

Contributions to the model theory of partial differential fields

by

Omar León Sánchez

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Pure Mathematics

Waterloo, Ontario, Canada, 2013

© Omar León Sánchez 2013

I hereby declare that I am the sole author of this thesis.

This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In this thesis three topics on the model theory of partial differential fields are considered: the generalized Galois theory for partial differential fields, geometric axioms for the theory of partial differentially closed fields, and the existence and properties of the model companion of the theory of partial differential fields with an automorphism. The approach taken here to these subjects is to relativize the algebro geometric notions of prolongation and D-variety to differential notions with respect to a fixed differential structure.

It is shown that every differential algebraic group which is not of maximal differential type is definably isomorphic to the sharp points of a relative D-group. Pillay's generalized finite dimensional differential Galois theory is extended to the possibly infinite dimensional partial setting. Logarithmic differential equations on relative D-groups are discussed and the associated differential Galois theory is developed. The notion of generalized strongly normal extension is naturally extended to the partial setting, and a connection between these extensions and the Galois extensions associated to logarithmic differential equations is established.

A geometric characterization, in the spirit of Pierce-Pillay, for the theory $DCF_{0,\ell+1}$ (differentially closed fields of characteristic zero with $\ell + 1$ commuting derivations) is given in terms of the differential algebraic geometry of $DCF_{0,\ell}$ using relative prolongations. It is shown that this characterization can be rephrased in terms of characteristic sets of prime differential ideals, yielding a first-order geometric axiomatization of $DCF_{0,\ell+1}$.

Using the machinery of characteristic sets of prime differential ideals it is shown that the theory of partial differential fields with an automorphism has a model companion. Some basic model theoretic properties of this theory are presented: description of its completions, supersimplicity and elimination of imaginaries. Differential-difference modules are introduced and they are used, together with jet spaces, to establish the canonical base property for finite dimensional types, and consequently the Zilber dichotomy for minimal finite dimensional types.

Acknowledgements

Throughout my years in Waterloo, I have been fortunate enough to be immersed in a welcoming, friendly and selfless environment. I am profoundly thankful to the faculty, staff, and my fellow graduate students at the Pure Mathematics Department of the University of Waterloo. Foremost, I am deeply indebted to Rahim Moosa, whose knowledge and guidance have been essential to the development of this thesis. I thank him for his patience, wise advise, generosity and frankness, which have been invaluable to me. He has not only been an outstanding supervisor, but also a great mentor and exceptional example of dedication and enthusiasm.

I am also compelled to thank the Model Theory community, which has been incredibly welcoming and supportive. In particular, I am thankful to Moshe Kamensky, Lynn Scow, Philipp Hieronymi, Lou van den Dries, Deirdre Haskell, Patrick Speissegger and Chris Miller for giving the opportunity to present my research, and for their hospitality. Special thanks go also to Anand Pillay and David Marker, whose own work, and clear exposition on both books and papers, have taught me so much, and inspired great part of the work presented here.

This thesis has also benefited from the many discussions and e-mail exchanges with James Freitag and Ronnie Nagloo. I thank James for being a familiar face in many of my talks, for many clear explanations about his ideas and for his infectious enthusiasm. I thank Ronnie for so many exciting conversations about math and future projects, and for his encouraging words.

I can certainly not overlook the generosity and support of the Differential Algebra community. I would like to thank the organizers and participants of the Kolchin seminar at CUNY and the DART V meeting. In particular, I thank Alexey Ovchinnikov, Phyllis Cassidy and William Sit for carefully listening while I talked about my research, and for having helpful advise and constructive feedback.

Finally, I have not enough words of gratitude to express my recognition to Ale, but I specially thank her for embarking us in this academic adventure. I thank her for her patience, for her love and for her friendship. I am also indebted to her for her always smart and honest advise, and for teaching me, with her own example, that one must always put both their brain and their heart in even the tiniest of tasks. I lastly thank her for sharing and enjoying with me these five years, in spite of them being filled with *derivations*, *model companions*, *characteristic sets* and, in some lucky days, *stability*.

Dedication

To Ale, her family and my family, for their continuous support.

Table of Contents

Introduction	1
1 Preliminaries	7
1.1 Model theory	7
1.2 Differential algebra	10
1.3 Differentially closed fields	13
1.4 Independence and U -rank	16
1.5 Affine differential algebraic geometry	18
1.6 Differential type and typical differential dimension	21
1.7 Abstract differential algebraic varieties	26
2 Relative prolongations, D-varieties and D-groups	30
2.1 A review of prolongations and D-varieties in DCF_0	30
2.2 Differential derivations	33
2.3 The relative prolongation of a differential algebraic variety	35
2.4 Relative D-varieties	42
2.5 Relative D-groups	47
3 Generalized Galois theory for partial differential fields	53
3.1 A review of Galois theory for ordinary differential fields	53
3.2 Generalized strongly normal extensions	57

3.3	Relative logarithmic differential equations and their Galois extensions . . .	66
3.4	Two examples	73
4	Geometric axioms for $DCF_{0,m}$	76
4.1	A review of the Pierce-Pillay axioms	76
4.2	Geometric characterizations and an axiomatization	79
5	The model theory of partial differential fields with an automorphism	85
5.1	A review of $ACFA$	86
5.2	The model companion $DCF_{0,m}A$	89
5.3	Basic model theory of $DCF_{0,m}A$	92
5.4	Zilber dichotomy for finite dimensional types	100
	Appendix: Some U-rank computations in $DCF_{0,m}$	107
	References	110

Introduction

Just as algebraic geometry can be viewed as the study of the set of solutions to systems of polynomial equations, *differential algebraic geometry* studies the set of solutions to systems of differential algebraic equations. Our approach to the subject is in the spirit of Weil, the analogue of his universal domain for algebraic geometry is provided by the model theoretic setting: our *differential algebraic varieties* are taken to live inside a large *differentially closed field* of characteristic zero. By now, it is well known that the interaction between model theory and differential algebraic geometry goes well beyond this. The theory of differentially closed fields is arguably the most interesting example of an ω -stable theory so far. For example, the study of *finite dimensional* differential algebraic varieties has led to a variety of interesting examples, and detailed descriptions, of notions from general stability (see for example [32] and [51]). On the other hand, the machinery of ω -stable theories has been successfully implemented in diverse aspects of differential algebra, such as differential Galois theory and the theory of differential algebraic groups (see [54], [45] and [50]). The model theory of finite dimensional differential algebraic varieties is behind Hrushovski's celebrated proof of the Mordell-Lang conjecture for function fields in characteristic zero [19].

While much has been done in the finite dimensional setting, it was not until the last decade that great interest in describing *infinite dimensional* differential algebraic varieties, from a model theoretic perspective, began [37]. These descriptions are mostly of interest in the partial case, i.e., in differential fields with several commuting derivations, since in this setting one has a rich supply of infinite dimensional definable sets. One example is the Heat equation $\delta_1 x = \delta_2^2 x$, which was studied in depth by Suer in [56]. This thesis can be viewed as a next step towards a better understanding of infinite dimensional differential algebraic geometry.

The basic approach we have in mind is the following: just as one thinks of ordinary (i.e., one derivation) differential algebraic geometry as adding structure to algebraic geometry (namely that induced by the derivation), we view differential algebraic geometry in two commuting derivations as adding structure to ordinary differential algebraic geometry, and

so on. Therefore, instead of treating pure fields as our base theory to which we then add some new structure, we treat differential fields as our base theory. This way of thinking of differential algebraic geometry is somewhat implicit in the work of Kolchin [25]; however, by formalizing this approach and building a theory systematically upon it, we are able to tackle several problems on the subject.

In the remainder of this introduction we will outline the contents of this thesis. More thorough introductions to these topics appear as the first sections of each chapter.

In the preliminary Chapter 1, we present the differential algebraic background needed for the rest of this thesis. We make special emphasis on the *differential type* of a definable set, as this notion yields a way to classify differential algebraic varieties according to their *transcendental dimension* (finite dimensional ones are precisely those of differential type zero).

Relative D-varieties and D-groups. The first step towards understanding finite dimensional differential algebraic varieties is simply to observe that, since their (differential) function field has finite transcendence degree, they can be generically described by algebraic polynomial equations together with first order differential equations of a very particular form. Loosely speaking, one obtains that all the derivations can be completely described in terms of algebraic relations. This was formalized through the concept of *algebraic D-variety* by Buium in [1], and later explored in model theoretic terms by Pillay in [51] and [48].

The situation is quite different for infinite dimensional differential algebraic varieties, their function fields have infinite transcendence degree and thus we need a way to measure how “transcendental” it really is. This is where the differential type comes into play, it gives us a way to differentiate between infinite dimensional differential algebraic varieties according to their (differential) transcendental properties. In loose terms, the differential type gives us the largest subset of the derivations that can be completely described in terms of the others. In particular, the differential type is a natural number between zero and the number of derivations.

We formalize the above ideas in Chapter 2 by introducing the category of *relative D-varieties*. Briefly, a *relative D-variety* is given by a partition $\mathcal{D} \cup \Delta$ of the total set of derivations, say Π , and a pair (V, s) where V is a Δ -algebraic variety and s is a Δ -algebraic section of a certain torsor of the Δ -tangent bundle of V called the *relative prolongation of V* and denoted by $\tau_{\mathcal{D}/\Delta}V$ (see Definition 2.3.3). The section s needs to satisfy a certain *integrability condition* (see Definition 2.4.1). Once in the right setting, we are able to prove a number of properties, including the following important characterization (Proposition 2.4.6):

Theorem A. *Let (K, Π) be a differential field in m commuting derivations Π , and suppose W is an irreducible Π -algebraic variety defined over K of Π -type $\ell < m$. Then there exist*

- *m linearly independent linear combinations $\{D_1, \dots, D_m\}$ of the derivations Π over the constant field of (K, Π) ,*
- *a $\{D_1, \dots, D_\ell\}$ -algebraic variety V defined over K of $\{D_1, \dots, D_\ell\}$ -type ℓ , and*
- *$\{D_1, \dots, D_\ell\}$ -polynomials $s_1, \dots, s_{m-\ell}$,*

such that W is Π -birationally equivalent to

$$\{a \in V : s_i(a) = D_{\ell+i}(a), i = 1, \dots, m - \ell\}.$$

As we have stated the theorem above, relative D-varieties did not appear explicitly. However, if we let $\Delta = \{D_1, \dots, D_\ell\}$, $\mathcal{D} = \{D_{\ell+1}, \dots, D_m\}$ and $s = (\text{Id}, s_1, \dots, s_{m-\ell})$, then (V, s) is a relative D-variety w.r.t. \mathcal{D}/Δ , and the displayed Π -algebraic variety in the conclusion is usually denoted by $(V, s)^\#$ and called the sharp points of (V, s) . The theorem says that every irreducible Π -algebraic variety is Π -birationally equivalent to the sharp points of a relative D-variety of the appropriate form. This is a precise generalization of the characterization of finite dimensional differential algebraic varieties in terms of algebraic D-varieties.

