in particular, for an open $D \subset R$, $S(D)$ is the (multiplicative) set of elements of R “invertible on D ”, meaning that their image (under the restriction homomorphism) in $\mathcal{O}_X(D)$ is invertible; and Q is one of our functors, i.e. a functor that takes a ring and returns a module over it, possibly with extra structure.*

The presheaf version \mathcal{Q} of Q applied to the structure sheaf \mathcal{O}_X produces a presheaf \mathcal{Q}_X of \mathcal{O}_X -modules.

On the other hand, there is the standard construction of a coherent sheaf $\widetilde{Q(R)}$ of \mathcal{O}_X -modules out of an R -module $Q(R)$. (For D an open subset of X , $\mathcal{O}_X(D) = (S(D))^{-1}R$ for some multiplicative subset $S(D)$ of R . $\widetilde{Q(R)}(D)$ is defined to be the localization of $Q(R)$ at $S(D)$.)

The presheaf \mathcal{Q}_X we constructed is closely related to the usual coherent sheaf $\widetilde{Q(R)}$:

Proposition 3.3. *If the functor Q commutes with localizations, then $\mathcal{Q}_X = \widetilde{Q(R)}$.*

Proof. For D an open subset of X ,

$$\mathcal{Q}_X(D) = Q(\mathcal{O}_X(D)) = Q((S(D))^{-1}R)$$

Since Q commutes with localizations,

$$Q((S(D))^{-1}R) = (S(D))^{-1}(Q(R))$$

By definition,

$$(S(D))^{-1}(Q(R)) = \widetilde{Q(R)}(D)$$

□

The purpose of this section is to show that J^m and A^m indeed commute with localizations.

3.3.1 Localizations and quasi-coherence

Let S be a multiplicative subset of a ring R . Let us apply the functors from section 3.1 to the canonical morphism $\eta : R \rightarrow S^{-1}R$ and use a few basic facts about localizations gathered in section 3.6.

Notation/Assumption 3.3.2. $S \otimes S := \{a \otimes b : a, b \in S\}$, a multiplicative subset of $T(R)$.

Lemma 3.4. • By definition $T(S^{-1}R) = S^{-1}R \otimes S^{-1}R = (S \otimes S)^{-1}(R \otimes R)$.

- $I(S^{-1}R) = (S \otimes S)^{-1}I(R)$ (Lemma 3.23).

- $I^m(S^{-1}R) = (S \otimes S)^{-1}I^m(R)$ (Lemma 3.25).

- $A^m(S^{-1}R) = (\delta_R^m(S \otimes S))^{-1}A^m(R)$ (Lemma 3.24).

(Remember, $\delta_R^m : T(R) \rightarrow A^m(R)$, $S \otimes S \subset T(R)$, so $\delta_R^m(S \otimes S) \subset A^m(R)$.)

- $J^m(S^{-1}R) = (\delta_R^m(S \otimes S))^{-1}J^m(R)$ because localizations commute with quotients.

Now it seems that we are in trouble. To lighten notation, let $S = S(D)$.

We just showed that $\mathcal{A}_X^m(D) := A^m(S^{-1}R) = (\delta_R^m(S \otimes S))^{-1}A^m(R)$.

On the other hand, $\widetilde{A(R)}(D) := S^{-1}(A(R)) = (\delta_R^m(S \otimes 1))^{-1}A^m(R)$.

To show that the two presheaves \mathcal{A}_X^m and $\widetilde{A(R)}$ are the same, we need one last lemma:

Lemma 3.5.

$$(\delta^m(S \otimes 1))^{-1}(A^m(R)) = (\delta^m(S \otimes S))^{-1}(A^m(R))$$

and

$$(\delta^m(S \otimes 1))^{-1}(J^m(R)) = (\delta^m(S \otimes S))^{-1}(J^m(R))$$

Proof. It is enough to show that every element of $\delta^m(S \otimes S)$ is already invertible in $(\delta^m(S \otimes 1))^{-1}(A^m(R))$.

$\delta^m(S \otimes S)$ is generated (as a multiplicative set) by $\delta^m(f \otimes 1)$ and $\delta^m(1 \otimes f)$, so it suffices to show that each of those is invertible in $(\delta^m(S \otimes 1))^{-1}(A^m(R))$.

$f \otimes 1 \in S \otimes 1$, hence $\delta^m(f \otimes 1)$ is invertible in $(\delta^m(S \otimes 1))^{-1}(A^m(R))$.

Note that $\delta^m(1 \otimes f - f \otimes 1)$ is nilpotent in $A^m(R)$ because it lies in $J^m(R)$.

Therefore, $\delta^m(1 \otimes f) = \delta^m(f \otimes 1) + \delta^m(1 \otimes f - f \otimes 1) = \text{invertible} + \text{nilpotent} = \text{invertible}$.

□

Similarly for J^m : we just showed that $\mathcal{J}_X^m(D) := J^m(R_S) = (\delta_R^m(S \otimes S))^{-1} J^m(R)$, where the localization is in the sense of $A^m(R)$ -modules.

On the other hand, $\widetilde{J^m(R)}(D) := (J^m(R))_S = S^{-1} J^m(R)$, where the localization is in the sense of R -modules, which is the same as $(\delta_R^m(S \otimes 1))^{-1} J^m(R)$, where the localization is in the sense of $A^m(R)$ -modules.

The same Lemma 3.5 now finishes the proof that $\mathcal{J}_X^m = \widetilde{J^m(R)}$.

3.3.2 Conclusions

Lemma 3.6. *If $X = \text{Spec}(R)$ is an affine variety, so that R is finitely-generated over k , then $A^m(R)$ and $J^m(R)$ are finitely-generated R -modules.*

Proof. This follows immediately from proposition 3.1. If $\{r_1, r_2, \dots, r_n\}$ are the finitely many generators of R as a k -algebra, then $\{\delta^m(\widehat{\mathbf{r}}^{\mathbf{j}}) \mid 1 \leq \sum \mathbf{j} \leq m\}$ generates $J^m(R)$ as an R -module, and $\{1, \delta^m(\widehat{\mathbf{r}}^{\mathbf{j}}) \mid \sum \mathbf{j} \leq m\}$ generates $A^m(R)$ as an R -module. □

Proposition 3.7. • $\mathcal{J}_X^m = \widetilde{J^m(R)}$. *In particular, it is a quasicoherent sheaf of \mathcal{O}_X -modules.*

• $\mathcal{A}_X^m = \widetilde{A^m(R)}$. *In particular, it is a quasicoherent sheaf of \mathcal{O}_X -algebras.*

• *on an affine variety, “quasicoherent” becomes “coherent”.*

Proof. The previous subsection contains most of the proof, and the last lemma allows us to say “coherent” instead of “quasicoherent”. □

Proposition 3.8. $\langle \text{Sp}(X), \mathcal{A}_X^m \rangle = \text{Spec}(A^m(R))$.

Proof. Remember that $A^m(R)$ is nilpotent over R , meaning that the ideal $J^m(R)$ of the R -algebra homomorphism $\pi^{(n,0)} : A^m(R) \rightarrow R$ is nilpotent. Therefore, the underlying topological spaces of $\text{Spec}(A^m(R))$ and $\text{Spec}(R)$ are canonically isomorphic; and the multiplicative subset of $A^m(R)$ corresponding to an open set D is precisely $\delta^m(\alpha(S))$, where S is the multiplicative subset of R corresponding to D . □

Theorem 6. For a scheme X associated to an affine variety from classical algebraic geometry, $\langle \mathrm{Sp}(X), \mathcal{A}_X^m \rangle = \mathrm{Spec} A^m(R)$ is of finite type over $X = \mathrm{Spec} R$, and \mathcal{J}_X^m is a coherent sheaf of \mathcal{O}_X -modules.

3.4 Jet Spaces

3.4.1 $\mathrm{Spec} \mathrm{Sym}$

Definition 21. For a scheme X , let the m th jet bundle of X be $J_X^m = \mathrm{Spec}(\mathrm{Sym}(\mathcal{J}_X^m))$, the standard construction of a scheme over X out of a quasi-coherent sheaf of \mathcal{O}_X -modules.

Notation abuse alert:

- $J^m(R)$ is an R -module,
- \mathcal{J}_X^m is a sheaf of \mathcal{O}_X -modules,
- J_X^m is a (possibly not locally-trivial) vector bundle over X .

This construction is described in detail in EGA.II.1.7: A coherent sheaf \mathcal{J} of \mathcal{O}_X -modules naturally produces a coherent (pre)sheaf of \mathcal{O}_X -algebras $\mathrm{Sym}(\mathcal{J})$, which associates the symmetric $\mathcal{O}_X(U)$ -algebra $\mathrm{Sym}(\mathcal{J}(U))$ to each open $U \subset X$. The coherent sheaf of \mathcal{O}_X -algebras produces (EGA.II.1.3.1) a scheme $\mathrm{Spec}(\mathrm{Sym}(\mathcal{J}))$ over X .

If $X = \mathrm{Spec}(R)$, then $J_X^m = \mathrm{Spec}(\mathrm{Sym}(J^m(R)))$, and this is how we begin to return to the classical, and first-order definable, setting.

3.4.2 Definability

Let R be the coordinate ring of a classical affine variety, and $M := J^m(R)$. If $X := \mathrm{Spec}(R)$ and \mathcal{M} is the coherent sheaf associated to M , then $\mathrm{Spec}(\mathrm{Sym}(\mathcal{M})) = \mathrm{Spec}(\mathrm{Sym}(M))$ and the morphism to X is precisely the one associated to the ring morphism $R \rightarrow \mathrm{Sym}(M)$.

In 3.1, we presented a uniform procedure to obtain generators of $J^m(R)$ as an R -module from the generators of R as a k -algebra. These are also the generators of $\mathrm{Sym}(J^m(R))$ as an R -algebra; combining them with the generators of R as a k -algebra, we get the generators of $\mathrm{Sym}(J^m(R))$ as a k -algebra. Thus, given an affine embedding of $\mathrm{Spec}(R)$, our uniform procedure gives an affine embedding of $\mathrm{Spec}(\mathrm{Sym}(J^m(R))) = J_{\mathrm{Spec}(R)}^m$

together with the morphism to the original, affinely embedded $\text{Spec}(R)$. The uniformity of the procedure means that several objects are definable in the first-order theory of algebraically closed fields. For example, the fibers of J_X^m above closed points of X form a definable family of varieties, parametrized by points of X . For another example, if $\{X_i\}_{i \in I}$ is a definable family of varieties, then $\{J_{X_i}^m\}_{i \in I}$ is also a definable family.

More details: let $R := k[r_1, \dots, r_a]/\langle p_1, \dots, p_b \rangle$, $M := (m_1 R \oplus \dots \oplus m_c R)/\langle l_1, \dots, l_d \rangle$, where

- the r_i 's are the generators of R as a k -algebra and the p_j 's are polynomials in the variables r_i , corresponding to an embedding of X in \mathbb{A}_k^a .
- the m_k 's are the generators of M as an R -module, and the l_n 's are R -linear relations among the m_k 's.

Then $\text{Sym}(M) = k[r_1, \dots, r_a, m_1, \dots, m_c]/\langle p_1, \dots, p_b, l_1, \dots, l_d \rangle$ corresponds to an embedding of $\text{Spec}(\text{Sym}(M))$ in $\mathbb{A}^{a+c}(k)$, and the canonical morphism $\pi : \text{Spec}(\text{Sym}(M)) \rightarrow X$ is induced by the projection onto the first a coordinates.

From this it is immediately evident that for a point $p \in X$ given by $r_i \mapsto a_i \in k$, the fiber in $\text{Spec}(\text{Sym}(M))$ above p is a subset of (an isomorphic copy of) k^c defined by k -linear l_n 's: fibers of π are finite-dimensional k -vector spaces. We cite a more precise result (Lemma 3.9) in the next section.

3.4.3 Jet spaces: fibers of jet bundles

Definition 22. *The fiber of J_X^m above a point $p \in X$ is denoted $J_{X,p}^m$ and called the m th jet space of X at p .*

One purpose of the construction is that $\text{Spec}(\text{Sym}(\mathcal{J}))$ remembers the linear structure of \mathcal{J} . In particular, each fiber of $\text{Spec}(\text{Sym}(\mathcal{J}))$ over a point of X has a linear structure on it. In the special case we're interested in, this linear structure can be described more precisely:

Proposition 3.9. *Suppose that R is a finitely-generated k -algebra, $X = \text{Spec}(R)$ is an affine scheme p is a closed k -point of X . For example, X could be a classical variety and R its coordinate ring.*

Then the fiber of $\text{Spec}(\text{Sym}(\mathcal{J}))$ over p (i.e. the k -rational points of $\text{Spec}(\text{Sym}(\mathcal{J}))$ over p) has the structure of a k -vector-space, canonically identified with the dual k -vector-space of $\mathcal{J}_{X,p} \otimes_{\mathcal{O}_{X,p}} k$.