In Theorem A the birational equivalence between W and $(V, s)^\#$ might indeed be cutting some proper subvarieties, and hence it might not be an isomorphism. However, in the presence of a group structure this can be refined. We consider the group objects in the category of relative D-varieties to obtain the notion of *relative D-group*. Then, by applying a general version of Weil's group chunk theorem, we are able to show that in the case when W is a connected differential algebraic group, then the birational equivalence becomes an isomorphism of groups between W and the set of sharp points of a relative D-group (see Theorem 2.5.7). This result is used in Chapter 3 when we establish a connection between our two differential Galois theories.

Galois theory for partial differential fields. As in the algebraic case, differential Galois theory is the study of differential field extensions with well behaved group of automorphisms. The subject has its roots in the work of Picard and Vessiot at the end of the 19th century when they studied the solutions of linear differential equations. After that, several formalizations and generalizations of their work were considered, but it was not until Kolchin's *strongly normal extensions* entered the picture that differential Galois theory was systematically built as a whole new subject [24].

In Kolchin's work on strongly normal extensions the field of constants played an essential role. However, in [45], Pillay successfully replaced the constants for an arbitrary definable set and introduced the notion of a *generalized strongly normal extension*. There, he showed that these extensions have a *good* Galois theory. Even though Pillay's setting was ordinary differential fields, his results can be extended more or less immediately to the partial case, as long as the definable sets considered are finite dimensional. One of the purposes of Chapter 3 is to extend his results to the possibly infinite dimensional setting.

In algebraic Galois theory the Galois extensions are associated to polynomial equations. In the Picard-Vessiot linear differential Galois theory this remains true: the Picard-Vessiot extensions are the Galois extensions associated to homogeneous linear differential equations. In [44], Pillay shows that, under certain assumptions on the base field, his generalized strongly normal extensions are the Galois extensions associated to certain differential equations on (possibly non-linear) algebraic groups. We develop a Galois theory associated to *relative logarithmic differential equations* on relative D-groups (Section 3.3), associating to each such equation a generalized strongly normal extension. We show (Proposition 3.3.9 and Corollary 3.3.10):

Theorem B. *The Galois group of the generalized strongly normal extension L/K associated to a relative logarithmic differential equation, on a relative D-group (G, s) , is of the form $(H, t)^\#$ where (H, t) is a relative D-subgroup of (G, s) defined over K . Moreover, there is a Galois correspondence between the intermediate differential fields of L/K and the relative D-subgroups of (H, t) defined over K .*

The above theory simultaneously generalizes the parametrized Picard-Vessiot theory of Cassidy and Singer [7], and Pillay's finite dimensional Galois theory on algebraic D-groups [44].

While Galois extensions corresponding to relative logarithmic differential equations are instances of generalized strongly normal extensions, it is not known in general (i.e., in the possibly infinite dimensional case) if the converse is true. However, we can show the following generalization of Remark 3.8 of [44] (see Theorem 3.3.12 below):

Theorem C. *Suppose (K, Π) is a differential field with m commuting derivations and Δ is a set of fewer than m linearly independent linear combinations of the derivations in Π over the constant field of (K, Π) . Suppose, moreover, that L is a generalized strongly normal extension of K whose Π -type is bounded by Δ (see Definition 1.6.7) and such that K is Δ -closed. Then L/K is the Galois extension associated to a logarithmic differential equation on a relative D-group.*

Geometric axioms for $DCF_{0,m}$. Axiomatizations of the existentially closed models (i.e., models in which each consistent system of equations has a solution) of several theories

of fields with operators have been formulated in terms of algebro geometric objects. For difference fields this was done by Chatzidakis and Hrushovski in [8], for ordinary differential fields we have the characterization of Pierce and Pillay in [43], and, generalizing both, for the theory of fields with free operators considered by Moosa and Scanlon we have [40]. These *geometric axiomatizations* usually involve a scheme of axioms referring to certain dominant projections of irreducible algebraic varieties.

In [42], Pierce attempted to extend the geometric axioms for ordinary differentially closed fields (known as the Pierce-Pillay axioms) to the partial setting in a very natural way; however, he found a mistake in his argument. Essentially, the problem is that the commuting of the derivations is too strong of a condition to impose using simple algebraic relations. Nonetheless, in [41], he does find a geometric axiomatization, but his new axioms do not formally specialize to the Pierce-Pillay axioms, and ultimately have a very different flavour.

In Chapter 4, using the approach discussed at the beginning of the introduction, we treat differential fields with $\ell + 1$ commuting derivations as adding one derivation to differential fields with ℓ derivations. Then we are able to characterize differentially closed fields in $\ell + 1$ derivations in terms of differential algebro geometric objects in ℓ derivations in precisely the same spirit as the Pierce-Pillay axioms (see Theorem 4.2.1).

Theorem D. *$(K, \Delta \cup \{D\})$ is differentially closed if and only if*

- (i) *(K, Δ) is differentially closed*
- (ii) *Suppose V and W are irreducible affine Δ -algebraic varieties defined over K such that $W \subseteq \tau_{D/\Delta}V$ and W projects Δ -dominantly onto V . If O_V and O_W are nonempty Δ -open subsets of V and W respectively, defined over K , then there exists a K -point $a \in O_V$ such that $(a, Da) \in O_W$.*

However, some complications arise in this setting. It is a well known open problem as to whether irreducibility in differential algebraic geometry is first order expressible (it is in algebraic geometry). A similar issue arises with the property of a projection being Δ -dominant. Thus, it is not clear if our characterization above yields, as it is, a first order axiomatization. We bypass these issues, of differential irreducibility and dominance, by reformulating our characterization using the differential algebra machinery of *characteristic sets of prime differential ideals*. Given a finite set Λ of Δ -polynomials, we let $\mathcal{V}(\Lambda)$ be the zeros of Λ and $\mathcal{V}^*(\Lambda)$ be the zeros of Λ that are not zeros of the *initials* or *separants* of Λ (see Section 1.2 for definitions). We are able to prove the following (Theorem 4.2.4):

Theorem E. *$(K, \Delta \cup \{D\})$ is differentially closed if and only if*

- (i) (K, Δ) is differentially closed.
- (ii) Suppose Λ and Γ are characteristic sets of prime Δ -ideals of the Δ -polynomial rings $K\{x\}_\Delta$ and $K\{x, y\}_\Delta$ respectively, such that

$$\mathcal{V}^*(\Gamma) \subseteq \mathcal{V}(f, d_{D/\Delta}f : f \in \Lambda)$$

(see Section 2.2 for the definition of the operator $d_{D/\Delta}$). Suppose O is a nonempty Δ -open subset of $\mathcal{V}^*(\Lambda)$ defined over K such that the projection of $\mathcal{V}^*(\Gamma)$ to $\mathcal{V}(\Lambda)$ contains O . Then there is a K -point $a \in \mathcal{V}^*(\Lambda)$ such that $(a, Da) \in \mathcal{V}^*(\Gamma)$.

By the results of Tressl from [58], one can express in a first order fashion when a finite set of differential polynomials is a characteristic set of a prime differential ideal. Thus, Theorem E yields a geometric axiomatization of partial differentially closed fields.

Partial differential fields with an automorphism. For over fifteen years, it has been known that the class of existentially closed ordinary differential fields with an automorphism has a first order axiomatization (see for example [3]). However, the proof does not extend to the theory of partial differential fields with an automorphism. In Chapter 5 we settle this open problem.

In [15], Guzy and Rivière observed that the existentially closed partial differential fields with an automorphism are characterized by a certain differential algebro geometric condition (see Fact 2.1), analogous to Theorem D above. However, due to the same definable questions there, it remains open as to whether their characterization is first order. Motivated by the ideas used in Chapter 4, we apply the machinery of characteristic sets of prime differential ideals in Chapter 5 to prove (c.f. Corollary 5.2.4):

Theorem F. *The class of existentially closed partial differential fields with an automorphism is axiomatizable.*

In the usual manner, various basic model theoretic properties of this model companion are derived. In addition, we extend the Pillay-Ziegler [52] theory of differential and difference modules to this setting and prove an appropriate form of the *canonical base property* for finite dimensional types. As a consequence we obtain the following Zilber dichotomy for finite dimensional types (Corollary 5.4.11):

Theorem G. *If p is a finite dimensional type with $SU(p) = 1$, then p is either one-based or almost internal to the intersection of the fixed and constant fields (see Section 5.4 for the definitions).*

We expect that the finite dimensional assumption is not necessary, and that by extending the methods of Moosa, Pillay, Scanlon [37], to this setting, we can remove this assumption. However, we have not checked this in detail and thus omit it from this thesis.

Chapter 1

Preliminaries

This chapter is intended to provide the reader with the necessary background in differential algebra and differential algebraic geometry. We include also a very cursory introduction to model theory, while the more advanced model theoretic notions are reviewed by specializing them to the theory of differentially closed fields. For the abstract definitions we suggest [33] and [46]. We do take this opportunity to sometimes introduce new and convenient terminology, for example that of *bounding* and *witnessing* the differential type in Section 1.6 below.

1.1 Model theory

This brief section is intended for the reader that is not at all familiar with model theory. While basic model theory is formally a prerequisite for this thesis, we hope that an ignorance on the subject will not be an insurmountable obstacle to reading this thesis, especially because we will be focusing on a few very particular theories and will not work in the full abstraction that model theory allows. Nevertheless, a basic understanding of the model theoretic approach to mathematics is necessary, and this is what we present now. For a detailed and contemporary introduction to model theory we suggest David Marker's textbook [33]. Our discussion here is informed by [36].

A *structure* \mathcal{M} consists of a nonempty underlying set M , called the universe of \mathcal{M} , together with:

- (i) a set $\{c_i : i \in I_C\}$ of distinguished elements of M , called *constants*

- (ii) a set $\{f_i : M^{n_i} \rightarrow M : i \in I_F\}$ of distinguished maps, called *basic functions*
- (iii) a set $\{R_i \subseteq M^{k_i} : i \in I_R\}$ of distinguished subsets of cartesian powers of M , called *basic relations*

For example any unitary ring R is made into a structure by considering $(R, 0, 1, +, -, \times)$ where 0 and 1 are the constants, $+$ and \times are the basic binary functions, $-$ is the basic unary function, and there are no basic relations.

Given a structure \mathcal{M} , the *definable sets* in \mathcal{M} are the subsets of various cartesian powers of M that are obtained in a finitary manner from the constants, basic relations, and graphs of basic functions, by closing off under

- intersections,
- complements,
- images under coordinate projections, and
- fibers of coordinate projections.

The definable sets have a syntactic characterization that we do not explain in detail here, but it suffices to say that definable sets are *defined by* first order formulas in a *language* that has symbols for the basic constants, functions, and relations, as well as the standard logical symbols. These latter correspond to the above operations on definable sets in the natural way: intersections of definable sets correspond to *conjunctions* of formulae, complements to *negation*, images under projections to *existential quantifications*, and taking fibers of projections correspond to *specializations* or *naming parameters*. Note that by combining these we get also *disjunctions* and *universal quantifications*.

In general, the class of definable sets can be very wild; for example, there is a rigorous sense in which all of mathematics can be definably encoded (via Gödel's numbering) in the ring $(\mathbb{Z}, 0, 1, +, -, \times)$. At the opposite extreme we have tame mathematics; for example, a theorem of Chevalley (and also Tarski) implies that in $(K, 0, 1, +, -, \times)$, where K is an algebraically closed field, the definable sets are precisely the Zariski-constructible sets. In particular, the third operation (existential quantification) becomes superfluous: this goes by the term *quantifier elimination* in model theory.

One central aspect of model theory is the study of definable sets in a given structure. Another aspect is the study of classes of structures satisfying a given set of *first order axioms* in a fixed language. For example, in the language of unitary rings as above, we

can express the axioms of algebraically closed fields. Each axiom is a formula without free variables, and hence finitary, thus one requires infinitely many axioms in this case. One such collection of axioms consists of the axioms of fields together with a formula of the form

$$\forall a_{n-1} \cdots \forall a_1 \forall a_0 \exists x x^n + a_{n-1}x^{n-1} + \cdots + a_1x + a_0 = 0,$$

for each $n > 0$. This collection of axioms is usually denoted by *ACF* and referred to as the *theory of algebraically closed fields*. A structure in which these axioms hold is a *model* of the theory. Of course, there are deep connections between the study of theories and their models, and the study of definable sets in a given structure.

We end this section by listing, more or less as a glossary, some basic model theoretic notation. Let \mathcal{M} be a structure, \mathcal{L} be a language having the appropriate symbols, and T be an \mathcal{L} -theory (i.e., a set of axioms in the language \mathcal{L}).

- $\mathcal{M} \models T$. This means that \mathcal{M} is a model of the theory T ; that is, all the axioms in T are true in \mathcal{M} .
- $\mathcal{M} \models \phi(a)$. Here a is a (finite) tuple from M , $\phi(x)$ is a first order \mathcal{L} -formula in free variables x , and the notation means that a realises ϕ in \mathcal{M} . In other words, if $X \subseteq M^{|x|}$ is the set defined by ϕ then $a \in X$.
- $X(A) = \{a \in A^n : a \in X\}$. Here $A \subseteq M$ and X is a definable subset of M^n .
- $S_n(A)$. Here $A \subseteq M$. This means the space of *complete n -types over A* . A complete n -type p over A is a set of \mathcal{L} -formulas over A in n variables such that each finite subset of p is realised in \mathcal{M} and p is maximal with this property.
- $a \models p$. Here a is a tuple from M and p is a complete type over a subset of M . This means that a realises in \mathcal{M} all the formulas from p .
- $p^{\mathcal{M}} = \{a \in M : a \models p\}$. Here p is a complete type over a subset of M .
- $tp(a/A) = \{\phi(x) : \phi \text{ is an } \mathcal{L}\text{-formula over } A \text{ and } \mathcal{M} \models \phi(a)\}$. Here a is a tuple from M and $A \subseteq M$. This is the *complete type of a over A* . Clearly, $tp(a/A) \in S_n(A)$ where $n = |a|$.
- $dcl(A) = \{a \in M : \{a\} \text{ is an } A\text{-definable set}\}$. This is the *definable closure of A* . Note that $a \in dcl(A)$ iff there is an \mathcal{L} -formula over A such that a is the unique realisation in \mathcal{M} .

- $\text{acl}(A) = \{a \in \mathcal{M} : a \text{ is contained in a finite } A\text{-definable set}\}$. This is the model theoretic *algebraic closure of A*.
- *A is interdefinable (interalgebraic) with B over C*. Here A, B, C are subsets of M . This means that $A \cup C$ has the same definable (algebraic) closure as $B \cup C$.

1.2 Differential algebra

In this section we review the differential algebraic notions required for this thesis. For a complete treatment of differential algebra we refer the reader to [24].

Let R be a ring (unital and commutative) and S a ring extension. An additive map $\delta : R \rightarrow S$ is called a *derivation* if it satisfies the Leibniz rule; i.e., $\delta(ab) = \delta(a)b + a\delta(b)$. A ring R equipped with a set of derivations $\Delta = \{\delta_1, \dots, \delta_m\}$, $\delta_i : R \rightarrow R$ for $i = 1, \dots, m$, such that the derivations commute with each other is called a Δ -ring. In the case when R is a field we say Δ -field. The notions of Δ -homomorphism, Δ -isomorphism and Δ -automorphism are defined naturally.

Let R be a Δ -ring. The ring of Δ -constants of R will be denoted by

$$R^\Delta = \{a \in R : \delta a = 0 \text{ for all } \delta \in \Delta\}.$$

If S is a Δ -ring extension of R and $A \subseteq S$, we denote by $R\{A\}_\Delta$ the Δ -subring of S generated by A over R . In the case when S is a field, we denote by $R\langle A \rangle_\Delta$ the Δ -subfield of S generated by A over R .

The set of *derivative operators* is defined as the commutative monoid

$$\Theta_\Delta = \{\delta_m^{e_m} \cdots \delta_1^{e_1} : e_i < \omega\}.$$

If $x = (x_1, \dots, x_n)$ is a tuple of (differential) indeterminates, we set $\Theta_\Delta x = \{\theta x_i : \theta \in \Theta_\Delta, i = 1, \dots, n\}$ and call these the corresponding *algebraic indeterminates*. The ring of Δ -polynomials over R is defined as $R\{x\}_\Delta = R[\Theta_\Delta x]$. One can equip $R\{x\}_\Delta$ with the structure of a Δ -ring by defining

$$\delta_j(\delta_m^{e_m} \cdots \delta_1^{e_1} x_i) = \delta_m^{e_m} \cdots \delta_j^{e_j+1} \cdots \delta_1^{e_1} x_i$$

and then extending to all of $R\{x\}_\Delta$ using additivity and the Leibniz rule. In the case when R is an integral domain, we define the field of Δ -rational functions over R , denoted

by $R\langle x \rangle_\Delta$, as the field of fractions of $R\{x\}_\Delta$, or equivalently $R\langle x \rangle_\Delta = R(\Theta_\Delta x)$. By the quotient rule, $R\langle x \rangle_\Delta$ admits a canonical Δ -field structure.

When the set of derivations Δ is understood, we will often omit the subscripts in the differential notation. For example, Θ and $R\{x\}$ stand for Θ_Δ and $R\{x\}_\Delta$, respectively.

Given a Δ -ring R , an ideal I of $R\{x\}$ is said to be a Δ -ideal (or differential ideal) if $\delta f \in I$ for all $f \in I$ and $\delta \in \Delta$. Given a set $B \subseteq R\{x\}$, the ideal and the Δ -ideal generated by B will be denoted by (B) and $[B]$, respectively.

Fact 1.2.1 (Differential basis theorem, [24], Chap. III, §4). *Suppose R is a \mathbb{Q} -algebra that satisfies the ACC on radical Δ -ideals. If I is a radical Δ -ideal of $R\{x\}$, then there are $f_1, \dots, f_s \in R\{x\}_\Delta$ such that $I = \sqrt{[f_1, \dots, f_s]}$.*

A set $X \subseteq R^n$ is said to be Δ -closed if there is a set $B \subseteq R\{x\}$ such that $X = \{a \in R^n : f(a) = 0 \text{ for all } f \in B\}$. The differential basis theorem shows that, if R is a field, the Δ -closed subsets of R^n are the closed sets of a Noetherian topology, called the Δ -topology of R^n .

Characteristic sets. In the remainder of this section we discuss the somewhat technical, but very useful, notion of characteristic sets. The motivation here is that despite the Noetherianity of the Δ -topology, it is not necessarily the case that prime Δ -ideals are finitely generated as Δ -ideals. However, they are determined by a finite subset in a canonical way.

The *canonical ranking* on the algebraic indeterminates Θx is defined by

$$\delta_m^{e_m} \cdots \delta_1^{e_1} x_i < \delta_m^{r_m} \cdots \delta_1^{r_1} x_j \iff \left(\sum e_k, i, e_m, \dots, e_1 \right) < \left(\sum r_k, j, r_m, \dots, r_1 \right)$$

in the lexicographical order. Let $f \in R\{x\} \setminus R$. The *leader* of f , v_f , is the highest ranking algebraic indeterminate that appears in f . The *order* of f , $ord(f)$, is defined as $\sum e_k$ where $v_f = \delta_m^{e_m} \cdots \delta_1^{e_1} x_i$. The *degree* of f , d_f , is the degree of v_f in f . The *rank* of f is the pair (v_f, d_f) . If $g \in R\{x\} \setminus R$ we say that g has *lower rank* than f if $rank(g) < rank(f)$ in the lexicographical order. By convention, an element of R has lower rank than all the elements of $R\{x\} \setminus R$.

Let $f \in R\{x\} \setminus R$. The *separant* of f , S_f , is the formal partial derivative of f with respect to v_f , that is

$$S_f := \frac{\partial f}{\partial v_f} \in R\{x\}.$$

The *initial* of f , I_f , is the leading coefficient of f when viewed as a polynomial in v_f , that is, if we write

$$f = \sum_{i=0}^{d_f} g_i(v_f)^i$$

where $g_i \in R\{x\}$ and $v_{g_i} < v_f$ then $I_f = g_{d_f}$. Note that both S_f and I_f have lower rank than f .