Proof. EGA.II.1.7.10 □

We work towards showing that, if $M \triangleleft R$ is the maximal ideal of the point $p \in \text{Spec}(R)$, then the fiber of $\text{Spec}(\text{Sym}(\mathcal{J}_X^m))$ over p is canonically identified with the dual of M/M^{m+1} as a k -vector-space. In other words,

Proposition 3.10. *The k -vector-spaces $\mathcal{J}_{X,p}^m \otimes_{\mathcal{O}_{X,p}} k$ and M/M^{m+1} are canonically isomorphic.*

The following commutative algebra lemma is lifted out of [MS], where it is proved in much greater generality.

Notation/Assumption 3.4.1. *Let k be a field, B a k -algebra, $e : B \rightarrow k$ a morphism of k -algebras, $N := \ker(e)$. Note that $B = N + k$. Let $T(B)$, $I(B)$, etc. be the B -modules defined above.*

Lemma 3.11. *Let $q : T(B) \rightarrow B : y \otimes z \mapsto e(y) \cdot z$.*

Then $\ker(q) = N \cdot T(B)$ and $q^{-1}(N) = I(B) + N \cdot T(B)$

The following concise proof, provided by Bjorn Poonen, is here for readers who are more comfortable with commutative algebra than the writer:

Proof. • Note that

$$0 \rightarrow N \rightarrow B \xrightarrow{e} k \rightarrow 0$$

is an exact sequence, and B is a free k -module. Tensoring with a free module is an exact functor, so

$$0 \rightarrow N \otimes B \rightarrow B \otimes B \xrightarrow{q} B \rightarrow 0$$

is an exact sequence. $N \otimes B = N \cdot T(B)$.

• The following commutes:

$$\begin{array}{ccc} B \otimes B & \xrightarrow{\quad} & B \\ \downarrow \delta & & \downarrow e \\ B & \xrightarrow{\quad} & k \end{array}$$

The two allegedly equal things are the kernel of the diagonal: $q^{-1}(N)$ is clearly the kernel of $e \circ q$; $I(B) + N \cdot T(B) = \ker(e \circ \delta)$ because $I(B)$ is the kernel of δ , and δ takes $N \cdot T(B)$ onto $N = \ker(e)$.

□

Another proof is provided for people who prefer ugly computation to exact functors:

Proof. • $N \cdot T(B) = \ker(q)$:

Let $x \in N \cdot T(B)$, meaning that for some $m \in N$, $y_i, z_i \in B$,

$$x = m \cdot \sum_i y_i \otimes z_i = \sum_i (my_i) \otimes z_i.$$

$$\text{Then } q(x) = \sum_i q((my_i) \otimes z_i) = \sum_i e(my_i) \cdot z_i = 0.$$

$$\text{Conversely, suppose that } q(\sum_i y_i \otimes z_i) = \sum_i e(y_i) \cdot z_i = 0.$$

Then $\sum_i e(y_i) \otimes z_i = 0$, so

$$\sum_i y_i \otimes z_i = \sum_i y_i \otimes z_i - \sum_i e(y_i) \otimes z_i = \sum_i (y_i - e(y_i)) \otimes z_i \in N \cdot T(B)$$

since each $y_i - e(y_i) \in N$.

• $q^{-1}(N) = \ker(e \circ q) = I(B) + N \cdot T(B)$:

Note that $y \otimes z =$

$$\begin{aligned} & \left(e(y) \otimes (z - e(z)) - (z - e(z)) \otimes e(y) \right) + \\ & \left((y - e(y)) \otimes z + (z - e(z)) \otimes e(y) \right) + \\ & \left(e(y) \otimes e(z) \right) \end{aligned}$$

So $T(B) = I(B) + N \cdot T(B) + k$, so $I(B) + N \cdot T(B)$ is a maximal ideal of $T(B)$. So it is sufficient to show that $I(B) + N \cdot T(B) \subset \ker(e \circ q)$. It is sufficient to show that $I(B) \subset \ker(e \circ q)$ and $N \cdot T(B) \subset \ker(e \circ q)$. Clearly, $N \cdot T(B) = \ker(q) \subset \ker(e \circ q)$.

To show that $I(B) \subset \ker(e \circ q)$, suppose $x = \sum_i y_i \otimes z_i \in I(B)$, meaning that $\sum_i y_i \cdot z_i = 0$.

Then

$$e(q(x)) = e\left(\sum_i e(y_i) \cdot z_i\right) = \sum_i e(y_i) \cdot e(z_i) = e\left(\sum_i y_i \cdot z_i\right) = e(0) = 0$$

so $q(x) \in \ker(e) = N$. Hence, $I(B) \subset \ker(e \circ q)$ as wanted. \square

The purpose of the previous lemma is to show that q induces an isomorphism $T(B)/(N \cdot T(B)) \cong B$. We want to think about $T(B)/(N \cdot T(B))$ as $T(B) \otimes_B k$, where k is viewed as a B -module via e . Then similarly, $(I(B) + N \cdot T(B))/N \cdot T(B)$ can be seen as $I(B) \otimes_B k$. Translating, we get

Lemma 3.12. *q induces an isomorphism*

$w : T(B) \otimes_B k \xrightarrow{\sim} B$, with $w(I(B) \otimes_B k) = N$. w in turn induces an isomorphism $(I(B) \otimes_B k)/((I(B))^{m+1} \otimes_B k) \cong N/N^{m+1}$.

Proof. The first claim is a restatement of the previous lemma.

For the second claim, note that $(I(B) \otimes_B k)^{m+1} = (I(B)^{m+1}) \otimes_B k$; and $w\left((I(B) \otimes_B k)^{m+1}\right) = N^{m+1}$ since $w(I(B) \otimes_B k) = N$ and w is an isomorphism. \square

Now we are ready to prove proposition 3.10. We use lemma 3.12 with $B = \mathcal{O}_{X,p}$, $e : \mathcal{O}_{X,p} \rightarrow k$ the evaluation homomorphism, so that $N = \mathcal{M}_{X,p}$. The lemma now says

$$(I(\mathcal{O}_{X,p}) \otimes_{\mathcal{O}_{X,p}} k)/((I(\mathcal{O}_{X,p}))^{m+1} \otimes_{\mathcal{O}_{X,p}} k) \cong (\mathcal{M}_{X,p})/(\mathcal{M}_{X,p})^{m+1} \quad (3.1)$$

Tensor products commute with quotients, so the left-hand side is canonically isomorphic to $\left((I(\mathcal{O}_{X,p})/(I(\mathcal{O}_{X,p}))^{m+1}) \otimes_{\mathcal{O}_{X,p}} k\right) = J^m(\mathcal{O}_{X,p}) \otimes_{\mathcal{O}_{X,p}} k$. The functor J^m commutes with localizations (section 3.3.1), so $J^m(\mathcal{O}_{X,p}) = J^m(R_p) = (J^m(R))_p = \mathcal{J}_{X,p}^m$, the last equality holding because \mathcal{J}_X^m is coherent.

On the right-hand side of equation 3.1, $(\mathcal{M}_{X,p})/(\mathcal{M}_{X,p})^{m+1}$ is canonically isomorphic to M/M^{m+1} .

With these substitutions, equation 3.1 now reads

$$\mathcal{J}_{X,p}^m \otimes_{\mathcal{O}_{X,p}} k \cong M/M^{m+1}$$

finishing the proof of proposition 3.10.

3.4.4 Linear maps induced on jet spaces

Remember, the construction of jet spaces is functorial, so a morphism of schemes $f : X \rightarrow Y$ induces a morphism of schemes $J^m(f) : J_X^m \rightarrow J_Y^m$ over f . For points p and $q := f(p)$, $J^m(f)$ restricts to a map $J^m(f)_p$ from the fiber of J_X^m above p to the fiber of $J^m(Y)$ above q .

Proposition 3.13. *Let $f : X \rightarrow Y$ be a morphism of schemes, $p \in X$ and $q \in Y$ be points such that $f(p) = q$; then the morphism $J^m(f)_p$ respects the k -vector-space structure described in Proposition 3.9.*

Proof. This is a local question, so we may assume that the schemes are affine and the morphism of schemes $f : \text{Spec}(R') \rightarrow \text{Spec}(R)$ comes from a morphism of k -algebras $g : R \rightarrow R'$.

Then $f(p) = q$ means that $p \triangleleft R$, $q \triangleleft R'$ are prime ideals with $q = g^{-1}(p)$.

Then g induces a k -linear map $q/q^{m+1} \rightarrow p/p^{m+1}$, whose dual gives the desired linear map between fibers, using the characterization of fibers in proposition 3.10. □

3.5 Information encoded in jet spaces

3.5.1 Constant morphisms crush jet spaces

First, an entirely obvious observation:

Lemma 3.14. *If k is the base field, then $J^m_{\text{Spec}(k)}$ is a single (closed) point, i.e. it has exactly one fiber which is zero-dimensional.*

Proof. $T(k) = k$, $I(k) = I^m(k) = \{0\}$, so $J^m(k) = \{0\}$, and $\text{Sym}_k(J^m(k)) = k$. □

Proposition 3.15. *If f is a constant morphism, then $J^m(f)_p$ is the zero map for all p .*

Proof. This is a local question, so it is enough to prove this for affine schemes.

A constant morphism between affine varieties

$$f : \text{Spec}(S) \rightarrow \text{Spec}(R)$$

comes from a k -algebra homomorphism that factors through the base field

$$S \xleftarrow{g} k \xleftarrow{h} R$$

.

Then the induced morphism

$$J^m(f) : J^m_{\text{Spec}(S)} \rightarrow J^m_{\text{Spec}(R)}$$

of jet spaces factors as

$$J_{\text{Spec}(S)}^m \xrightarrow{J^m(g)} J_{\text{Spec}(k)}^m \xrightarrow{J^m(h)} J_{\text{Spec}(R)}^m$$

Then for each $p \in \text{Spec}(S)$, the induced linear morphism of fibers $J^m(f)_p$ factors as $J^m(g)_p \circ J^m(h)_{g(p)}$ through the (unique, zero-dimensional) fiber of the jet space $J_{\text{Spec}(k)}^m$ from previous lemma. □

3.5.2 Relations between jet bundles

We have constructed vector bundles J_X^m over X , one for each $m \in \mathbb{N}$. These bundles are equipped, not only with a canonical projection to X , but also with canonical injections:

Lemma 3.16. • *There are canonical injections $j_X^{m,n} : J_X^m \rightarrow J_X^n$ for $m \leq n$.*

- *These restrict to injections $j_{X,p}^{m,n} : J_{X,p}^m \rightarrow J_{X,p}^n$ for points $p \in X$.*

Proof. • We build these injections locally, on affine pieces of X ; the canonicity insures compatibility on the intersections of the pieces. Suppose now that $X = \text{Spec}(R)$ is an affine variety. In diagram 3.1.2, the R -modules $J^m(R)$ are equipped with canonical surjective R -module homomorphisms $\pi_0^{n,m} : J^n(R) \rightarrow J^m(R)$ for $n \geq m$. These induce canonical surjective R -algebra homomorphisms $\pi^{n,m} : \text{Sym}_R(J^n(R)) \rightarrow \text{Sym}_R(J^m(R))$, which in turn induce canonical injections $j_{\text{Spec}(R)}^{m,n} : J_{\text{Spec}(R)}^m \rightarrow J_{\text{Spec}(R)}^n$.

- Since $\pi^{n,m}$ is not only a k -algebra homomorphism, but also an R -algebra homomorphism, $j_{\text{Spec}(R)}^{m,n}$ is a morphism of schemes over $\text{Spec}(R)$. □

We will later need the following technical lemma:

Lemma 3.17. *Suppose $f, g : X \rightarrow Y$ are morphisms of schemes, $f(p) = g(p) = q$ for some points $p \in X$ and $q \in Y$, and $J^n(f)_p = J^n(g)_p$. Then $J^m(f)_p = J^m(g)_p$ for all $m \leq n$.*

Proof. This diagram commutes

$$\begin{array}{ccc} J_{X,p}^n & \xrightarrow{J^n(f)_p} & J_{Y,q}^n \\ \uparrow j_{X,p}^{m,n} & & \uparrow j_{Y,q}^{m,n} \\ J_{X,p}^m & \xrightarrow{J^m(f)_p} & J_{Y,q}^m \end{array}$$

and $j_{X,p}^{m,n}$ is an injection, so $J^n(f)_p$ completely determines $J^m(f)_p$. \square

We will need a special case:

Proposition 3.18. *If $f : X \rightarrow X$ fixes a point p and $J^n(f)_p$ is the identity isomorphism of $J_{X,p}^n$, then $J^m(f)_p = \text{id}$ for all $m \leq n$.*

Proof. Set $g = \text{id}_X$, $q = p$ in the previous lemma. \square

3.5.3 A morphism is determined by its jets at a point

We work towards proving:

Proposition 3.19. *Suppose that X is an irreducible scheme, and that two morphisms from X to Y agree at a point p , and agree on all jet spaces at p . Then the two morphisms agree almost everywhere on X .*

Since this is a local question, we may assume that X is affine; let $X = \text{Spec}(R)$. Using Proposition 3.10, we view the m th jet space of X above a closed point $p \in X$ as the dual of the finite-dimensional k -vector space $M_{X,p}/M_{X,p}^{m+1}$. This immediately implies that this k -vector space is independent of the affine embedding of X , i.e. that the vector spaces resulting from different embeddings are canonically isomorphic. therefore we may and do we assume that p is the origin of the affine space, so if r_i are the corresponding affine coordinates (generators of R as a k -algebra), $M_{X,p}$ is generated by $\{r_i\}$. Having established this, we will return to viewing $J_{X,p}^m$ as the dual of the k -vector space p/p^{m+1} , thinking of p as a maximal ideal of R .

Therefore, it suffices to prove

Proposition 3.20. • *Let R be a k -algebra generated by $\{r_i\}$;*

- *Let S be a k -algebra generated by $\{s_j\}$;*
- *Let $X := \text{Spec}(R)$ and $Y := \text{Spec}(S)$ be affinely embedded, with the embeddings given by the k -algebra generators above, and suppose that the origin of the corresponding affine space belongs to each;*
- *Let R and S be intergral domains, so that X and Y are classical, irreducible varieties;*
- *Let $f, g : S \rightarrow R$ be k -algebra homomorphisms;*

- Let $\tilde{f}, \tilde{g} : X \rightarrow Y$ be the morphisms of schemes induced by f and g ;
- Suppose $\tilde{f}(0) = \tilde{g}(0) = 0$, where 0 is the origin of an affine space;
- Let M be the ideal of 0 in R , N be the ideal of 0 in S ;
- Let $f_p^m : N/N^{m+1} \rightarrow M/M^{m+1}$ be induced by f , and similarly for g_p^m ;
- Let \hat{f}_p^m (resp., \hat{g}_p^m) be the dual of f_p^m (resp., g_p^m);
- Suppose that $\hat{f}_0^m = \hat{g}_0^m$ for all m .

Then $f = g$ and therefore $\tilde{f} = \tilde{g}$.

Proof. First note that two linear maps $f_0^m, g_0^m : (N/N^{m+1}) \rightarrow (M/M^{m+1})$ induce the same dual maps if and only if they are themselves the same map. (This is just linear algebra, unrelated to the origin of the linear maps of the vector spaces.) So $f_0^m = g_0^m$ for all m .

Now consider $s \in N$ and consider $d(s) = f(s) - g(s) \in M$. We know that for all m , $f(s) \equiv g(s) \pmod{(M^{m+1})}$, so $d(s) \in \bigcap_m M^m$. It is a standard result of commutative algebra (Corollary 10.18 of [AM69]) that this intersection is trivial, as long as R is a Noetherian integral domain and M is a proper ideal. Therefore, $d(s) = 0$ for all s in N , so f and g agree on N .

Also note that R/M and S/N are both canonically isomorphic to k , and both f and g respect that isomorphism, so we have:

$$\begin{array}{ccccc} N & \longrightarrow & S & \longrightarrow & k \\ f=g \downarrow & & f \stackrel{?}{=} g \downarrow & & \text{id} \downarrow \\ M & \longrightarrow & R & \longrightarrow & k \end{array}$$

This is, among other things, a diagram of k -vector spaces. So S is canonically isomorphic to a direct sum of k -vector spaces N and k . So indeed f and g are the same map on all of S .

□

3.5.4 An étale morphism induces an isomorphism on jet spaces.

Proposition 3.21. *Suppose that $f : X \rightarrow Y$ is étale at a point p . Then the induced linear map between jet spaces, $J^m(f)_p$, is an isomorphism.*

Proof. Since f is étale at p , the induced map

$$(\mathcal{M}_{Y,f(p)})/(\mathcal{M}_{Y,f(p)})^{m+1} \longrightarrow (\mathcal{M}_{X,p})/(\mathcal{M}_{X,p})^{m+1}$$

is an isomorphism, and its dual is the isomorphism between fibers of jet spaces. \square

3.6 Commutative algebra lemmata

This section is a collection of easy results from commutative algebra that have been used in this chapter. Their proofs can be found in any introductory commutative algebra book such as [AM69], and have nothing to do with the particulars of the construction of jet spaces.

Lemma 3.22. *Let A, B be rings; let $f : A \rightarrow B$ be a ring homomorphism; let $I \leq A$, $J \leq B$ be ideals.*

Then there is a morphism $\alpha : A/I \rightarrow B/J$ making the right square in the following diagram commute iff $f(I) \subset J$, i.e. iff there is an A -module morphism $\beta : I \rightarrow J$ making the left square of the diagram commute.

$$\begin{array}{ccccc} I & \xrightarrow{\iota_A} & A & \xrightarrow{\pi_A} & A/I \\ \beta \downarrow & & f \downarrow & & \alpha \downarrow \\ J & \xrightarrow{\iota_B} & B & \xrightarrow{\pi_B} & B/J \end{array}$$

Proof. Suppose β exists; then $I \subset \ker(\pi_B \circ f)$: if $a \in I$, then $f(a) = \beta(a) \in J = \ker(\pi_B)$, so $\pi_B(f(a)) = 0$.

Suppose α exists. If $a \in I$, then $\pi_A(a) = 0$, then $\alpha(\pi_A(a)) = 0$, then $\pi_B(f(a)) = 0$, then $f(a) \in \ker(\pi_B) = J$. \square

Lemma 3.23. *let $f : A \rightarrow B$ be a ring homomorphism, let $R \subset A$, $S \subset B$ be multiplicative subsets such that $f(R) = S$. Then there exists a canonical ring homomorphism $f^+ : R^{-1}A \rightarrow S^{-1}B$ ($a/r \mapsto f(a)/f(r)$) commuting with the canonical maps $A \rightarrow R^{-1}A$, $B \rightarrow S^{-1}B$ (for this, $f(R) \subset S$ is sufficient). And $\ker(f^+) = R^{-1}\ker(f)$.*

now it follows immediately, as a special case, that:

Corollary 3.24. *for I an ideal of A , R multiplicative subset of A , $R^{-1}A/R^{-1}I = (R/I)^{-1}(A/I)$.*

Proof. By “ R/I ”, we mean the image of R in A/I under the canonical projection from A . Use previous lemma, letting $B := A/I$, f be the canonical projection $A \rightarrow A/I$, $S = R/I$. \square

Lemma 3.25. *if S is a multiplicative subset of A , I is an ideal of A , then $(S^{-1}I)^m = S^{-1}I^m$ as ideals of $S^{-1}A$.*

Chapter 4

Main story

We showed in chapter 2 that a minimal set defined by $\sigma(x) = f(x)$, where $f : \mathbb{P}^1 \rightarrow \mathbb{P}^1$ is a separable morphism, is group-like if and only if f is not an isomorphism, and f fits into the following commutative diagram of algebraic curves and non-constant morphisms

$$\begin{array}{ccccc}
 G & \xleftarrow{\psi} & \Gamma & \xrightarrow{\phi} & G^\sigma \\
 \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
 C & \longleftarrow & D & \longrightarrow & C^\sigma \\
 \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
 \mathbb{P}^1 & \xleftarrow{id} & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array} \tag{4.1}$$

where G and $\Gamma \leq G \times G^\sigma$ are algebraic groups. This diagram is obtained in section 2.4 and a little more is known about it.

The first two parts of this chapter are devoted to proving that f fits into that diagram if and only if f fits into a much simpler diagram, possibly with different G , ϕ , and π :

$$\begin{array}{ccc}
 G & \xrightarrow{\phi} & G^\sigma \\
 \downarrow \pi & & \downarrow \pi^\sigma \\
 \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array} \tag{4.2}$$

The third part of this chapter shows that the only way for f to fit into this diagram is for π to be a quotient by a finite group of automorphisms of the algebraic curve G .

The last part puts a bound on the degree of the cover π .

4.1 Part 1, in which the group correspondence and the intermediate correspondence are replaced by morphisms.

4.1.1 Summary

The purpose of this part is to prove that:

Theorem 7. *A separable morphism f fits into the first of the following diagrams (for some algebraic group curves G and Γ , algebraic group morphisms ϕ and ψ , curves C and D , etc.) if and only if it fits into the second of the following diagrams, for some other algebraic group G_2 , algebraic group morphism ϕ_2 , curve C_1 , and morphisms of curves π_1 , ρ_2 , and α_1 .*

$$\begin{array}{ccccc}
 G & \xleftarrow{\psi} & \Gamma & \xrightarrow{\phi} & G^\sigma \\
 \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
 C & \xleftarrow{\beta} & D & \xrightarrow{\alpha} & C^\sigma \\
 \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
 \mathbb{P}^1 & \xleftarrow{id} & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \\
 & & & & \\
 & & G_2 & \xrightarrow{\phi_2} & G_2^\sigma \\
 & & \uparrow \rho_2 & & \uparrow \rho_2^\sigma \\
 & & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \\
 & & \downarrow \pi_1 & & \downarrow \pi_1^\sigma \\
 & & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array} \tag{4.3}$$

First, in section 4.1.2, we find another diagram shaped like the first one, satisfying an extra assumption (\star) . Then, in section 4.1.3, we show that, under this assumption, β_1 is an isomorphism, so f must fit into

$$\begin{array}{ccccc}
 G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\
 \uparrow \rho_1 & & \uparrow & & \uparrow \rho_1^\sigma \\
 C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \\
 \downarrow \pi_1 & & \downarrow & & \downarrow \pi_1^\sigma \\
 \mathbb{P}^1 & \xrightarrow{id} & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array} \tag{4.4}$$

Finally in section 4.1.4, we examine the top of this new diagram

$$\begin{array}{ccccc}
G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\
\uparrow \rho_1 & & \uparrow & & \uparrow \rho_1^\sigma \\
C & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C^\sigma
\end{array}$$

and find G_2 , Γ_2 , ϕ_2 , and ρ_2 such that the following commutes:

$$\begin{array}{ccccc}
G_2 & \xrightarrow{id} & G_2 & \xrightarrow{\phi_2} & G_2^\sigma \\
\uparrow \rho_2 & & \uparrow & & \uparrow \rho_2^\sigma \\
C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma
\end{array}$$

Putting this back on top of the bottom half of diagram 4.4, we get

$$\begin{array}{ccccc}
G_2 & \xrightarrow{id} & G_2 & \xrightarrow{\phi_2} & G_2^\sigma \\
\uparrow \rho_2 & & \uparrow & & \uparrow \rho_2^\sigma \\
C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \\
\downarrow \pi_1 & & \downarrow & & \downarrow \pi_1^\sigma \\
\mathbb{P}^1 & \xrightarrow{id} & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

Forgetting the left half of it, we finish the proof of the theorem.

4.1.2 The first diagram chase, to get (\star)

This part is done in unnecessary generality, in hopes that it will be useful not only for the special case examined in this thesis (minimal sets encoded by morphisms on curves), but also for analyzing all σ -degree 1 minimal sets, which are encoded by correspondences on curves.

The diagrams provided are sadly monochromatic. The reader is urged to grab colored writing implements and draw a multicolored diagram as it is constructed, labeling everything. I don't see how one can follow the construction otherwise.

The notation will be moderately functorial; colors refer to a picture the reader should be drawing. In the diagrams provided, previously existing arrows are dashed, while the newly defined ones are solid. When the piece of the new diagram corresponding to x in the old diagram appears, it is labeled x_1 .

Starting diagram (black)

Unnecessary generality: in this thesis, $g = \text{id}$, but we are not making this assumption until the next section 4.1.3.

$$\begin{array}{ccccc}
 G & \xleftarrow{\psi} & \Gamma & \xrightarrow{\phi} & G^\sigma \\
 \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
 C & \xleftarrow{\beta} & D & \xrightarrow{\alpha} & C^\sigma \\
 \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
 A & \xleftarrow{g} & B & \xrightarrow{f} & A^\sigma
 \end{array} \tag{4.5}$$

G and Γ are (the projective closures of) algebraic groups; ϕ and ψ are morphisms of algebraic group curves. This commutative diagram of (smooth, irreducible, projective) algebraic curves was created in section 2.4 by normalizing another diagram. In that other diagram, $\psi \boxtimes \phi$, $g \boxtimes f$, $\rho \boxtimes \pi$, and $\beta \boxtimes \alpha$ were inclusions. In this diagram, they are still generically injective.

The purpose of this section is to produce another such diagram, with the same A , B , f , and g , which satisfies an additional assumption

$$(\star) \quad (\alpha_1) \boxtimes (\pi_1 \circ \beta_1) \text{ is generically injective.}$$

Reminder: “generically injective” is defined in section 1.4.3.

Once the diagram satisfies (\star) , a quick count of the degrees of the morphisms gives the desired result for the special case in this thesis: if g is an isomorphism, then so is β_1 .

This is an inductive argument: we assume that (\star) fails and produce a new diagram of the same kind with a lower-degree π . Since the degree of π can only decrease a finite number of times, after a finite number of iterations of this construction we will have a diagram satisfying (\star) .