Definition 1.2.2. Let $f, g \in R\{x\} \setminus R$. We say g is *partially reduced* with respect to f if no (proper) derivative of v_f appears in g . If in addition the degree of v_f in g is less than d_f we say that g is *reduced* with respect to f .

A finite set $\Lambda = \{f_1, \dots, f_s\} \subset R\{x\} \setminus R$ is said to be *autoreduced* if for all $i \neq j$ we have that f_i is reduced with respect to f_j . We will always write autoreduced sets in order of increasing rank, i.e., $\text{rank}(f_1) < \dots < \text{rank}(f_s)$. The canonical ranking on autoreduced sets is defined as follows: $\{g_1, \dots, g_r\} < \{f_1, \dots, f_s\}$ if and only if, either there is $i \leq r, s$ such that $\text{rank}(g_i) = \text{rank}(f_j)$ for $j < i$ and $\text{rank}(g_i) < \text{rank}(f_i)$, or $r > s$ and $\text{rank}(g_j) = \text{rank}(f_j)$ for $j \leq s$.

Fact 1.2.3 ([24], Chap. I, §10). *Every Δ -ideal I of $R\{x\}$ contains a lowest ranking autoreduced set. Any such set is called a characteristic set of I .*

Remark 1.2.4. If Λ is a characteristic set of I and $g \in R\{x\} \setminus R$ is reduced with respect to all $f \in \Lambda$, then $g \notin I$. Indeed, if g were in I then

$$\{f \in I : \text{rank}(f) < \text{rank}(g)\} \cup \{g\}$$

would be an autoreduced subset of I of lower rank than Λ , contradicting the minimality of Λ .

Suppose $\Lambda = \{f_1, \dots, f_s\}$ is a subset of $R\{x\} \setminus R$. We let $H_\Lambda := \prod_{i=1}^s I_{f_i} S_{f_i}$. Also, for each $\theta \in \Theta$, we let $(\Lambda)_\theta$ denote the ideal of $R\{x\}$ generated by the f_i 's and their derivatives whose leaders are strictly below θ . If I is an ideal of $R\{x\}$ and $g \in R\{x\}$, we denote the *saturated ideal of I over g* by

$$I : g^\infty = \{f \in R\{x\} : g^\ell f \in I \text{ for some } \ell \geq 0\}.$$

Definition 1.2.5. Let $\Lambda = \{f_1, \dots, f_s\}$ be a subset of $R\{x\}$. The set Λ is said to be *coherent* if the following two conditions are satisfied:

1. Λ is autoreduced.
2. For $i \neq j$, suppose there are θ_i and $\theta_j \in \Theta$ such that $\theta_i v_{f_i} = \theta_j v_{f_j} = v$ where v is the least such in the ranking. Then

$$S_{f_j} \theta_i f_i - S_{f_i} \theta_j f_j \in (\Lambda)_\theta : H_\Lambda^\infty.$$

Even though differential ideals of $R\{x\}$ are not in general generated by their characteristic sets, when R is a field of characteristic zero, prime differential ideals are determined by these. Moreover, we have a useful criterion to determine when a finite set of differential polynomials is the characteristic set of a prime differential ideal.

Fact 1.2.6 (Rosenfeld's criterion, [24], Chap. IV, §9). *Let K be a Δ -field of characteristic zero. If Λ is a characteristic set of a prime Δ -ideal \mathcal{P} of $K\{x\}$, then*

$$\mathcal{P} = [\Lambda] : H_\Lambda^\infty,$$

Λ is coherent, and $(\Lambda) : H_\Lambda^\infty$ is a prime ideal not containing a nonzero element reduced with respect to all $f \in \Lambda$. Conversely, if Λ is a coherent subset of $K\{x\}$ such that $(\Lambda) : H_\Lambda^\infty$ is prime and does not contain a nonzero element reduced with respect to all $f \in \Lambda$, then Λ is a characteristic set of a prime Δ -ideal of $K\{x\}$.

Example 1.2.7. Suppose K is a Δ -field of characteristic zero and $a_1, \dots, a_m \in K$ satisfy $\delta_j a_i = \delta_i a_j$ for $i, j = 1, \dots, m$. We use Rosenfeld's criterion to show that the Δ -ideal

$$[\delta_i x - a_i x : i = 1, \dots, m] \subseteq K\{x\}$$

is prime. Let $\Lambda = \{\delta_i x - a_i x : i = 1, \dots, m\}$. Note that $H_\Lambda = 1$, and that the assumption on the a_i 's imply that Λ is coherent. Clearly (Λ) is a prime ideal of $K\{x\}$. If $g \in K\{x\}$ is reduced with respect to all the elements of Λ , then g is an (algebraic) polynomial. Hence, (Λ) does not contain a nonzero element reduced with respect to the elements of Λ . By Rosenfeld's criterion, Λ is a characteristic set of the prime Δ -ideal

$$[\Lambda] : H_\Lambda^\infty = [\delta_i x - a_i x : i = 1, \dots, m].$$

1.3 Differentially closed fields

In this section we review the basic model theory of differential fields of characteristic zero with several commuting derivations. We will work in the first order language of partial differential rings with m derivations

$$\mathcal{L}_m = \{0, 1, +, -, \times\} \cup \{\delta_1, \dots, \delta_m\},$$

where $\Delta = \{\delta_1, \dots, \delta_m\}$ are unary function symbols. The first order \mathcal{L}_m -theory $DF_{0,m}$ consists of the axioms of fields of characteristic zero together with axioms asserting that Δ is a set of m commuting derivations.

Note that a Δ -field (K, Δ) is existentially closed (in the sense of model theory) if every system of Δ -polynomial equations with a solution in a Δ -field extension already has a solution in K . In [35], McGrail showed that the class of existentially closed Δ -fields is axiomatizable; this is the theory $DCF_{0,m}$ of *differentially closed fields with m commuting derivations of characteristic zero*, also called simply *Δ -closed fields*. We refer the reader to [35] for McGrail's original axioms, but let us give an alternative axiomatization given by Tressl in [58]. He observed that $(K, \Delta) \models DCF_{0,m}$ if and only if given a characteristic set $\Lambda = \{f_1, \dots, f_s\}$ of a prime differential ideal of $K\{x\}_\Delta$ then there exists a tuple a from K such that

$$f_1(a) = 0 \wedge \dots \wedge f_s(a) = 0 \wedge H_\Lambda(a) \neq 0.$$

In order to conclude that this is indeed a first order axiomatization, Tressl proved the following property (which will play an essential role for us in Chapters 4 and 5):

Fact 1.3.1. *Suppose (K, Δ) is a Δ -field of characteristic zero. The condition that “ $\Lambda = \{f_1, \dots, f_s\}$ is a characteristic set of a prime differential ideal of $K\{x\}_\Delta$ ” is a definable property (in the language \mathcal{L}_m) of the coefficients of f_1, \dots, f_s .*

This means that if $f_1(u, x), \dots, f_s(u, x)$ is a set of Δ -polynomials over \mathbb{Q} in the variables u and x , then the set

$$\{a \in K^{|u|} : \{f_1(a, x), \dots, f_s(a, x)\} \text{ is a characteristic set of a prime } \Delta\text{-ideal of } K\{x\}_\Delta\}$$

is definable in (K, Δ) .

Fact 1.3.1 is essentially an application of Rosenfeld's criterion (Fact 1.2.6 above) which reduces the problem to the classical problem of checking primality in polynomial rings in finitely many variables where uniform bounds are well known [59].

In Chapter 4 below we will give a new geometric axiomatization for $DCF_{0,m}$.

We now outline the model theoretic properties of $DCF_{0,m}$ that will be used throughout the thesis. For proofs and more details the reader is referred to [35] and [56].

Fact 1.3.2.

(i) *Every Δ -field embeds into a Δ -closed field.*

- (ii) If (K, Δ) is differentially closed, then the field of constants K^Δ is a pure algebraically closed field.
- (iii) $DCF_{0,m}$ eliminates quantifiers; that is, every definable set is a finite union of sets of solutions to systems of Δ -polynomial equations and inequations. In other words, the definable sets are precisely the Δ -constructible sets.
- (iv) $DCF_{0,m}$ eliminates imaginaries; that is, for every definable set X with a definable equivalence relation E there is a definable set Y and a definable function $f : X \rightarrow Y$ such that xEy if and only if $f(x) = f(y)$.
- (v) Suppose $(K, \Delta) \models DCF_{0,m}$ and $A \subseteq K$. Then $\text{dcl}(A) = \mathbb{Q}\langle A \rangle_\Delta$, the Δ -field generated by A , and $\text{acl}(A) = \mathbb{Q}\langle A \rangle_\Delta^{\text{alg}}$, the field theoretic algebraic closure of the Δ -field generated by A .

In somewhat the same way as we can consider the algebraic closure of a field, there exist *differential closures* of differential fields. A Δ -field extension \bar{K} of K is a Δ -closure of K if $(\bar{K}, \Delta) \models DCF_{0,m}$ and for every $(L, \Delta) \models DCF_{0,m}$ extending K there is an Δ -isomorphism from \bar{K} into L fixing K pointwise. The following result follows from general model theoretic facts (see [47]).

Fact 1.3.3. Δ -closures always exist and are unique up to Δ -isomorphism over K . Moreover, if \bar{K} is a Δ -closure of K then for every $a \in \bar{K}$ the type $\text{tp}(a/K)$ is isolated; that is, there is an \mathcal{L}_m -formula over K such that for all tuples b from \mathbb{U} we have

$$\mathbb{U} \models \phi(b) \iff \text{tp}(b/K) = \text{tp}(a/K).$$

A model (\mathbb{U}, Δ) of $DCF_{0,m}$ is said to be *saturated* if whenever $\Phi(x)$ is a (possibly infinite) system of Δ -polynomial equations and inequations over a Δ -subfield of cardinality less than $|\mathbb{U}|$, any finite subsystem of which has a solution in \mathbb{U} , then $\Phi(x)$ has a solution in \mathbb{U} . It follows from general model theoretic facts (see [46]) that for arbitrarily large cardinal κ there is a saturated differentially closed field of cardinality κ .

Fact 1.3.4. Let (\mathbb{U}, Δ) be a saturated model of $DCF_{0,m}$.

- (i) If $(K, \Delta) \models DCF_{0,m}$ and $|K| \leq |\mathbb{U}|$, then there is a Δ -isomorphism from K into \mathbb{U} .
- (ii) Suppose K is a Δ -subfield of \mathbb{U} with $|K| < |\mathbb{U}|$ and f is a Δ -isomorphism from K into \mathbb{U} , then there is a Δ -automorphism of \mathbb{U} extending f .