D', η, α', π' (green)


Suppose the given diagram does not satisfy (\star) . Consider the morphism of varieties

$$(\alpha \boxtimes (\pi \circ \beta)) : D \rightarrow C^\sigma \times A$$

and let D' be the normalization of the normalization of its image. Let $\eta : D \rightarrow D'$ be the normalization of $(\alpha \boxtimes (\pi \circ \beta))$. Since η is an initial compositional factor of both α and $(\pi \circ \beta)$, there exist α' and π' such that


$$\alpha = \alpha' \circ \eta \text{ and } \pi \circ \beta = \pi' \circ \eta$$

Now we have:



The new group G_1 and some group morphisms (blue)

Two group homomorphisms ϕ and ψ from the same abelian group can be completed to a diamond, i.e. we can find group homomorphisms γ and δ such that $\gamma \circ \psi = \delta \circ \phi$ by letting $\ker(\gamma) = \psi(\ker(\phi))$ and $\ker(\delta) = \phi(\ker(\psi))$. For algebraic abelian groups, this determines γ and δ up to composition with the Frobenius automorphism; compose one of them with the right power of the Frobenius to match inseparable degrees. It is worth noting that, unlike most of our diagram chases, this “matching inseparable degrees” would be problematic for higher-dimensional varieties in positive characteristic. Let $\theta: \Gamma \rightarrow G_1$ be the diagonal of this diagram ($= \gamma \circ \psi = \delta \circ \phi$). Let $\zeta: D \rightarrow G_1$ be the composition of θ with the unnamed arrow $D \rightarrow \Gamma$.



λ , C_1 , $\hat{\pi}$, π_1 , and ρ_1 (red)

Now ζ factors through both β and η , and so does $\pi \circ \beta$. In other words, ζ and $\pi \circ \beta$ share non-trivial initial compositional factors β and η . Let C_1 be the normalization of the image of $(\zeta \boxtimes (\pi \circ \beta)) : D \rightarrow (G_1 \times A)$; let $\lambda : D \rightarrow C_1$ be the normalization of $(\zeta \boxtimes (\pi \circ \beta))$. From Lemma 1.17, we know that

- λ factors through both β and η , and
- ζ and $\pi \circ \beta$ both factor through λ .


Together, “ $\pi \circ \beta$ factors through λ ” and “ λ factors through β ” mean that there are morphisms $\hat{\pi}$ and π_1 such that $\pi = \pi_1 \circ \hat{\pi}$ and $\lambda = \hat{\pi} \circ \beta$.

For future use, let us note that

Lemma 4.1. *$\hat{\pi}$ is non-trivial.*

Proof. One of the initial assumptions was that β and α have no non-trivial common factors; η is a non-trivial factor of α ; hence η is not a factor of β ; however, η is a factor of λ ; hence λ has strictly higher degree than β ; hence $\hat{\pi}$ is non-trivial.

□



Done! (purple)

Add C_1^σ , π_1^σ , $\dot{\pi}^\sigma$, G_1^σ , γ^σ and ρ_1^σ to the diagram.

Let

D_1 be the normalization of $(\dot{\pi} \times \dot{\pi}^\sigma)(D) \subset C_1 \times C_1^\sigma$

Γ_1 be the normalization of $(\rho_1 \times \rho_1^\sigma)(D_1) \subset G_1 \times G_1^\sigma$

Lemma 4.2. Γ_1 is a subgroup of $G_1 \times G_1^\sigma$.

Proof. Previously established identities $\rho_1 \circ \lambda = \zeta$, $\zeta = \gamma \circ \rho \circ \beta$, and $\lambda = \dot{\pi} \circ \beta$, and the fact that β is surjective imply that

$$\rho_1 \circ \dot{\pi} = \gamma \circ \rho$$

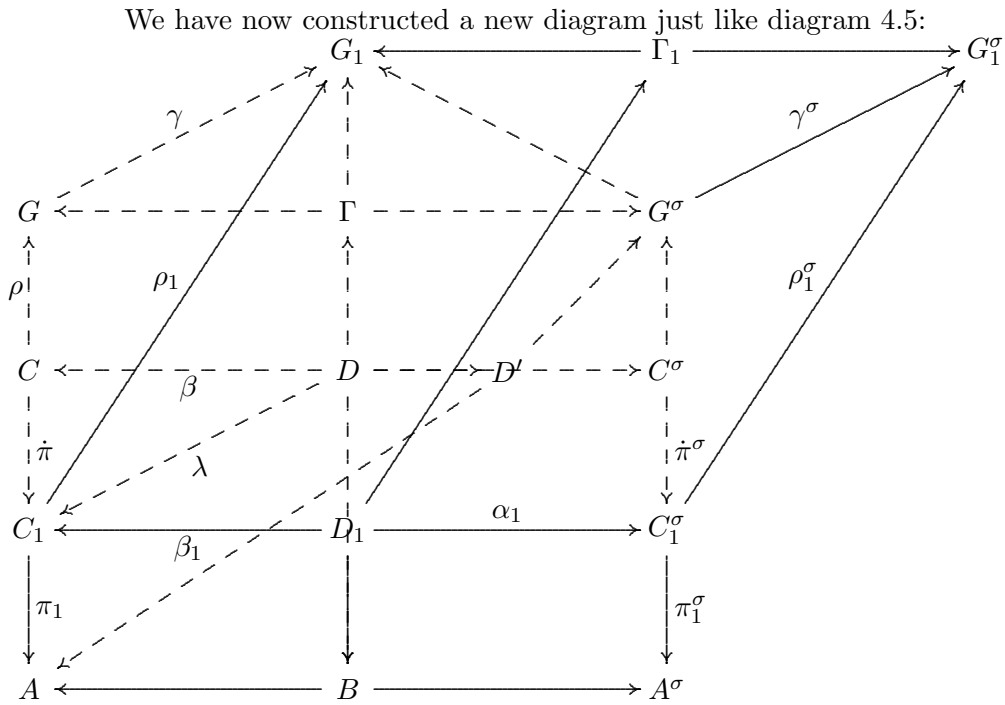
Hence,

$$\Gamma_1 = ((\rho_1 \circ \dot{\pi}) \times (\rho_1 \circ \dot{\pi})^\sigma)(D) = ((\gamma \circ \rho) \times (\gamma \circ \rho)^\sigma)(D) = (\gamma \times \gamma^\sigma)(\Gamma)$$

The image of the subgroup Γ under the group homomorphism $\gamma \times \gamma^\sigma$ is a subgroup. \square

Lemma 4.3. The degree of π_1 is strictly lower than the degree of π :

Proof. we showed in Lemma 4.1 that $\dot{\pi}$ is non-trivial. \square



This construction can be repeated as long as (\star) fails, but it can only be repeated a finite number of times, because the finite $\deg(\pi)$ decreases each time. Therefore, after finitely many repetitions of this construction (\star) will be satisfied.

4.1.3 The purpose of (\star)

(\star) allows us to count degrees of morphisms.

If the separable degrees in the new diagram are as follows

$$\begin{array}{ccccc}
 G_1 & \xleftarrow{m'} & \Gamma_1 & \xrightarrow{n'} & G_1^\sigma \\
 \uparrow e & & \uparrow & & \uparrow e \\
 C_1 & \xleftarrow{m''} & D_1 & \xrightarrow{n''} & C_1^\sigma \\
 \downarrow d & & \downarrow & & \downarrow d \\
 A & \xleftarrow{m} & B & \xrightarrow{n} & A^\sigma
 \end{array}$$

then (\star) says that a generic $\pi_1\beta_1$ -fiber containing dm'' points is mapped injectively onto m π_1^σ -fibers containing d points each, so $\deg_s(\beta_1) \leq \deg_s(g)$.

Counting inseparable degrees (addressing separately the case where $\deg_i(\alpha_1) = 1$ and the case where $\deg_i(\beta_1) = 1$) shows that $\deg_i(\beta_1) = \deg_i(g)$.

In particular, if $g = \text{id}$, and therefore $\deg_s(g) = 1 = \deg_i(g)$, then $\deg_s(\beta_1) = 1 = \deg_i(\beta_1)$ as well, and β_1 is an isomorphism. Then we can replace D by C and each of the arrows arr coming out of D by $arr \circ \beta_1^{-1}$ and maintain the commutativity of the diagram.

So if the original diagram was:

$$\begin{array}{ccccc}
 G & \xleftarrow{\psi} & \Gamma & \xrightarrow{\phi} & G^\sigma \\
 \uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
 C & \xleftarrow{\beta} & D & \xrightarrow{\alpha} & C^\sigma \\
 \downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
 A & \xrightarrow{id} & A & \xrightarrow{f} & A^\sigma
 \end{array}$$

then the new one will be:

$$\begin{array}{ccccc}
G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\
\uparrow \rho_1 & & \uparrow & & \uparrow \rho_1^\sigma \\
C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \\
\downarrow \pi_1 & & \downarrow & & \downarrow \pi_1^\sigma \\
A & \xrightarrow{id} & A & \xrightarrow{f} & A^\sigma
\end{array}$$

4.1.4 The second diagram chase

The first diagram chase is now complete; instead of the original diagram 4.1

$$\begin{array}{ccccc}
G & \xleftarrow{\psi} & \Gamma & \xrightarrow{\phi} & G^\sigma \\
\uparrow \rho & & \uparrow & & \uparrow \rho^\sigma \\
C & \longleftarrow & D & \longrightarrow & C^\sigma \\
\downarrow \pi & & \downarrow & & \downarrow \pi^\sigma \\
\mathbb{P}^1 & \xrightarrow{id} & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

we can examine diagram 4.4

$$\begin{array}{ccccc}
G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\
\uparrow \rho_1 & & \uparrow & & \uparrow \rho_1^\sigma \\
C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \\
\downarrow \pi_1 & & \downarrow & & \downarrow \pi_1^\sigma \\
\mathbb{P}^1 & \xrightarrow{id} & \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

Let us now examine the top half of the diagram above:

$$\begin{array}{ccccc}
G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\
\uparrow \rho_1 & & \uparrow & & \uparrow \rho_1^\sigma \\
C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma
\end{array}$$

We warp this top half until ψ_1 becomes an isomorphism. The warping is again inductive; the induction step is:

Lemma 4.4. *Unless ψ_1 is an isomorphism, we can produce a new diagram of the same shape*

$$\begin{array}{ccccc} G_2 & \xleftarrow{\psi_2} & \Gamma_2 & \xrightarrow{\phi_2} & G_2^\sigma \\ \uparrow \rho_2 & & \uparrow & & \uparrow \rho_2^\sigma \\ C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \end{array}$$

with the degree of ρ_2 strictly less than the degree of ρ_1 .

Proof. So we have

$$\begin{array}{ccccc} G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\ \uparrow \rho_1 & & r \uparrow & & \uparrow \rho_1^\sigma \\ C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \end{array}$$

where ψ_1 is not an isomorphism. Then ρ_1 factors as $\psi_1 \circ r$. Let us draw that into the diagram:

$$\begin{array}{ccccc} G_1 & \xleftarrow{\psi_1} & \Gamma_1 & \xrightarrow{\phi_1} & G_1^\sigma \\ \uparrow \psi_1 & & \uparrow & & \uparrow \psi_1^\sigma \\ \Gamma_1 & \longleftarrow & A & \longrightarrow & \Gamma_1^\sigma \\ \uparrow r & & \uparrow & & \uparrow r^\sigma \\ C_1 & \xrightarrow{id} & C_1 & \xrightarrow{\alpha_1} & C_1^\sigma \end{array}$$

where A is the normalization of the image of $r \boxtimes (r^\sigma \circ \alpha_1)$; this image is a subset of $\Gamma_1 \times \Gamma_1^\sigma$ and the normalizations of the projections give the arrows in the middle row of the diagram.

Let $B := (\psi_1 \times \psi_1^\sigma)^{-1}(\Gamma_1)$ and note that $A \subset B$. Since Γ_1 is a 1-dimensional subgroup of $G_1 \times G_1^\sigma$ and $(\psi_1 \times \psi_1^\sigma)$ is a finite group morphism, B is a 1-dimensional subgroup of $\Gamma_1 \times \Gamma_1^\sigma$. Then A must be a coset of B° , the connected component of B .