(iii) Suppose K is a Δ -subfield of \mathbb{U} with $|K| < |\mathbb{U}|$. If a and a' are tuples from \mathbb{U} , then $tp(a/K) = tp(a'/K)$ if and only if there is $f \in \text{Aut}_\Delta(\mathbb{U}/K)$ such that $f(a) = a'$. Here $\text{Aut}_\Delta(\mathbb{U}/K)$ is the set of Δ -automorphisms of \mathbb{U} fixing K pointwise.

(iv) Let $A \subset \mathbb{U}$ with $|A| < |\mathbb{U}|$. Then,

$$\text{dcl}(A) = \{a \in \mathbb{U} : f(a) = a \text{ for all } f \in \text{Aut}_\Delta(\mathbb{U}/A)\}$$

and

$$\text{acl}(A) = \{a \in \mathbb{U} : \text{the orbit of } a \text{ under } \text{Aut}_\Delta(\mathbb{U}/A) \text{ is finite}\}.$$

(v) Let $A \subset \mathbb{U}$ with $|A| < |\mathbb{U}|$. Then a definable set is definable over A if and only if every automorphism of \mathbb{U} fixing A pointwise fixes the definable set setwise.

The previous facts explain why it is convenient to work in a saturated model. It serves as a universal domain in which to study differential fields. Thus, we will usually work in a fixed sufficiently large saturated $(\mathbb{U}, \Delta) \models \text{DCF}_{0,m}$. Unless stated otherwise, all the parameter subsets of \mathbb{U} considered will be assumed to be small, i.e., of cardinality less than \mathbb{U} .

1.4 Independence and U -rank

We work in fixed sufficiently large saturated differentially closed field (\mathbb{U}, Δ) .

We have a robust notion of independence that is analogous to algebraic independence in algebraically closed fields. Let A, B, C be (small) subsets of \mathbb{U} . We say that A is *independent of B over C* , denoted by $A \downarrow_C B$, if and only if $\mathbb{Q}\langle A \cup C \rangle_\Delta$ is algebraically disjoint from $\mathbb{Q}\langle B \cup C \rangle_\Delta$ over $\mathbb{Q}\langle C \rangle_\Delta$. Abstractly this notion of independence comes from Shelah's nonforking in stable theories [55]; that is, if $C \subseteq B$ and a is a tuple, then $tp(a/B)$ *does not fork over C* , or *is a nonforking extension of $tp(a/C)$* , if and only if $a \downarrow_C B$. For our purposes, we may as well take the previous sentence as a definition of nonforking extension. We collect in the following fact the essential properties of independence.

Fact 1.4.1. *Let a be a tuple from \mathbb{U} and A, B, C subsets of \mathbb{U} .*

1. (Invariance) *If ϕ is an automorphism of \mathbb{U} and $A \downarrow_C B$, then $\phi(A) \downarrow_{\phi(C)} \phi(B)$.*
2. (Local character) *There is a finite subset $B_0 \subseteq B$ such that $A \downarrow_{B_0} B$.*

3. (Extension) There is a tuple b such that $tp(a/C) = tp(b/C)$ and $b \perp_C B$.
4. (Symmetry) $A \perp_C B$ if and only if $B \perp_C A$
5. (Transitivity) Assume $C \subseteq B \subseteq D$. Then $A \perp_C D$ if and only if $A \perp_C B$ and $A \perp_B D$.
6. (Stationarity) If C is an algebraically closed Δ -field and $C \subseteq B$, then there is a unique nonforking extension $p \in S_n(B)$ of $tp(a/C)$.

The notion of nonforking extension induces a partial ordering in the set of types, and the foundation rank of this partial ordering is a very important model theoretic invariant.

Definition 1.4.2. Let $p \in S_n(K)$. Then $U(p) \geq 0$, and for an ordinal α we define $U(p) \geq \alpha$ inductively as follows:

- 1) For a successor ordinal $\alpha = \beta + 1$, $U(p) \geq \alpha$ if there is a forking extension q of p such that $U(q) \geq \beta$.
- 2) For a limit ordinal α , $U(p) \geq \alpha$ if $U(p) \geq \beta$ for all $\beta < \alpha$.

We define $U(p) = \alpha$ if $U(p) \geq \alpha$ and $U(p) \not\geq \alpha + 1$. In case $p = tp(a/K)$ we sometimes write $U(a/K)$ for $U(p)$.

We now list some of the basic properties of U -rank.

Fact 1.4.3. Let a and b be tuples from \mathbb{U} .

- (i) $U(a/K)$ is an ordinal.
- (ii) $U(a/K) = 0$ if and only if $a \in \text{acl}(K)$. Consequently, $U(a/K) = 1$ if and only if $a \notin \text{acl}(K)$ but every forking extension of $tp(a/K)$ is algebraic (i.e., has only finitely many realisations).
- (iii) If L is a Δ -field extension of K , then $a \perp_K L$ if and only if $U(a/K) = U(a/L)$.
- (iv) Lascar inequalities:

$$U(a/K\langle b \rangle) + U(b/K) \leq U(a, b/K) \leq U(a/K\langle b \rangle) \oplus U(b/K).$$

Here \oplus stands for the Cantor sum of ordinals.

The U -rank of an arbitrary definable set $X \subseteq \mathbb{U}^n$ is defined by

$$U(X) = \sup\{U(a/F) : a \in X\}$$

where F is any (small) Δ -field over which X is defined. Using Fact 1.4.1, it is not hard to show that this is well defined (i.e., it is independent of the choice of F). By the same token, if $f : X \rightarrow Y$ is a definable bijection between definable sets then $U(X) = U(Y)$.

Example 1.4.4. Since the field of constants \mathbb{U}^Δ is a pure algebraically closed field, every definable subset of \mathbb{U}^Δ is either finite or cofinite. That is, by definition, the field of constants forms a *strongly minimal set*. It follows from Fact 1.4.3 (ii) that the constant field is of U -rank 1.

The U -rank of \mathbb{U} is ω^m . See [34] for a detailed proof of this in the case when $m = 1$. As we could not find a correct proof of the general case in the literature, we have included one in the appendix of this thesis. We also include in the appendix some U -rank computations of certain definable subgroups of the additive group.

1.5 Affine differential algebraic geometry

We fix a sufficiently large saturated model $(\mathbb{U}, \Delta) \models DCF_{0,m}$ and a base (small) Δ -subfield K .

Definition 1.5.1. An *affine Δ -algebraic variety* V (or *affine Δ -variety* for short) *defined over* K is a Δ -closed subset of \mathbb{U}^n (for some n) defined over K . That is, for some set of Δ -polynomials $B \subseteq K\{x\}_\Delta$,

$$V = \{a \in \mathbb{U}^n : f(a) = 0 \text{ for all } f \in B\}.$$

Remark 1.5.2.

- (i) Our approach here is in the spirit of classical rather than contemporary algebraic geometry: for us affine Δ -varieties are a set of points rather than a scheme. The fact that we are working in a saturated model (so a universal domain in the spirit of Weil) means that this is not an unreasonable way to proceed. Scheme theoretic approaches to differential algebraic geometry have been certainly developed (see for example [27]), but there are subtleties there that we wish to avoid.

(ii) Every affine Δ -variety is of course a definable set in (\mathbb{U}, Δ) . On the other hand, quantifier elimination (Fact 1.3.2 (iii)) tells us that every definable set is a finite boolean combination of affine Δ -varieties.

Definition 1.5.3. A map $f : V \rightarrow W$ between affine Δ -varieties defined over K is Δ -regular over K if there is a finite cover $\{O_i\}$ of V by Δ -open subsets defined over K and a family $\{g_i\}$ of tuples of Δ -rational functions over K such that $f|_{O_i} = g_i|_{O_i}$.

Given $A \subseteq K\{x\}$ and $V \subseteq \mathbb{U}^n$, we let

$$\mathcal{V}(A) := \{a \in \mathbb{U}^n : f(a) = 0 \text{ for all } f \in A\}$$

and

$$\mathcal{I}(V/K) := \{f \in K\{x\} : f(a) = 0 \text{ for all } a \in V\}.$$

The set $\mathcal{V}(A)$ is called the *vanishing set of A* , and $\mathcal{I}(V/K)$ is called the *defining Δ -ideal of V over K* . It is easy to check that $\mathcal{I}(V/K)$ is a radical differential ideal, and that $\mathcal{V}(\mathcal{I}(V/K))$ is the Δ -closure of V over K ; that is, the smallest Δ -closed set defined over K containing V . If a is a tuple from \mathbb{U} we define the *Δ -locus of a over K* as $\mathcal{V}(\mathcal{I}(a/K))$.

Fact 1.5.4 (Differential Nullstellensatz, [24], Chap. IV, §2). *Let $A \subseteq K\{x\}$. Then $\mathcal{I}(\mathcal{V}(A)/K) = \sqrt{[A]}$.*

An affine Δ -variety V defined over K is said to be *irreducible over K* , or *K -irreducible*, if it is not the union of two proper Δ -subvarieties defined over K ; equivalently, the defining Δ -ideal of V over K is prime. By the differential basis theorem (Fact 1.2.1), every affine Δ -variety V defined over K has a unique irredundant decomposition into K -irreducible Δ -varieties, which are called the *K -irreducible components of V* .

Fact 1.5.5 ([34], Chap. II, §5). *If V is K^{alg} -irreducible then V is (absolutely) irreducible; that is, V is F -irreducible for all Δ -field extensions F of K .*

Working in a saturated model gives us Δ -generic points:

Definition 1.5.6. Let V be a K -irreducible affine Δ -variety. A point $a \in V$ is called a *Δ -generic point of V over K* (or simply generic when Δ and K are understood) if a is not contained in any proper Δ -subvariety of V defined over K .

Note, for example, that any a is a Δ -generic point over K of its own Δ -locus over K .

Let us show that Δ -generic points always exist. Let V be K -irreducible. Consider the set of \mathcal{L}_m -formulas over K

$$\Phi = \{x \in V\} \cup \{x \notin W : W \text{ is a proper } \Delta\text{-subvariety of } V \text{ defined over } K\}.$$

If this set were finitely inconsistent, then V would not be K -irreducible. Hence, by saturation, Φ is realised in \mathbb{U} , and such a realisation is the desired Δ -generic point.