If $A = B^\circ$, we make the desired new diagram by letting $G_2 = \Gamma_1$, $\rho_2 = r$, and $\Gamma_2 = A$. Otherwise, we need to argue as in the proof of Lemma 4.5. Since we are working over a submodel of parameters, we can find some $a \in (\Gamma_1, A)^\sharp$. Let $t : \Gamma_1 \rightarrow \Gamma_1$ be the subtraction of a in the sense of the group law of Γ_1 , i.e. $t(g) = g - a$, and note that the following commutes:

$$\begin{array}{ccccc}
\Gamma_1 & \longleftarrow & A & \longrightarrow & \Gamma_1^\sigma \\
\uparrow t^{-1} & & \uparrow & & \uparrow (t^{-1})^\sigma \\
\Gamma_1 & \longleftarrow & B^\circ & \longrightarrow & \Gamma_1^\sigma \\
\uparrow t & & \uparrow & & \uparrow t^\sigma \\
\Gamma_1 & \longleftarrow & A & \longrightarrow & \Gamma_1^\sigma
\end{array}$$

Sticking this into the middle of the diagram above we get back to the case where $A = \circ$. \square

Since the degree of ρ can only decrease a finite number of times, iterating this construction eventually produces a commutative diagram where ψ_2 is an isomorphism. Then we can replace Γ_2 by G_2 and each of the arrows arr coming out of Γ_2 by $arr \circ \psi_2^{-1}$ and maintain the commutativity of the diagram:

$$\begin{array}{ccccc}
G_2 & \xrightarrow{\text{id}} & G_2 & \xrightarrow{\phi_2} & G_2^\sigma \\
\uparrow \rho_2 & & \uparrow & & \uparrow \rho_2^\sigma \\
C & \xrightarrow{\text{id}} & C & \xrightarrow{\alpha} & C^\sigma
\end{array}$$

4.1.5 After the chases

We now examine the diagram 4.3; to lighten notation, we drop subscripts and use g instead of α :

$$\begin{array}{ccc}
G & \xrightarrow{\phi} & G^\sigma \\
\uparrow \rho & & \uparrow \rho^\sigma \\
C & \xrightarrow{g} & C^\sigma \\
\downarrow \pi & & \downarrow \pi^\sigma \\
\mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

4.2 Part 2, in which the intermediate minimal set turns out to be a group.

4.2.1 Setup

Diagram chasing showed that, for a separable f , $(\mathbb{P}^1, f)^\sharp$ is group-like if and only if f is not an isomorphism, and there is an algebraic group curve G , an algebraic group morphism ϕ , a curve C and some more morphisms of curves such that the following commutes:

$$\begin{array}{ccc}
 G & \xrightarrow{\phi} & G^\sigma \\
 \uparrow \rho & & \uparrow \rho^\sigma \\
 C & \xrightarrow{g} & C^\sigma \\
 \downarrow \pi & & \downarrow \pi^\sigma \\
 \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array}$$

This section examines the top half and proves that

Theorem 8. *C is a group of the same kind (additive, multiplicative, or elliptic) as G , and g is a group morphism.*

Separability

We already know that f are separable; therefore, g and ϕ are separable also. We may further assume without loss of generality that ρ is also separable: if it were not, we could write it as a composition of a separable ρ' and a power of the Frobenius, so that the diagram would become

$$\begin{array}{ccc}
 G & \xrightarrow{\phi} & G^\sigma \\
 \uparrow F^m & & \uparrow F^m \\
 G^{F^{-m}} & \xrightarrow{\phi^{F^{-m}}} & (G^{F^{-m}})^\sigma \\
 \uparrow \rho' & & \uparrow \rho'^\sigma \\
 C & \xrightarrow{g} & C^\sigma \\
 \downarrow \pi & & \downarrow \pi^\sigma \\
 \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array}$$

We prove the theorem for the diagram

$$\begin{array}{ccc}
G^{F^{-m}} & \xrightarrow{\phi^{F^{-m}}} & (G^{F^{-m}})^\sigma \\
\uparrow \rho' & & \uparrow \rho'^\sigma \\
C & \xrightarrow{g} & C^\sigma \\
\downarrow \pi & & \downarrow \pi^\sigma \\
\mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

with a separable morphism in place of ρ , and then immediately conclude that the theorem holds for the original diagram, since F^m is a field automorphism.

Notation/Assumption 4.2.1. *For the rest of Part 2 of this chapter, ρ is separable.*

A technical lemma

Lemma 4.5. *If G is an algebraic group and $f : G \rightarrow G^\sigma$ is a group morphism composed with a translation (in the sense of the group law), then there is an isomorphism of varieties $t : G \rightarrow G$ and a group homomorphism $h : G \rightarrow G^\sigma$ such that the following commutes:*

$$\begin{array}{ccc}
G & \xrightarrow{h} & G^\sigma \\
\uparrow t & & \uparrow t^\sigma \\
G & \xrightarrow{f} & G^\sigma
\end{array}$$

In particular, $(G, f)^\sharp$ is group-like iff f is not a bijection.

Proof. Since we are working over a model of parameters, we can find some $a \in (G, f)^\sharp$. Let $t(g) := g - a$. Note that $h := t^\sigma \circ f \circ t^{-1}$ is also a group homomorphism composed with a translation. Note also that

$$h(0) = (t^\sigma \circ f \circ t^{-1})(0) = f(0 + a) - a^\sigma = 0$$

the last equality holding because $a \in (G, f)^\sharp$. So h is a group homomorphism, and the diagram commutes by definition of h .

Now $(G, f)^\sharp$ and $(G, h)^\sharp$ are definably isomorphic, so one is as group-like as the other. \square

4.2.2 Elliptic curves

If G is an elliptic curve, then the genus of C is at least 1. Since C admits a separable, non-constant morphism g of degree greater than 1 to a curve C^σ of the same genus, the genus of C is at most 1. So C is a genus 1 curve.

Lemma 4.6. *If the genus of C is 1 and $g : C \rightarrow C^\sigma$ is not a bijection, then $(C, g)^\sharp$ is group-like.*

Proof. Any morphism of elliptic curves is a composition of an isogeny and a translation, so Lemma 4.5 applies. \square

Theorem 8 is now proved for elliptic curves.

4.2.3 \mathbb{G}_a

Reminder: we are studying the diagram

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \uparrow \rho & & \uparrow \rho^\sigma \\ C & \xrightarrow{g} & C^\sigma \end{array}$$

G is the (projective closure of the) additive group, we are in positive characteristic p , and ϕ is given by a (separable) additive polynomial of degree $p^n > 1$. On the affine open subset where G has a group law, fibers of ϕ are cosets of its kernel, all of size p^n . We may and do identify G with \mathbb{P}^1 in the usual way, so that the affine open on which it has a group law correspond to \mathbb{A}^1 . Then $\phi^{-1}(\infty) = \{\infty\}$: ϕ ramifies at infinity with ramification index p^n , and does not ramify anywhere else.

Lemma 4.7. *The genus of C is 0 and g ramifies at the unique point in $\rho^{-1}(\infty)$ with ramification index p^n .*

Proof. Let $\{a_0, a_1, \dots, a_m\} := \rho^{-1}(\infty)$.

Note that g must be a bijection from $\{a_0, a_1, \dots, a_m\}$ to $\{a_0^\sigma, a_1^\sigma, \dots, a_m^\sigma\}$, since g must take the ρ -fiber above ∞ to the ρ^σ -fiber above $\phi(\infty) = \infty$, and cannot take any other points into that fiber since ϕ does not take any other points to ∞ . Let τ be the permutation of $\{0, 1, \dots, m\}$ such that $g(a_i) = a_{\tau(i)}^\sigma$.

Let us compare the two ways to compute the ramification index of the diagonal of the diagram at a_i :

$$e_\rho(a_i) \cdot e_\phi(\rho(a_i)) = e_g(a_i) \cdot e_{\rho^\sigma}(g(a_i))$$

Since $\rho(a_i) = \infty$ and $e_\phi(\infty) = p^n$; $g(a_i) = a_{\tau(i)}^\sigma$ and $e_{\rho^\sigma}(a_{\tau(i)}^\sigma) = e_\rho(a_{\tau(i)})$; and $e_g(\text{any point}) \leq \deg(g) = p^n$, the equation becomes

$$e_\rho(a_i) \cdot p^n = e_g(a_i) \cdot e_\rho(a_{\tau(i)}) \leq p^n \cdot e_\rho(a_{\tau(i)})$$

So for all i , $e_\rho(a_i) \leq e_\rho(a_{\tau(i)})$, with equality iff $e_g(a_i) = p^n$.

But $\sum_i e_\rho(a_i) = \sum_i e_\rho(a_{\tau(i)})$ since τ is a permutation. (This observation, suggested by Bjorn Poonen, makes the proof much cuter.)

Therefore, for all i , $e_\rho(a_i) = e_\rho(a_{\tau(i)})$ and $e_g(a_i) = p^n > 1$.

Notice that, in particular, g ramifies: C admits a ramified, separable, non-constant morphism g of degree greater than 1 to a curve C^σ of the same genus, so the genus of C must be 0.

Hurwitz's formula for a separable morphism g of degree p^n from C to C^σ , both of genus 0, is

$$2(0 - 1) \geq 2p^n(0 - 1) + \sum_a (e_g(a) - 1)$$

where equality happens only if the ramification of g is prime-to- p , which it is not. Simplifying:

$$-2 + 2p^n \geq |\rho^{-1}(\infty)|(p^n - 1)$$

which forces $|\rho^{-1}(\infty)| = 1$, as wanted. \square

Lemma 4.8. *We may identify C with \mathbb{P}^1 (as we have already done with G), in such a way that ρ and g become polynomials.*

Proof. Let $y := \rho^{-1}(\infty)$. Note that $\sigma(y) = g(y)$. Take an isomorphism $C \rightarrow \mathbb{P}^1$ that sends y to ∞ ; then the corresponding isomorphism $C^\sigma \rightarrow \mathbb{P}^1$ sends $g(y)$ to ∞ . Using this isomorphism to identify C with \mathbb{P}^1 for the rest of the discussion, we see that g totally ramifies at ∞ , i.e. it is given by a polynomial. Since $\rho^{-1}(\infty) = \{\infty\}$, ρ is given by a polynomial also. \square

So now our diagram is

$$\begin{array}{ccc} \mathbb{P}^1 & \xrightarrow{\phi} & \mathbb{P}^1 \\ \uparrow \rho & & \uparrow \rho^\sigma \\ \mathbb{P}^1 & \xrightarrow{g} & \mathbb{P}^1 \end{array}$$

where ϕ is given by a separable additive polynomial, and g and ρ are given by separable polynomials.

Lemma 4.9. *Suppose that $\rho \boxtimes g$ is generically injective. Then g is an additive group morphism.*

Actually, this is a bit of a lie: we only prove that g is a composition of an additive polynomial with a translation, and then use the proof of Lemma 4.5 to further tweak the identification of C with \mathbb{P}^1 to remove the translation.

Proof. We construct the kernel of the additive group morphism g to show that g must already be a composition of an additive polynomial with a translation.

For each $u \in \ker(\phi)$ let $t_u(x) := x + u$.

Since $u \in \ker(\phi)$, $\phi(\rho(a) + u) = \phi(\rho(a)) = \rho^\sigma(g(a))$ for every a .

Since $\rho \boxtimes g$ is generically injective, the lower left \mathbb{P}^1 is the full fiber product of the rest of the diagram: for almost any c, d such that $\rho^\sigma(c) = \phi(d)$ there exists a unique b such that $\rho(b) = d$ and $g(b) = c$ (see section 1.4.3 if this sentence is confusing).

In particular, for a generic a there exists a unique b such that $\rho(b) = \rho(a) + u$ and $g(b) = g(a)$.

Therefore, “ $t_u(\rho(a)) = \rho(b)$ and $g(a) = g(b)$ ” defines a one-to-one function $f_u : a \mapsto b$ on a Zariski-dense subvariety of the lower left \mathbb{P}^1 . Let s_u be the unique morphism on \mathbb{P}^1 that agrees with f_u almost everywhere. (See section 1.4.2 for why this exists.) One should be concerned that the definable function giving rise to s_u involves negative powers of the Frobenius. However, the fact that $g(a) = g(s_u(a))$ on generic points, and hence everywhere, precludes this. This fact also implies that $\deg(s_u) = 1$ and $s_u^{-1}(\infty) = \{\infty\}$. Therefore, $s_u(x) = m_u \cdot x + b_u$ for some $m_u, b_u \in k$.

It follows immediately from the definition of s_u that $s_u \circ s_v = s_{u+v}$ and that $s_u = s_v$ iff $u = v$. Since $\ker(\phi)$ is a group isomorphic to $(\mathbb{Z}/p\mathbb{Z})^n$ for some n , it follows that $S := \{s_u \mid u \in \ker(\phi)\}$ is the same group under composition.

Each s_u has order p , so $s_u^p(x) = m_u^p \cdot x + \text{constant term} = x$, so $m_u^p = 1$, so $m_u = 1$.

Then $b_{u+v} = b_u + b_v$, so $B := \{b_u | u \in \ker \phi\}$ is a group under addition.

Thus, fibers of g are (generically, and therefore, everywhere) orbits of S are additive cosets of B . Therefore, g is an additive group morphism, possibly composed with a translation. The translation can be removed exactly as in Lemma 4.5. □

Lemma 4.10. *Even if $\rho \boxtimes g$ is not generically injective, g is still an additive polynomial.*

Proof. If $\rho \boxtimes g$ is not generically injective, there is an α such that $\rho = \rho' \circ \alpha$ and $g = h \circ \alpha$. Let C' be the image of α and let $g' = \alpha^\sigma \circ h$ be a morphism from C' to C'^σ ; note that $(\alpha \times \alpha^\sigma)(\text{graph of } g) = \text{graph of } g'$. So the diagram above breaks into

$$\begin{array}{ccc}
 \mathbb{P}^1 & \xrightarrow{\phi} & \mathbb{P}^1 \\
 \uparrow \rho' & & \uparrow \rho'^\sigma \\
 C' & \xrightarrow{g'} & C'^\sigma \\
 \uparrow \alpha & & \uparrow \alpha^\sigma \\
 \mathbb{P}^1 & \xrightarrow{g} & \mathbb{P}^1
 \end{array}$$

By construction (taking maximal-degree α), $\rho' \boxtimes g'$ is generically injective, the previous lemma applies to them, and g' is a separable additive polynomial. Remember, $g' = \alpha^\sigma \circ h$ and $g = h \circ \alpha$. To show that g is a group morphism, it is enough to show that h and α are, which was done in Lemma 1.19 long long ago. Again, we may need Lemma 4.5 to remove a stray translation introduced by Lemma 1.19. □

Theorem 8 is now proved for additive groups.