If a and a' are Δ -generic points over K of a K -irreducible affine Δ -variety V , then, by quantifier elimination for $DCF_{0,m}$ (Fact 1.3.2), $tp(a/K) = tp(a'/K)$. Hence, we can define the Δ -generic type of V over K to be the complete type over K of any Δ -generic point of V over K . In fact, this yields a bijective correspondence between K -irreducible affine Δ -varieties and complete types over K , given by

$$V \mapsto \text{the } \Delta\text{-generic type of } V \text{ over } K.$$

The inverse of this correspondence is given by taking the Δ -locus over K of any realisation of the type.

The differential analogue of the tangent bundle of an algebraic variety is given by the following construction of Kolchin (c.f. [24], Chap. VIII, §2).

Definition 1.5.7. Let V be an affine Δ -variety defined over K . The Δ -tangent bundle of V , denoted by $T_\Delta V$, is the affine Δ -variety defined by

$$f(x) = 0 \quad \text{and} \quad \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(x) \theta u_i = 0,$$

for all $f \in \mathcal{I}(V/K)$, together with the projection onto the first n coordinates $\rho : T_\Delta V \rightarrow V$.

An *affine Δ -algebraic group defined over K* is a group object in the category of affine Δ -varieties defined over K . In other words, an affine Δ -algebraic group over K is a group whose underlying universe is an affine Δ -algebraic variety and the group operation $p : G \times G \rightarrow G$ is a Δ -regular map, both over K .

The following property of affine Δ -algebraic groups, which is not necessarily true for Δ -varieties, will be important for us in the appendix.

Fact 1.5.8 ([12], §2). *If G is a connected affine Δ -algebraic group defined over K , then $U(G) = U(p)$ where p is the Δ -generic type of G over K .*

We conclude this section with the following proposition describing the geometric meaning of characteristic sets.

Proposition 1.5.9. *Suppose V is a K -irreducible affine Δ -variety. Let Λ be a characteristic set of the prime Δ -ideal $\mathcal{I}(V/K)$.*

(i) *Let $\mathcal{V}^*(\Lambda) = \mathcal{V}(\Lambda) \setminus \mathcal{V}(H_\Lambda)$, where H_Λ is the product of the initials and separants of Λ (see Section 1.2). Then $\mathcal{V}^*(\Lambda)$ is nonempty.*

(ii) $\mathcal{V}^*(\Lambda) = V \setminus \mathcal{V}(H_\Lambda)$.

(iii) V is a K -irreducible component of $\mathcal{V}(\Lambda)$.

Proof.

(i) If $\mathcal{V}(\Lambda) \subseteq \mathcal{V}(H_\Lambda)$ then, by the differential Nullstellensatz (Fact 1.5.4), H_Λ would be in $\sqrt{[\Lambda]} \subseteq \mathcal{I}(V/K)$ contradicting Remark 1.2.4. Thus $\mathcal{V}(\Lambda) \setminus \mathcal{V}(H_\Lambda) \neq \emptyset$.

(ii) Let $a \in \mathcal{V}^*(\Lambda)$. We need to show that $f(a) = 0$ for all $f \in \mathcal{I}(V/K)$. Clearly if $f \in [\Lambda]$ then $f(a) = 0$. Let $f \in \mathcal{I}(V/K)$, since $\mathcal{I}(V/K) = [\Lambda] : H_\Lambda^\infty$, we can find ℓ such that $H_\Lambda^\ell f \in [\Lambda]$. Hence, $H_\Lambda(a)f(a) = 0$, but $H_\Lambda(a) \neq 0$, and so $f(a) = 0$. The other containment is clear.

(iii) Let W be an irreducible component of $\mathcal{V}(\Lambda)$ containing V . Since $\mathcal{V}^*(\Lambda) = V \setminus \mathcal{V}(H_\Lambda)$, we have that $W \setminus \mathcal{V}(H_\Lambda) = V \setminus \mathcal{V}(H_\Lambda) \subseteq V$. Hence V contains a nonempty Δ -open set of W , and so, by irreducibility of W , $V = W$. \square

1.6 Differential type and typical differential dimension

We continue to work in a fixed sufficiently large saturated $(\mathbb{U}, \Delta) \models DCF_{0,m}$ and over a (small) base Δ -subfield K of \mathbb{U} . In this section we discuss the Kolchin polynomial and the associated notions of differential type and typical differential dimension. As these are central to our work in this thesis we will give more details than in previous sections.

Let L be a Δ -subfield of \mathbb{U} containing K . A set $A \subseteq L$ is said to be Δ -algebraically independent over K if for every (finite) tuple a from A we have that $\mathcal{I}(a/K) = \{0\}$. A maximal subset of L that is Δ -algebraically independent over K is called a Δ -transcendence basis of L over K . An element $a \in L$ is said to be Δ -transcendental over K if $\{a\}$ is Δ -algebraically independent over K , and Δ -algebraic over K otherwise.

Fact 1.6.1 ([24], Chap. II, §9). *There exists a Δ -transcendence basis of L over K , and any two such basis have the same cardinality.*

By virtue of the previous fact one defines the Δ -transcendence degree of L over K as the cardinality of any Δ -transcendence basis of L over K .

By a *numerical polynomial* we mean a polynomial (in one variable) $f \in \mathbb{R}[x]$ such that if $e = \deg(f)$ then there are unique $d_0, \dots, d_e \in \mathbb{Z}$ with $f(x) = \sum_{i=0}^e d_i \binom{x+i}{i}$. Note that if f is a numerical polynomial then $f(h) \in \mathbb{Z}$ for sufficiently large $h \in \mathbb{N}$.

We denote by $\Theta_\Delta(h)$ the set of derivative operators of order less than or equal to $h \in \mathbb{N}$.

Fact 1.6.2 ([24], Chap.2, §12). *Let $a = (a_1, \dots, a_n)$ be a tuple from \mathbb{U} . There exists a numerical polynomial $\omega_{a/K}$ with the following properties:*

- (i) *For sufficiently large $h \in \mathbb{N}$, $\omega_{a/K}(h)$ equals the transcendence degree of $K(\theta a_i : i = 1, \dots, n, \theta \in \Theta_\Delta(h))$ over K .*
- (ii) *$\deg \omega_{a/K} \leq m$.*
- (iii) *If we write $\omega_{a/K} = \sum_{i=0}^m d_i \binom{x+i}{i}$ where $d_i \in \mathbb{Z}$, then d_m equals the Δ -transcendence degree of $K\langle a \rangle_\Delta$ over K .*
- (iv) *If b is another tuple with each coordinate in $K\langle a \rangle$, then there is $h_0 \in \mathbb{N}$ such that for sufficiently large $h \in \mathbb{N}$ we have $\omega_{b/K}(h) \leq \omega_{a/K}(h + h_0)$.*

The polynomial $\omega_{a/K}$ is called the *Kolchin polynomial of a over K* . Even though the Kolchin polynomial is not in general a Δ -birational invariant, (iv) of Fact 1.6.2 shows that its degree is. We call this degree the Δ -type of a over K and denote it by $\Delta\text{-type}(a/K)$. Similarly, if we write $\omega_{a/K} = \sum_{i=0}^m d_i \binom{x+i}{i}$ where $d_i \in \mathbb{Z}$, the coefficient d_τ , where $\tau = \Delta\text{-type}(a/K)$, is a Δ -birational invariant. We call d_τ the *typical Δ -dimension of a over K* and denote it by $\Delta\text{-dim}(a/K)$. We will adopt the convention that if $\omega_{a/K} = 0$ then $\Delta\text{-type}(a/K) = 0$ and $\Delta\text{-dim}(a/K) = 0$.

Given a complete type $p = tp(a/K)$ its Kolchin polynomial ω_p is defined to be $\omega_{a/K}$. Thus it makes sense to talk about the Δ -type and typical Δ -dimension of complete types.

Now we extend these concepts to definable sets. First recall that we can put a total ordering on numerical polynomials by eventual domination, i.e., $f \leq g$ if and only if $f(h) \leq g(h)$ for all sufficiently large $h \in \mathbb{N}$.

Definition 1.6.3. Let X be a definable set. The *Kolchin polynomial* of X is defined by

$$\omega_X = \sup \{ \omega_{a/F} : a \in X \},$$

where F is any Δ -field over which X is defined (the fact that this does not depend on the choice of F is a consequence, for example, of Theorem 4.3.10 of [35]). In Lemma 1.6.4 below we will see that $\omega_X = \omega_{a/F}$ for some $a \in X$. We define the Δ -type of X and its typical Δ -dimension in the obvious way. Also, X is said to be *finite dimensional* if $\Delta\text{-type}(X) = 0$.

Lemma 1.6.4. *Let X be a K -definable set and let V_1, \dots, V_s be the K -irreducible components of the Δ -closure of X over K (in the Δ -topology). If p_i is the Δ -generic type of V_i over K , then*

$$\omega_X = \max \{ \omega_{p_i} : i = 1, \dots, s \}.$$

Proof. Let $b \in X$, then b is in some V_i . Let a_i be a realisation of p_i . By Δ -genericity, there is a Δ -ring homomorphism $f_i : K\{a_i\}_\Delta \rightarrow K\{b\}_\Delta$ over K such that $a_i \mapsto b$. Thus, for any $h \in \mathbb{N}$ the transcendence degree of $K(\theta a_{i,j} : j = 1, \dots, n, \theta \in \Theta_\Delta(h))$ over K , where $a_i = (a_{i,1}, \dots, a_{i,n})$, is greater than or equal to the transcendence degree of $K(\theta b_j : j = 1, \dots, n', \theta \in \Theta_\Delta(h))$ over K , where $b = (b_1, \dots, b_{n'})$. This implies that, for sufficiently large $h \in \mathbb{N}$, $\omega_{a_i/K}(h) \geq \omega_{b/K}(h)$ and hence $\omega_{a_i/K} \geq \omega_{b/K}$. \square

In the ordinary case of $\Delta = \{\delta\}$, if an element a is $\{\delta\}$ -algebraic over K then $\delta^{k+1}a \in K(a, \delta a, \dots, \delta^k a)$ for some k (see for example Chapter 1.1 of [34]). In general, for arbitrary Δ , it may occur that a is Δ -algebraic over K but $\delta_m^{k+1}a \notin K\langle a, \delta_m a, \dots, \delta_m^k a \rangle_{\Delta'}$ for every k , where $\Delta' = \Delta \setminus \{\delta_m\}$. However, if we allow linearly independent transformations of Δ over K^Δ , an analogous result holds, which we explain now.