4.2.4 \mathbb{G}_m

Now $G \cong \mathbb{P}^1$, $\phi(x) = x^n$ for some n prime to p . ϕ ramifies with index n at 0 and ∞ , both in $(G, \phi)^\sharp$. Exactly the same analysis as for the additive group shows that for every $x \in \rho^{-1}(0) \cup \rho^{-1}(\infty)$, the ramification index of g at x is n . Again, Hurwitz's formula dictates that there can be at most one $x_0 \in \rho^{-1}(0)$ and at most one $x_\infty \in \rho^{-1}(\infty)$. Again, since g ramifies, C must have genus 0.

Use an isomorphism that takes x_0 to 0, and x_∞ to ∞ to identify C with \mathbb{P}^1 . Now as a morphism from \mathbb{P}^1 to \mathbb{P}^1 , g ramifies completely at 0 and at ∞ (and nowhere else). The

only such morphisms are $x \mapsto a \cdot x^n$. Again, we can get rid of the a exactly as in the proof of Lemma 4.5 and obtain a group homomorphism of \mathbb{G}_m .

Theorem 8 is now proved for the last case, \mathbb{G}_m .

4.2.5 Conclusion and theorem

Since C itself is a group, we shall call it G and forget the original G altogether. We have now shown that $(\mathbb{P}^1, f)^\sharp$ is group-like if and only if f is not a bijection, and there is a group G , a group morphism ϕ , and a (curve) morphism π such that diagram 4.2 commutes:

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \pi & & \downarrow \pi^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

A harmless assumption

We can also assume without loss of generality that π is separable: otherwise, there is a separable morphism π' and a power m of the Frobenius automorphism such that the following commutes

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow F^m & & \downarrow F^m \\ G^{F^m} & \xrightarrow{\phi^{F^m}} & (G^{F^m})^\sigma \\ \downarrow \pi' & & \downarrow \pi'^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

Since F^m is a field automorphism, G^{F^m} is another algebraic group, of the same kind as G , and ϕ^{F^m} is a group morphism.

Another harmless assumption

Lemma 4.11. *Without loss of generality, $\pi \boxtimes \phi$ is generically injective.*

Proof. If not, both ϕ and π factor through some morphism θ : $\phi = \phi' \circ \theta$ and $\pi = \pi' \circ \theta$. Then the diagram becomes:

$$\begin{array}{ccccc}
G & \xrightarrow{\theta} & & \xrightarrow{\phi'} & G^\sigma \\
\downarrow \theta & & & & \downarrow \theta^\sigma \\
G' & \xrightarrow{\phi'} & & \xrightarrow{\theta^\sigma} & G'^\sigma \\
\downarrow \pi' & & & & \downarrow \pi'^\sigma \\
\mathbb{P}^1 & \xrightarrow{\quad} & & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

The whole top half of the diagram above is a diagram of groups by lemma 1.19, so the middle row is a group morphism, so we can examine the bottom half instead of the original diagram. Note that the degree of π' is strictly less than the degree of π , so after iterating this construction a finite number of times we obtain a diagram where $\phi \boxtimes \pi$ is generically injective. \square

Since π and π^σ have the same degree, with this extra assumption diagram 4.2 becomes a full fiber product, as described in section 1.4.3.

We have just proved:

Theorem 9. *The minimal set defined by $\sigma(x) = f(x)$ for a separable morphism $f : \mathbb{P}^1 \rightarrow \mathbb{P}^1$, is group-like if and only if f fits into diagram 4.2*

$$\begin{array}{ccc}
G & \xrightarrow{\phi} & G^\sigma \\
\downarrow \pi & & \downarrow \pi^\sigma \\
\mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
\end{array}$$

where G is a group curve, ϕ a group morphism, all morphisms separable, and $\pi \boxtimes \phi$ is generically injective so the diagram is a full fiber product.

4.3 Part 3, in which jet spaces are analyzed and the group of automorphisms of G is found.

4.3.1 Outline

Our goal in this part is to show that (without loss of generality we may assume that) π is a quotient of G by a (necessarily finite) group of algebraic-curve automorphisms of G . We are looking for automorphisms of the curve G that permute π -fibers. We want to show that one point a in a generic π -fiber can identify each of the other points b in its fiber,

and this identification is uniform across fibers. To this end, we will construct a definable function $\chi(x, y)$ on $G \otimes_{\pi} G$ (so $\chi(a, b)$ is defined only when $\pi(a) = \pi(b)$) with the following two properties:

1. Across all π -fibers, there are only finitely many possible values for $\chi(a, b)$, and each value is attained at almost all a .
2. Different b 's in the π -fiber containing a have different values of $\chi(a, b)$.

Each possible (generically occurring) value v of $\chi(a, b)$ will correspond to an automorphism of G that generically takes a point a to the unique point b in the π -fiber containing a that satisfies $\chi(a, b) = v$.

4.3.2 Jets

In chapter 3, we describe the construction of jet spaces; here, we summarize the results of that chapter. In section 3.4.1, for each $m \in \mathbb{N}$, we construct a covariant functor J^m that takes a variety X and returns a finite-dimensional vector bundle J_X^m over X , called the m th jet bundle of X . The fiber $J_{X,p}^m$ of J_X^m above a point $p \in X$ is called the m th jet space of X at p ; it is a finite-dimensional vector space over k . Notation abuse alert: the fibers of J_X^m above singular points of X may not have the same dimension as generic fibers, so perhaps it isn't fair to use the word bundle. However, we will apply this functor only to smooth varieties. $J_X^0 = X$ and J_X^1 is the usual tangent bundle of X . A morphism of varieties $f : X \rightarrow Y$ induces a morphism of varieties $J^m(f) : J_X^m \rightarrow J_Y^m$, and a k -linear function $J^m(f)_p : J_{X,p}^m \rightarrow J_{Y,f(p)}^m$ for each point $p \in X$. In this chapter, we write X_p^m instead of $J_{X,p}^m$, and f_p^m instead of $J^m(f)_p$ to lighten notation.

Properties of jet spaces proved in chapter 3 for use in this part of this chapter:

- If f is a (separable) non-constant morphism of smooth curves that does not ramify at p , then f_p^m is an isomorphism for each m (section 3.5.4).
- For any two morphisms of curves $f, g : X \rightarrow Y$, for any point $p \in X$, if $f(p) = g(p)$ and for all m , $f_p^m = g_p^m$, then $f = g$ (section 3.5.3).
- If f is a constant morphism, then all f_p^m are zero-maps (section 3.5.1).

- If $f : X \rightarrow X$ fixes a point p and f_p^n is the identity isomorphism of X_p^n , then $f_p^m = \text{id}$ for all $m \leq n$ (section 3.18).
- These functors are first-order definable in the following senses: the jet bundle of a variety is a variety; the projection from the jet bundle to the underlying variety is a morphism of varieties; the k -vector space structure on jet spaces of a variety is definable; and the function between jet bundles induced by a morphism of varieties is itself a morphism of varieties. (section 3.4.2).

This is all we need to know about jet spaces.

4.3.3 ν_{ab}^m

Let us go back to our diagram 4.2:

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \pi & & \downarrow \pi^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

In order for the following definitions and constructions to make sense, we need to avoid ramification and points of G that are not part of the group (e.g. ∞ for the multiplicative group), and also their images under functions in this diagram. For this purpose we define:

Definition 23. For $r \in \mathbb{N}$, let $F_r := f^{\sigma^r} \circ \dots \circ f^\sigma \circ f$.

Let $A(r)$ be a subvariety of G obtained by removing $F_r \circ \pi$ -fibers which contain points at which $F_r \circ \pi$ ramifies or which are not in the group of G (see section 1.4.3). (Remove whole fiber even if only one point in it is bad.)

Definition 24. (For $a, b \in A(0)$ and $c, d \in \phi(A(0))$)

- Suppose that $\pi(a) = \pi(b)$, and let $t : G \rightarrow G$ be the translation (in the sense of the group law of G) by $b - a$ that takes a to b . Then we define

$$\mu_{ab}^m := (\pi_a^m)^{-1} \circ (\pi_b^m) \circ (t_a^m)$$

μ_{ab}^m is a k -vector space automorphism of G_a^m , the m th jet space of G at a .

- Conjugating μ_{ab}^m by $(+a)_e^m$, where $(+a)$ is the translation by a and e is the identity of G , gives an automorphism of G_e^m :

$$\nu_{ab}^m := (-a)_a^m \circ \mu_{am}^m \circ (+a)_e^m$$

- The same construction for $c, d \in G^\sigma$ such that $\pi^\sigma(c) = \pi^\sigma(d)$ produces $\hat{\mu}_{cd}^m$, an automorphism of $(G^\sigma)_c^m$, and $\hat{\nu}_{cd}^m$, an automorphism of $(G^\sigma)_e^m$.

Note that (for $a \in A(0)$) $\phi_a^m : G_a^m \rightarrow (G^\sigma)_a^m$ is an isomorphism, since ϕ is separable and non-constant, and unramified because it is a group morphism.

Lemma 4.12. (For $a \in A(0)$) $\mu_{ab}^m = (\phi_a^m)^{-1} \circ \hat{\mu}_{\phi a, \phi b}^m \circ \phi_a^m$.

Proof. $\mu_{ab}^m := (\pi_a^m)^{-1} \circ (\pi_b^m) \circ (t_a^m)$, by definition of μ .

$t_a^m = (\phi_b^m)^{-1} \circ s_{\phi a}^m \circ \phi_a^m$, where s is the translation by $\phi(b) - \phi(a)$ in G^σ , because ϕ is a group morphism, so $\phi \circ t = s \circ \phi$.

$$\text{so } \mu_{ab}^m := (\pi_a^m)^{-1} \circ (\pi_b^m) \circ (\phi_b^m)^{-1} \circ s_{\phi a}^m \circ \phi_a^m.$$

$$(\pi_b^m) \circ (\phi_b^m)^{-1} = (f_{\pi a}^m)^{-1} \circ (\pi^\sigma)_{\phi b}^m \text{ because } f \circ \pi = \pi^\sigma \circ \phi$$

$$\text{so } \mu_{ab}^m := (\pi_a^m)^{-1} \circ (f_{\pi a}^m)^{-1} \circ (\pi^\sigma)_{\phi b}^m \circ s_{\phi a}^m \circ \phi_a^m.$$

$$(\pi_a^m)^{-1} \circ (f_{\pi a}^m)^{-1} = (\phi_a^m)^{-1} \circ ((\pi^\sigma)_{\phi a}^m)^{-1} \text{ because } f \circ \pi = \pi^\sigma \circ \phi.$$

$$\text{so } \mu_{ab}^m := (\phi_a^m)^{-1} \circ ((\pi^\sigma)_{\phi a}^m)^{-1} \circ (\pi^\sigma)_{\phi b}^m \circ s_{\phi a}^m \circ \phi_a^m.$$

$$\text{and finally } ((\pi^\sigma)_{\phi a}^m)^{-1} \circ (\pi^\sigma)_{\phi b}^m \circ s_{\phi a}^m = \hat{\mu}_{\phi a, \phi b}^m \text{ by the definition of } \hat{\mu}.$$

□

Corollary 4.13. (For $a \in A(0)$) $\nu_{ab}^m = (\phi_e^m)^{-1} \circ \hat{\nu}_{\phi a, \phi b}^m \circ \phi_e^m$.

Proof. this follows immediately from the previous lemma and the definitions of ν^m and $\hat{\nu}^m$, using the fact that ϕ is a group morphism, so $\phi \circ (+a) = (+\phi(a)) \circ \phi$. □

The purpose of the previous two results is

Corollary 4.14. (For $a \in A(0)$) If $\phi(a) = \phi(a')$ and $\phi(b) = \phi(b')$, then $\nu_{ab}^m = \nu_{a'b'}^m$.

Proof. Observe that in Corollary 4.13, the automorphism ν_{ab}^m depends only on $\phi(a)$ and $\phi(b)$. \square

We iterate this idea to obtain more and more pairs (a, b) with the same ν_{ab}^m .

Definition 25. For $r \in \mathbb{N}$, let $\Phi_r = \phi^{\sigma^r} \circ \dots \circ \phi^\sigma \circ \phi$.

Corollary 4.15. For any r , for $a, a', b, b' \in A(r)$, if $\Phi_r(a) = \Phi_r(a')$ and $\Phi_r(b) = \Phi_r(b')$, then $\nu_{ab}^m = \nu_{a'b'}^m$.