Observe that any linear combination of the derivations in Δ over the constant field \mathbb{U}^Δ is again a derivation on \mathbb{U} . As we are working over K we will be considering the K^Δ -vector space $\text{span}_{K^\Delta} \Delta$ of derivations of \mathbb{U} obtained by taking linear combinations of the derivations in Δ over the constant field of (K, Δ) . Note that if $\Delta' \subseteq \text{span}_{K^\Delta} \Delta$, then any Δ -field extension of K is also a Δ' -field. One consequence of \mathbb{U} being differentially closed is that Δ is a basis for $\text{span}_{K^\Delta} \Delta$. Moreover, by verifying the definition of existentially closed, it is not hard to see that if $\Delta' \subseteq \text{span}_{K^\Delta} \Delta$ is a K^Δ -linearly independent set, then (\mathbb{U}, Δ') is itself differentially closed and in fact saturated. Shifting from Δ to some other basis for $\text{span}_{K^\Delta} \Delta$ is something that turns out to be quite useful. For example:

Fact 1.6.5 ([24], Chap. 2, §11). *Let $a = (a_1, \dots, a_n)$ be a tuple from \mathbb{U} such that each a_i is Δ -algebraic over K . Then there is a natural number $k > 0$ and a basis $\Delta' \cup \{D\}$ of the K^Δ -vector space $\text{span}_{K^\Delta} \Delta$ such that $K\langle a \rangle_\Delta = K\langle a, D, \dots, D^k a \rangle_{\Delta'}$.*

Fact 1.6.5 implies the following fact which will be of central importance to us.

Fact 1.6.6 ([24], Chap. 2, §13). *Let a be a tuple from \mathbb{U} . Then there is a set Δ' of linearly independent elements of the K^Δ -vector space $\text{span}_{K^\Delta} \Delta$ with $|\Delta'| = \Delta\text{-type}(a/K)$, and a tuple α from \mathbb{U} , such that*

$$K\langle a \rangle_\Delta = K\langle \alpha \rangle_{\Delta'}$$

with $\Delta'\text{-type}(\alpha/K) = \Delta\text{-type}(a/K)$ and $\Delta'\text{-dim}(\alpha/K) = \Delta\text{-dim}(a/K)$.

Fact 1.6.6 says that a tuple a has Δ -type less than or equal to $\tau \in \mathbb{N}$ if and only if there is a set Δ' of linearly independent elements of $\text{span}_{K^\Delta} \Delta$ with $|\Delta'| = \tau$ such that $K\langle a \rangle_\Delta$ is finitely Δ' -generated over K . If in addition the Δ' -transcendence degree of $K\langle a \rangle_\Delta$ over K is positive then $\Delta\text{-type}(a/K) = \tau$. This motivates the following terminology (which we take the liberty to introduce).

Definition 1.6.7 (*Bounding and witnessing the Δ -type*). Let a be a (finite) tuple. A set Δ' of linearly independent elements of $\text{span}_{K^\Delta} \Delta$ is said to *bound the Δ -type of a over K* if the Δ -field $K\langle a \rangle_\Delta$ is finitely Δ' -generated over K . Moreover, if Δ' bounds the Δ -type of a over K and $|\Delta'| = \Delta\text{-type}(a/K)$ then we will say that Δ' *witnesses the Δ -type of a over K* .

So Fact 1.6.6 says that given any finite tuple a we can always find a Δ' witnessing the Δ -type of a over K .

Lemma 1.6.8. *Let a be a tuple, L a Δ -field extension of K , and $\Delta' \subseteq \text{span}_{K^\Delta} \Delta$ linearly independent. Suppose $a \perp_K L$. Then Δ' bounds $\Delta\text{-type}(a/L)$ if and only if it bounds $\Delta\text{-type}(a/K)$.*

Proof. Assume Δ' bounds $\Delta\text{-type}(a/L)$. Then we can find α a tuple of the form

$$(a, \theta_1 a, \dots, \theta_s a : \theta_i \in \Theta_\Delta, s < \omega)$$

such that $L\langle a \rangle_\Delta = L\langle \alpha \rangle_{\Delta'}$; here if $a = (a_1, \dots, a_n)$ and $\theta \in \Theta_\Delta$ then $\theta a = (\theta a_1, \dots, \theta a_n)$. It suffices to prove the following:

Claim. $K\langle a \rangle_\Delta = K\langle \alpha, \delta_1 \alpha, \dots, \delta_m \alpha \rangle_{\Delta'}$.

Let $w \in K\langle \alpha, \delta_1 \alpha, \dots, \delta_m \alpha \rangle_{\Delta'}$, we need to show that $\delta w \in K\langle \alpha, \delta_1 \alpha, \dots, \delta_m \alpha \rangle_{\Delta'}$ for all $\delta \in \Delta$. Since $w \in L\langle \alpha \rangle_\Delta = L\langle \alpha \rangle_{\Delta'}$, then $w = f(\alpha)$ for some Δ' -rational function f over L . Thus, there is a tuple β of the form $(\alpha, \theta'_1 \alpha, \dots, \theta'_r \alpha : \theta'_i \in \Theta_{\Delta'}, r < \omega)$ such that

$w = g(\beta)$ for some rational function g over L . Let ρ be a tuple whose entries form a maximal algebraically independent (over K) subset of the set consisting of the entries of β . Then (w, ρ) is algebraically dependent over L . But since $a \downarrow_K L$, $K\langle a \rangle_\Delta$ is algebraically disjoint from L over K . Thus, since (w, ρ) is from $K\langle a \rangle_\Delta$, we get that (w, ρ) is algebraically dependent over K . Thus, there is $h \in K[x, y]$ such that $h(w, \rho) = 0$. Moreover, since ρ is algebraically independent over K , w appears non trivially in $h(w, \rho)$. Hence, w is algebraic over $K(\rho)$ and so $\delta w \in K(w, \rho, \delta\rho) \subseteq K\langle \alpha, \delta_1\alpha, \dots, \delta_m\alpha \rangle_{\Delta'}$. \square

Definition 1.6.9. Let X be a K -definable set. A linearly independent subset $\Delta' \subseteq \text{span}_{K\Delta} \Delta$ is said to *bound the Δ -type of X* if Δ' bounds $\Delta\text{-type}(a/K)$ for all $a \in X$. We say Δ' *witnesses the Δ -type of X* if Δ' bounds the Δ -type of X and $|\Delta'| = \Delta\text{-type}(X)$.

By Lemma 1.6.8, this definition does not depend on the choice of K . In other words, Δ' bounds the Δ -type(X) if and only if there is a Δ -field F , over which X and Δ' are defined, such that Δ' bounds the Δ -type(a/F) for all $a \in X$. Indeed, suppose that Δ' bounds the Δ -type of X , F is a Δ -field over which X and Δ' are defined, and $a \in X$. One can assume, without loss of generality, that $F < K$. Let b be a tuple from \mathbb{U} such that $b \downarrow_F K$ and $tp(b/F) = tp(a/F)$. By assumption, Δ' bounds the Δ -type(b/K), and so, by Lemma 1.6.8, Δ' bounds the Δ -type(b/F). Hence, Δ' bounds Δ -type(a/F).

Remark 1.6.10. It is not clear (at least to the author) that every definable set has a witness to its Δ -type. However, for definable groups one can always find such a witness. To see this let G be a K -definable group. Note that, since the property “ Δ' witnesses the Δ -type” for definable sets is preserved under definable bijection, a witness to the Δ -type of the connected component of G will be a witness to the Δ -type of G . Hence, we may assume that G is connected. Let p be the generic type of G over K . By Fact 1.6.6, there is Δ' , a linearly independent subset of $\text{span}_{K\Delta} \Delta$, such that for any $a \models p$ the Δ -field $K\langle a \rangle_\Delta$ is finitely Δ' -generated over K . Let us check that this Δ' witnesses the Δ -type of G . Let $g \in G$, then there are $a, b \models p$ such that $g = a \cdot b$. Indeed, choose $a \models p$ such that $a \downarrow_K g$ and let $b = a^{-1} \cdot g$ (see §7.2 of [33] for details). Thus $K\langle g \rangle_\Delta \leq K\langle a, b \rangle_\Delta$, but since the latter is finitely Δ' -generated over K the former is as well.

We conclude this section with a remark on how the differential type interacts with the model theoretic notion of internality. Let X be a K -definable set. A complete type $tp(a/K)$ is said to be *internal to X* if there is Δ -field extension L of K with $a \downarrow_K L$ and a tuple c from X such that $a \in \text{dcl}(L, c) = L\langle c \rangle_\Delta$.

Lemma 1.6.11. *Let X be a K -definable set and suppose that $tp(a/K)$ is X -internal. Then*

1. $\Delta\text{-type}(a/K) \leq \Delta\text{-type}(X)$.

2. If Δ' bounds Δ -type(X) then Δ' also bounds Δ -type(a/K).

Proof.

(1) For any $d = (d_1, \dots, d_r)$, Δ -type(d/K) = $\max_{i \leq r} \{\Delta$ -type(d_i/K) $\}$ (see for example Lemma 3.1 of [37]). Now let L be a Δ -field extension of K such that $a \downarrow_K L$, and c from X such that $a \in L\langle c \rangle_\Delta$. Hence,

$$\Delta\text{-type}(a/K) = \Delta\text{-type}(a/L) \leq \Delta\text{-type}(c/L) \leq \Delta\text{-type}(X),$$

where the equality follows from Lemma 1.6.8 and the first inequality holds by Fact 1.6.2 (iv).

(2) Assume Δ' bounds Δ -type(X), then $L\langle c \rangle_\Delta$ is finitely Δ' -generated over L . Since $L\langle a \rangle_\Delta \leq L\langle c \rangle_\Delta$, then $L\langle a \rangle_\Delta$ is also finitely Δ' -generated over L . But $a \downarrow_K L$, so, by Lemma 1.6.8, Δ' bounds the Δ -type of a over K . \square

1.7 Abstract differential algebraic varieties

In Section 1.5 we defined the notion of an affine differential algebraic variety. While these affine objects are quite concrete and provide the proper setting for the presentation of most of this thesis, in Sections 2.5 and 3.3 we will need a more general notion of differential algebraic variety. In this section we discuss abstract differential algebraic varieties as presented in [50] (see Chap. 7 of [33] for the algebraic case). We then naturally extend the concepts of Sections 1.5 and 1.6 to this more general setting.