Proof. Induct on r :

For $r = 0$, this is Corollary 4.14.

For the induction step, assume the result for r .

Let $c = \phi(a)$, $c' = \phi(a')$, $d = \phi(b)$, and $d' = \phi(b')$.

$(\Phi_r)^\sigma(c) = \Phi_{r+1}(a) = \Phi_{r+1}(a') = (\Phi_r)^\sigma(c)$, and similarly for d and d' .

Then σ (the induction assumption at r) implies that $\hat{\nu}_{cd}^m = \hat{\nu}_{c'd'}^m$, and then Corollary 4.13 implies that $\nu_{ab}^m = \nu_{a'b'}^m$. \square

4.3.4 N_a^m

Definition 26. (For $a, b \in A(0)$ and $c, d \in \phi(A(0))$)

Let $N_a^m := \{\nu_{ab}^m \mid \pi(a) = \pi(b)\}$ and $\hat{N}_c^m := \{\hat{\nu}_{cd}^m \mid \pi^\sigma(c) = \pi^\sigma(d)\}$.

Lemma 4.16. (For $a, b \in A(0)$) $N_a^m = (\phi_e^m)^{-1} \circ \hat{N}_{\phi a}^m \circ \phi_e^m$

Proof. Corollary 4.13 says that

$$\begin{aligned} N_a^m &= \{(\phi_e^m)^{-1} \circ \hat{\nu}_{\phi a, \phi b}^m \circ \phi_e^m \mid \pi(a) = \pi(b)\} = \\ &(\phi_e^m)^{-1} \circ \{\hat{\nu}_{\phi a, \phi b}^m \mid \pi(a) = \pi(b)\} \circ \phi_e^m \end{aligned}$$

So we need to show that

$$\{\hat{\nu}_{\phi a, \phi b}^m \mid \pi(a) = \pi(b)\} = \hat{N}_{\phi a}^m$$

By definition,

$$\hat{N}_{\phi a}^m = \{\hat{\nu}_{\phi a, d}^m \mid \pi^\sigma(\phi(a)) = \pi^\sigma(d)\}$$

So we need to show that

$$\{\phi(b) \mid \pi(a) = \pi(b)\} = \{d \mid \pi^\sigma(d) = \pi^\sigma(\phi a)\}$$

This follows from the assumption (in force since Lemma 4.11) that $\pi \boxtimes \phi$ is generically injective (see section 1.4.3 for details). □

Lemma 4.17. *For any r , for $a, a', b, b' \in A(r)$, if $\Phi_r(a) = \Phi_r(a')$, then $N_a^m = N_{a'}^m$.*

Proof. Induct on r .

For $r = 0$, $\phi(a) = \phi(a')$, and the result follows immediately from Lemma 4.16.

For the induction step, assume the result for r .

Let $c = \phi(a)$, $c' = \phi(a')$.

$$(\Phi_r)^\sigma(c) = \Phi_{r+1}(a) = \Phi_{r+1}(a') = (\Phi_r)^\sigma(c').$$

Then σ (the induction assumption at r) implies that $\hat{N}_c^m = \hat{N}_{c'}^m$, and then Lemma 4.16 implies that $N_a^m = N_{a'}^m$. □

Now we can establish something resembling property (1) from the outline:

Proposition 4.18. *For each m , there is a Zariski-dense subvariety of G on which the function N^m is defined and constant.*

Proof. We know that N^m is defined on a subvariety $A(0)$; we show that it is a pure field language definable function, and that it is constant on an infinite (therefore, Zariski-dense) subset of G .

G_e^m is a finite-dimensional k vector space. The linear group $\text{Aut}(G_e^m)$ of its k -vector-space automorphisms is an algebraic group.

N^m is a (pure field language!) definable function from an irreducible curve G to the symmetric product of $\deg(\pi)$ -many copies of $\text{Aut}(G_e^m)$.

Let a be a sufficiently generic point of G , i.e. suppose $a \notin A(r)$ for any r ; then $\bigcup_r (\Phi_r\text{-fiber of } a)$ is an infinite set on which N^m is constant. (This is the last lemma.)

A definable function on a curve which is constant on an infinite, and therefore Zariski-dense, subset, must be constant almost everywhere. □

4.3.5 Another reduction, towards uniqueness

Equivalence relations

Now we know that there is a finite set of possible values for ν_{ab}^m . We would like each value to occur exactly once for each generic a , so that a together with the value of ν_{ab}^m would uniquely determine b . In other words, we wish to eliminate the possibility that $\nu_{ab}^m = \nu_{ac}^m$ for distinct b, c in the π -fiber containing a .

Lemma 4.19. *For $a, b, c \in A(0)$,*

1. *If $\pi(a) = \pi(b) = \pi(c)$, then $\nu_{ac}^m = \nu_{ab}^m \circ \nu_{bc}^m$.*
2. *(If $\pi(a) = \pi(b) = \pi(c)$,) $\nu_{ab}^m = \nu_{ac}^m$ if and only if $\nu_{bc}^m = \text{id}$.*
3. *For each m , “ $\pi(b) = \pi(c)$ and $\nu_{bc}^m = \text{id}$ ” is an equivalence relation E^m on A_0 .*
4. *For $m \leq n$, E^n refines E^m .*

Proof. (1) is clear from the definition of ν_{ab}^m . (2) and (3) follow immediately from (1). (4) follows from Proposition 3.18. \square

Since we would like at least one of these equivalence relations to be just equality, we shall quotient out by them:

Definition 27. *Let E^m be the equivalence relation defined by $\nu_{bc}^m = \text{id}$ on $A(0)$.*

Lemma 4.20. *There exists an integer M such that*

- *For each $m \geq M$, there is a Zariski-dense variety of G on which E^M and E^m are both defined and equal.*
- *There is a Zariski-dense subset on which E^M and E^m are both defined and equal, for all $m \geq M$.*

Proof. The equivalence classes of E^0 are simply π -fibers: ν_{bc}^0 is an automorphism of G_e^0 , the fiber of J_G^0 above e . Since J_G^0 is the trivial, 0-dimensional vector bundle over G , its fibers are 0-dimensional k -vector-spaces with only one automorphism: the identity.

We just proved that for $m \leq n$, E^n refines E^m .

E^0 already has finite equivalence classes (of size $\deg(\pi)$), so $\{E^m\}_m$ is an infinite chain of refining equivalence relations with finite classes. Since each E^m is definable, the

size n_m of its fibers is generically constant. The function $m \mapsto n_m$ is a non-increasing function from positive integers to positive integers, hence it is eventually constant: there exists some M so that for all $m \geq M$, the fibers of E^m generically have the same size as fibers of E^M . Since E^m refines E^M and both are definable, this implies that E^m and E^M agree on a Zariski-dense subvariety of G .

□

Quotient by E^M

Let $\zeta : G \rightarrow H$ be the separable morphism of curves whose generic fibers are equivalence classes of E^M , constructed in Lemma 1.15. Since fibers of ζ (equivalence classes of E^M) are always subsets of the fibers of π , π can be written as $\eta \circ \zeta$ for some morphism η .

Now diagram 4.2 becomes

$$\begin{array}{ccc}
 G & \xrightarrow{\phi} & G^\sigma \\
 \downarrow \zeta & & \downarrow \zeta^\sigma \\
 H & & H^\sigma \\
 \downarrow \eta & & \downarrow \eta^\sigma \\
 \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1
 \end{array} \tag{4.6}$$

ζ is a group morphism.

Definition 28. By analogy with μ_{ab}^m , for $a, b \in A(0)$ such that $\zeta(a) = \zeta(b)$, let

$$\lambda_{ab}^m := (\zeta_a^m)^{-1} \circ (\zeta_b^m) \circ (t_a^m)$$

where t is again translation by $b - a$ in the sense of the group law of G .

Lemma 4.21. For $a \in A(0)$, $\lambda_{ab}^m = \text{id}$ for all m .

Proof. By definition of ζ , if $\zeta(a) = \zeta(b)$, then for all m , $\nu_{ab}^m = \text{id}$.

If $\nu_{ab}^m = \text{id}$, then $\mu_{ab}^m = \text{id}$; and if in addition, $\zeta(a) = \zeta(b)$ so that it makes sense to speak of λ_{ab}^m , then $\lambda_{ab}^m = \text{id}$ also:

$$\begin{aligned}
 \lambda_{ab}^m &= (\zeta_a^m)^{-1} \circ (\zeta_b^m) \circ (t_a^m) = \\
 &= (\zeta_a^m)^{-1} \circ (\eta_{\zeta b}^m)^{-1} \circ (\eta_{\zeta a}^m) \circ (\zeta_b^m) \circ (t_a^m) = \nu_{ab}^m = \text{id}
 \end{aligned}$$

□

Lemma 4.22. *If $\zeta(a) = \zeta(a + u)$ for one sufficiently generic a (i.e. $a \in A(r)$ for all r), then $\zeta(a) = \zeta(a + u)$ for all a .*

Proof. Let $b = a + u$. We just showed that for all m

$$\text{id} = \lambda_{ab}^m = (\zeta_a^m)^{-1} \circ (\zeta_b^m) \circ (t_a^m)$$

where t is the translation by $b - a = u$. Applying ζ_a^m to the equation above and noting that $b = t(a)$, we see that

$$\zeta_a^m = (\zeta_{t(a)}^m) \circ (t_a^m) = (\zeta \circ t)_a^m$$

for all m . Since the linear maps induced on fibers of jet spaces determine the morphism (section 3.5.3), $\zeta = \zeta \circ t$, as wanted. \square

Take a sufficiently generic a and let $U := \{u \mid \zeta(a) = \zeta(a + u)\}$. U is a finite subgroup of G ; let $\zeta' : G \rightarrow G/U$ be the group morphism whose kernel is U .

Proposition 4.23. *$\zeta = \zeta'$, i.e. all fibers of ζ are cosets of the group U , which is independent of the choice of generic a .*

Proof. We just showed that this holds on a Zariski-dense subset, so it holds on all of G . \square

The intermediate group morphism ϕ_1

We now know that ζ in diagram 4.6 is a group morphism and H is a group curve.

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \zeta & & \downarrow \zeta^\sigma \\ H & & H^\sigma \\ \downarrow \eta & & \downarrow \eta^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

Lemma 4.24. *There is group morphism $\phi_1 : H \rightarrow H^\sigma$ that fits into the above diagram.*

Proof. If such a ϕ_1 existed, its graph would be $(\zeta \times \zeta^\sigma)$ (the graph of ϕ); and if $(\zeta \times \zeta^\sigma)$ (the graph of ϕ) is a graph of some morphism, it is necessarily a group morphism. So all we really need to show is that for all g and h , if $\zeta(g) = \zeta(h)$, then $\zeta^\sigma(\phi(g)) = \zeta^\sigma(\phi(h))$. This follows immediately from the definition of ζ in terms of ν_{ab}^M and from Corollary 4.15:

1. $\zeta(g) = \zeta(h)$, which implies
2. $\pi(g) = \pi(h)$ and $\nu_{gh}^M = \text{id}$, which implies
3. $\pi^\sigma(\phi(g)) = f(\pi(g)) = f(\pi(h)) = \pi^\sigma(\phi(h))$ (the middle equality follows from previous step, the first and last from commutativity of diagram 4.2); and $\hat{\nu}_{\phi a \phi b}^M = \text{id}$ (Corollary 4.15).
4. (3) says $\zeta^\sigma(\phi(g)) = \zeta^\sigma(\phi(h))$: this is simply the definition of ζ^σ . \square

4.3.6 Automorphisms

Now the bottom half of the diagram above

$$\begin{array}{ccc} H & \xrightarrow{\phi_1} & H^\sigma \\ \downarrow \eta & & \downarrow \eta^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

is just like the original diagram 4.2 with one new property. Let o_{ab}^m be defined as ν_{ab}^m with η in place of π .

Lemma 4.25. (*New Property*) For $m \geq M$ (the M from Lemma 4.20) and generic $a \in H$, if $b \neq a$, then $o_{ab}^m \neq \text{id}$.

Proof. $o_{ab}^m = \text{id}$ iff $(\eta_a^m)^{-1} \circ (\eta_b^m) \circ (t_a^m) = \text{id}$ where t is the translation in H by $b - a$. Take $c, d \in G$ such that $\zeta(c) = a$ and $\zeta(d) = b$. Since ζ is a group morphism, $t \circ \zeta = \zeta \circ s$, where s is the translation in G by $(d - c)$. So

$$\begin{aligned} & (\eta_a^m)^{-1} \circ (\eta_b^m) \circ (t_a^m) = \\ & (\eta_a^m)^{-1} \circ (\eta_b^m) \circ \zeta_d^m \circ s_d^m \circ (\zeta_c^m)^{-1} \end{aligned}$$

Therefore, $o_{ab}^m = \text{id}$ iff

$$(\zeta_d^m)^{-1} \circ (\eta_a^m)^{-1} \circ (\eta_b^m) \zeta_c^m \circ s_d^m = \text{id}$$

Remembering that $\pi = \eta \circ \zeta$, we see that the equation above says

$$\mu_{cd}^m = (\pi_d^m)^{-1} \circ \pi_c^m \circ s_d^m = \text{id}$$

which means that, by definition of ζ , $a = \zeta(c) = \zeta(d) = b$. \square

We now forget all about η and ζ and re-summarize what we know so far, including the new property:

Proposition 4.26. *The minimal set defined by $\sigma(x) = f(x)$ for a separable morphism $f : \mathbb{P}^1 \rightarrow \mathbb{P}^1$, is group-like if and only if f fits into diagram 4.2*

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \pi & & \downarrow \pi^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

where G is a group curve; ϕ is a group morphism; all morphisms are separable; $\pi \boxtimes \phi$ is generically injective (so the diagram is a full fiber product); and the new property holds: for all $m \geq M$ (the M from Lemma 4.20), for all generic a , and for all b such that $\pi(a) = \pi(b)$ but $a \neq b$, $\nu_{ab}^m \neq \text{id}$.

Let $\chi(a, b) := \nu_{ab}^M$; we have shown that

- There is a finite set S of possible values for $\chi(a, b)$ (Proposition 4.18).
- If a value $s \in S$ is taken on at *some* sufficiently generic point a in G (i.e. if for some sufficiently generic point a in G there exists a point $b \in G$ with $\chi(a, b) = s$), then this value is taken on at *all* sufficiently generic points of G (also Proposition 4.18).
- For sufficiently generic $a \in G$, if $b \neq c$ and $\pi(a) = \pi(b) = \pi(c)$, then $\chi(a, b) \neq \chi(a, c)$. (This is the new property)

Therefore, for each $s \in S$ and for each sufficiently generic $a \in G$, there is a unique b such that $\pi(b) = \pi(a)$ and $\chi(a, b) = s$. Therefore, the formula $\pi(y) = \pi(x) \wedge \chi(x, y) = s$ defines an injective function $\alpha_s : x \mapsto y$ on generic points of G . Counting separable degrees in $\pi = \pi \circ \alpha_s$ shows that this definable function cannot involve negative powers of the Frobenius, so it is equal to a morphism almost everywhere; we abuse notation and call that morphism α_s also. The same equation $\pi = \pi \circ \alpha_s$ shows that α_s is birational, hence extends to an automorphism α_s of G (the curve, not necessarily the group); and the π -fiber of a generic a is $\{\alpha_s(a) \mid s \in S\}$. It follows immediately that $A := \{\alpha_s\}_{s \in S}$ is a group under composition, and that the fibers of π are (generically, and therefore everywhere) orbits of A .

Theorem 10. *The minimal set defined by $\sigma(x) = f(x)$ for a separable morphism $f : \mathbb{P}^1 \rightarrow \mathbb{P}^1$ is group-like if and only if there exist*

- *a group curve G ,*
- *a group morphism ϕ ,*
- *and a finite group S of automorphisms of the curve G which does not include any translations of G ,*

such that f fits into diagram 4.2

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \pi & & \downarrow \pi^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

where π is the separable morphism whose generic fibers are orbits of S .

Proof. The only part not yet proved is the lack of G -translations in S ; assume otherwise: let t be a G translation in S . Then $\pi = \pi \circ t$, which immediately implies that $\nu_{a,t(a)}^m = \text{id}$ for almost all a and for all m , contradicting Proposition 4.26. What really happened is that all the G -translations in S have already been absorbed into ζ in section 4.3.5. \square

4.4 Part 4, in which the degree of the cover is bounded.

The purpose of this chapter is to bound the degree of π in diagram 4.2 in theorem 10

$$\begin{array}{ccc} G & \xrightarrow{\phi} & G^\sigma \\ \downarrow \pi & & \downarrow \pi^\sigma \\ \mathbb{P}^1 & \xrightarrow{f} & \mathbb{P}^1 \end{array}$$

4.4.1 Elliptic curves

Any morphism from an elliptic curve G to itself is an isogeny composed with a translation. So any automorphism of G is a group automorphism γ composed with a translation τ . If S contains two elements $\gamma \circ \tau_1$ and $\gamma \circ \tau_2$ with the same group-automorphism part, composing them will put a translation into S , which is not supposed to happen.

Therefore, S can only have as many elements as G has group automorphisms: at most 24 [Sil86].

4.4.2 \mathbb{G}_m

$G = \mathbb{P}^1$, $\phi(x) = x^m$ for some m prime to the characteristic and greater than 1. Automorphisms of \mathbb{P}^1 are fractional linear transformations $l : x \mapsto \frac{ax+b}{cx+d}$.

If $l \in S$, then for all x there is an $l' \in S^\sigma$ such that $l(x)^m = l'(x^m)$. Because there are only finitely many choices for l' and \mathbb{P}^1 is irreducible, this l' is in fact the same for all x . So $l(x)^m = l'(x^m)$ holds for all x ; expanding,

$$\frac{(ax+b)^m}{(cx+d)^m} = \frac{a'x^m + b'}{c'x^m + d'}$$

The poles of the left side of the equation are:

- if $c \neq 0$, there is a pole of order m at $-d/c$ and no other poles
- if $c = 0$, there are no poles at all.

The poles of the right hand side of the equation are:

- if $c' \neq 0$ and $d' \neq 0$, there are m distinct poles of order 1, one at each $\zeta \cdot e$ where e is one particular m th root of $-d'/c'$ and ζ is any m th root of unity.
- if $c' \neq 0$ but $d' = 0$, there is a pole of order m at 0.
- if $c' = 0$, there are no poles at all.

Since m is not 1, these cannot match unless either $c = c' = 0$ (no poles on either side), or $d = d' = 0$ (pole of order m at 0 on both sides).

The same analysis of zeros of the two sides shows that either $a = a' = 0$ or $b = b' = 0$.

So either $l(x) = (a/d)x$ is a \mathbb{G}_m -translation and cannot be in S ; or $l(x) = (b/cx)$, so $l^2(x) = (b/c)^2x$, so $(b/c) = \pm 1$. Note that if S contains both $x \mapsto 1/x$ and $x \mapsto -1/x$, then S also contains their composition, $x \mapsto -x$, a \mathbb{G}_m -translation. So S can only have size 1 or 2.

4.4.3 \mathbb{G}_a

The additive group is relevant only in positive characteristic; then group maps are additive polynomials of the form $\phi(x) = \sum_i r_i x^{p^i}$. In the same way as for the multiplicative group, the underlying curve is \mathbb{P}^1 , whose automorphisms are fractional linear transformations. Again, for every $l \in S$ there must be an $l' \in S^\sigma$ such that $\phi(l(x)) = l'(\phi(x))$:

$$\sum_i r_i \left(\frac{ax + b}{cx + d} \right)^{p^i} = \frac{a' \sum_i (r_i x^{p^i}) + b'}{c' \sum_i (r_i x^{p^i}) + d'}$$

The same analysis of poles shows that $c' \sum_i (r_i x^{p^i}) + d'$ must have a unique root, of multiplicity $\deg(\phi)$, which is impossible unless $c' = 0$, or $d' = 0$ and ϕ is a single monomial. But if ϕ is a monomial, it is purely inseparable, so $(\mathbb{P}^1, \phi)^\sharp$ is field-like, not group-like. So $c = c' = 0$, and d can be absorbed into a and b , and d' into a' and b' .

Now the equation above becomes

$$\sum_i r_i (a^{p^i} x^{p^i} + b_i^{p^i}) = a' (\sum_i r_i x^{p^i}) + b'$$

Matching coefficients for the same powers of x , we get that for all i , $r_i a^{p^i} = a' r_i$; in other words, whenever r_i and r_j are non-zero, $a^{p^i} = a^{p^j}$; note that there are at least two indices i such that $r_i \neq 0$, because otherwise ϕ would be either constant or purely inseparable. Each of the restrictions $a^{p^i} = a^{p^j}$ is a polynomial (in a) of degree at most p^m , so has at most p^m -many solutions.

As before, if there are two elements of S with the same a , then there is a translation in S , which there isn't. So if the degree of f is p^m , the degree of π is also at most p^m .

4.4.4 The End

Theorem 11. • *The minimal set $(\mathbb{P}^1, f)^\sharp$ is non-orthogonal to a minimal group living on an elliptic curve if and only if there exists an elliptic curve G , an isogeny $\phi : G \rightarrow G^\sigma$ of the same degree as f , and a map $\pi : G \rightarrow \mathbb{P}^1$ of degree at most 24 making diagram 4.2 commute.*

- *It is non-orthogonal to a minimal group living on \mathbb{G}_m if and only if $\phi : x \mapsto x^{\deg(f)}$ and π of degree at most 2 make diagram 4.2 commute.*

- *It is non-orthogonal to a minimal group living on \mathbb{G}_a if and only if the degree of f is a power of the characteristic of the field, and there is an additive polynomial ϕ of the same degree as f , and there is a map π of degree at most $\deg(f)$ making diagram 4.2 commute.*

Bibliography

- [AM69] M. F. Atiyah and I. G. Macdonald, *Introduction to commutative algebra*, Addison-Wesley Publishing Co., Reading, Mass.-London-Don Mills, Ont., 1969. MR MR0242802 (39 #4129)
- [CH99] Zoé Chatzidakis and Ehud Hrushovski, *Model theory of difference fields*, Trans. Amer. Math. Soc. **351** (1999), no. 8, 2997–3071. MR MR1652269 (2000f:03109)
- [CHP02] Zoé Chatzidakis, Ehud Hrushovski, and Ya'acov Peterzil, *Model theory of difference fields. II. Periodic ideals and the trichotomy in all characteristics*, Proc. London Math. Soc. (3) **85** (2002), no. 2, 257–311. MR MR1912052 (2004c:03047)
- [Coh65] Richard M. Cohn, *Difference algebra*, Interscience Publishers John Wiley & Sons, New York-London-Sydney, 1965. MR MR0205987 (34 #5812)
- [Gro60] A. Grothendieck, *Éléments de géométrie algébrique. I. Le langage des schémas*, Inst. Hautes Études Sci. Publ. Math. (1960), no. 4, 228. MR MR0217083 (36 #177a)
- [Gro61] ———, *Éléments de géométrie algébrique. II. Étude globale élémentaire de quelques classes de morphismes*, Inst. Hautes Études Sci. Publ. Math. (1961), no. 8, 222. MR MR0217084 (36 #177b)
- [Gro62] ———, *Techniques de construction en géométrie analytique. VII: Étude locale des morphismes; éléments de calcul infinitésimal.*, Familles d'Espaces Complexes et Fondements de la Geom. Anal., Sem. H. Cartan 13 (1960/61), No.13, 13 p.; No.14, 27 p.; No.15, 10 p.; No.16, 20 p.; No.17, 20 p. (1962)., 1962.
- [Har77] Robin Hartshorne, *Algebraic geometry*, Springer-Verlag, New York, 1977, Graduate Texts in Mathematics, No. 52. MR MR0463157 (57 #3116)

- [HI03] E. Hrushovski and M. Itai, *On model complete differential fields*, Trans. Amer. Math. Soc. **355** (2003), no. 11, 4267–4296 (electronic). MR MR1990753 (2004m:03139)
- [Hod97] Wilfrid Hodges, *A shorter model theory*, Cambridge University Press, Cambridge, 1997. MR MR1462612 (98i:03041)
- [Hru01] Ehud Hrushovski, *The Manin-Mumford conjecture and the model theory of difference fields*, Ann. Pure Appl. Logic **112** (2001), no. 1, 43–115. MR MR1854232 (2003d:03061)
- [HS99] Ehud Hrushovski and Thomas Scanlon, *Lascaz and Morley ranks differ in differentially closed fields*, J. Symbolic Logic **64** (1999), no. 3, 1280–1284. MR MR1779761 (2002g:03082)
- [Mac97] Angus Macintyre, *Generic automorphisms of fields*, Ann. Pure Appl. Logic **88** (1997), no. 2-3, 165–180, Joint AILA-KGS Model Theory Meeting (Florence, 1995). MR MR1600899 (99c:03046)
- [Moo04] Rahim Moosa, *Jet spaces in complex analytic geometry: an exposition*, 2004.
- [MS] Rahim Moosa and Thomas Scanlon, *Differential jets and arcs.*, in preparation.
- [Pil96] Anand Pillay, *Geometric stability theory*, Oxford Logic Guides, vol. 32, The Clarendon Press Oxford University Press, New York, 1996, , Oxford Science Publications. MR MR1429864 (98a:03049)
- [Sil86] Joseph H. Silverman, *The arithmetic of elliptic curves*, Graduate Texts in Mathematics, vol. 106, Springer-Verlag, New York, 1986. MR MR817210 (87g:11070)
- [vdDS84] L. van den Dries and K. Schmidt, *Bounds in the theory of polynomial rings over fields. A nonstandard approach*, Invent. Math. **76** (1984), no. 1, 77–91. MR MR739626 (85i:12016)
- [Wag00] Frank O. Wagner, *Simple theories*, Mathematics and its Applications, vol. 503, Kluwer Academic Publishers, Dordrecht, 2000. MR MR1747713 (2001b:03035)