We continue to work in a sufficiently large saturated $(\mathbb{U}, \Delta) \models DCF_{0,m}$ and over a base Δ -subfield K .

By a *quasi-affine Δ -variety* defined over K we mean a Δ -open subset of an affine Δ -variety, where both are defined over K . We define Δ -regular maps between quasi-affine Δ -varieties in the natural way (i.e., as in Definition 1.5.3). The idea behind the construction of abstract Δ -varieties is to build them up from affine ones (or even quasi-affine) in the same way that manifolds are built from open sets of \mathbb{R}^n or \mathbb{C}^n . This is the analogue of Weil's definition of abstract algebraic varieties.

Definition 1.7.1. An *abstract Δ -algebraic variety* (or *Δ -variety* for short) is a topological space V equipped with a finite open cover $\{V_i\}_{i=1}^s$ and homeomorphisms $f_i : V_i \rightarrow U_i$, where U_i is an affine Δ -variety, such that

- (i) $U_{ij} = f_i(V_i \cap V_j)$ is a Δ -open subset of U_i , and

(ii) $f_{ij} = f_i \circ f_j^{-1} : U_{ji} \rightarrow U_{ij}$ is a Δ -regular map of quasi-affine varieties.

The pairs (V_i, f_i) are called the charts of the Δ -variety. Hence, a Δ -variety V is given by the information $(V, \{V_i\}, \{f_i\})$. We say V is defined over K if the U_i 's and the f_{ij} 's are defined over K .

It is easy to see that affine, and more generally quasi-affine, Δ -varieties are abstract Δ -varieties.

Definition 1.7.2. Suppose $(V, \{V_i\}, \{f_i\})$ and $(W, \{W_j\}, \{g_j\})$ are Δ -varieties defined over K . We say that a map $h : V \rightarrow W$ is Δ -regular if $g_j \circ h \circ f_i^{-1} : f_i(V_i) \rightarrow g_j(W_j)$ is a Δ -regular map of affine Δ -varieties for each i, j . We say that h is defined over K if all the $g_j \circ h \circ f_i^{-1}$ are defined over K .

Let $(V, \{V_i\}, \{f_i\})$ be a Δ -variety. By an *affine open set* of V we mean an open subset of V of the form $f_i^{-1}(W)$ for some i and W a Δ -open subset of $f_i(V_i)$. Any open subset of V is a finite union of affine open subsets. This implies that the topology of V is Noetherian. We can define irreducibility and K -irreducibility of V in the natural way. It follows, from Noetherianity, that if V is defined over K then it has a unique irredundant decomposition into K -irreducible Δ -varieties.

As the reader might expect, since abstract Δ -varieties are modelled by affine ones, in most cases one can reduce constructions and proofs of statements about abstract Δ -varieties to the affine case. More precisely, one chooses an affine cover and works with coordinates in affine space using a local chart. We will however, in some cases, carry out certain constructions and proofs to make the concepts clear.

We now extend some of the concepts that were defined for affine Δ -varieties.

Let $(V, \{V_i\}, \{f_i\})$ be a K -irreducible Δ -variety. A point $a \in V$ is said to be a *Δ -generic point of V over K* if it is not contained in any proper Δ -subvariety defined over K . As in the affine case, Δ -generic points always exist. Any a of the form $f_i^{-1}(b)$, where b is a Δ -generic point over K of the affine K -irreducible Δ -variety $f_i(V_i)$, is a Δ -generic point of V over K . Moreover, for every i , any Δ -generic point of V will be of this form. Hence, if we identify any two Δ -generic points over K with their images under the same f_i , they will have the same complete type over K . We define the *Δ -generic type of V over K* to be this complete type. Notice that the Δ -generic type depends on the embedding f_i . However, if p_i and p_j are the Δ -generic types of V over K with respect to f_i and f_j , then there are $a_i \models p_i$ and $a_j \models p_j$ such that a_i and a_j are interdefinable over K (to see this simply use the transition functions f_{ij}).

Let $(V, \{V_i\}, \{f_i\})$ be a Δ -variety defined over K and $a \in V$. We define the *Kolchin polynomial* of a over K to be $\omega_{a/K} = \max\{\omega_{f_i(a)/K} : i = 1, \dots, s\}$. The Δ -type and typical Δ -dimension of a over K are defined in the obvious way. It follows, using the transition functions f_{ij} , that $\Delta\text{-type}(a/K) = \Delta\text{-type}(f_i(a)/K)$ and $\Delta\text{-dim}(a/K) = \Delta\text{-dim}(f_i(a)/K)$ for any i . That is, the notions of Δ -type and typical Δ -dimension of a over K do not depend on the embedding.

The *Kolchin polynomial* of V is defined as

$$\omega_V = \max\{\omega_{f_i(V_i)} : i = 1, \dots, s\}.$$

We define the Δ -type and typical Δ -dimension of V in the obvious way. It follows that if V is K -irreducible then $\Delta\text{-type}(V) = \Delta\text{-type}(p)$ and $\Delta\text{-dim}(V) = \Delta\text{-dim}(p)$ where p is the Δ -generic type of V over K with respect to any embedding.

Let $a \in V$. A set Δ' of linearly independent elements of $\text{span}_{K\Delta} \Delta$ is said to *bound the Δ -type of a over K* if Δ' bounds the Δ -type of $f_i(a)$ over K for some i . Since the property “ Δ' bounds the Δ -type” is preserved under interdefinability over K , we get that Δ' bounds $\Delta\text{-type}(a/K)$ if and only if it bounds $\Delta\text{-type}(f_j(a)/K)$ for every j . We say that Δ' *bounds the Δ -type of V* if Δ' bounds the Δ -type of each $f_i(V_i)$. The notion of *witness* of the Δ -type is defined in the obvious way (see Definitions 1.6.7 and 1.6.9).

Given an abstract Δ -variety $(V, \{V_i\}, \{f_i\})$, we can construct the Δ -tangent bundle of V as an abstract Δ -variety. One patches together all the affine pieces $T_\Delta f_i(V_i)$ to form a Δ -variety where the transitions are given by the differentials of the f_{ij} 's. For a more detailed description of this type of constructions, we refer the reader to the end of Section 2.3 below, where we perform the construction of the *relative prolongation* of V .

Definition 1.7.3. An (abstract) Δ -algebraic group is a group object in the category of Δ -varieties. In other words, a Δ -algebraic group is a group whose underlying universe G is a Δ -variety and the group operation $p : G \times G \rightarrow G$ is a Δ -regular map. We say the Δ -algebraic group is defined over K , if G and p are over K .

Let us finish this section with a couple of remarks about Δ -algebraic groups.

Remark 1.7.4.

- (i) By Remark 1.6.10, every Δ -algebraic group has a witness to its Δ -type.
- (ii) Every Δ -algebraic group defined over K is isomorphic to a K -definable group in (\mathbb{U}, Δ) . Indeed, suppose $(G, \{V_i\}, \{f_i\})$ is a Δ -algebraic group over K . Let H be

the disjoint union of the $f_i(V_i)$ quotiented by the K -definable equivalence relation E induced by the transitions f_{ij} . By elimination of imaginaries (Fact 1.3.2 (iv)) for $DCF_{0,m}$, H is K -definable group which is isomorphic to G . Due to this fact, we think of Δ -algebraic groups as being definable sets in (\mathbb{U}, Δ) . Conversely, every K -definable group is definably isomorphic to a Δ -algebraic group over K . The proof of this requires much more effort and we refer the reader to [50] for details.

Chapter 2

Relative prolongations, D-varieties and D-groups

In this chapter we introduce the notions of *relative* prolongation, D-variety, and D-group. We present their basic properties and show that every differential algebraic variety which is not of maximal differential type is differentially birational to the sharp points of a relative D-variety. This material will be used in Chapters 3 and 4. We begin, however, with a review of this theory in the ordinary case.

2.1 A review of prolongations and D-varieties in DCF_0

In this section we briefly review the notions/properties of prolongations and D-varieties over an ordinary differential field. For more details (and proofs) on the material reviewed in this section we suggest [43], [32], [51], [48] and [29].

Fix a sufficiently large saturated $(\mathbb{U}, \delta) \models DCF_0$ and a small δ -subfield K . Recall that we identify all affine algebraic varieties with their set of \mathbb{U} -points and that we denote by \mathbb{U}^δ and K^δ the field of constants of \mathbb{U} and K , respectively.

Definition 2.1.1. Let $V \subseteq \mathbb{U}^n$ be an algebraic variety defined over K . The *prolongation* of V , denoted by τV , is the subvariety of \mathbb{U}^{2n} defined by

$$f(x) = 0 \quad \text{and} \quad \sum_{i=1}^n \frac{\partial f}{\partial x_i}(x) u_i + f^\delta(x) = 0 \quad (2.1.1)$$

for all f in $\mathcal{I}(V/K)$, where $x = (x_1, \dots, x_n)$, $u = (u_1, \dots, u_n)$. Here f^δ is the polynomial obtained by applying δ to the coefficients of f . The map $\pi : \tau V \rightarrow V$ denotes the projection onto the first n coordinates and, for each $a \in V$, τV_a denotes the fibre (with respect to π) of τV above a .

Remark 2.1.2. It is not hard to verify that in the definition of τV it suffices to consider the equations 2.1.1 as f ranges over any set of polynomials generating $\mathcal{I}(V/K)$, rather than all of $\mathcal{I}(V/K)$.

Note that if $a \in V$ then $(a, \delta a) \in \tau V$. This allows us to define a map $\nabla = \nabla_\delta^V : V \rightarrow \tau V$ given on points by $a \mapsto (a, \delta a)$. It is clear that ∇ is a δ -regular section of $\pi : \tau V \rightarrow V$; that is, $\pi \circ \nabla = \text{Id}_V$.

In the case that V is defined over the constants K^δ , inspecting the equations 2.1.1 and using Remark 2.1.2 above, we see that τV is simply the tangent bundle TV . In general, $\pi : \tau V \rightarrow V$ is a *torsor* under $\rho : TV \rightarrow V$; that is, there is a regular map $TV \times_V \tau V \rightarrow \tau V$ inducing a regular action (by translation) of TV_a on τV_a for all $a \in V$. For this reason the prolongation is often thought of as a *twisted* version of the tangent bundle.

The points of the prolongation also have the following differential algebraic interpretation.

Fact 2.1.3. *Suppose $a \in V$ is a generic point of V over K (in the sense of algebraic geometry). Then the points $b \in \tau V_a$ correspond to derivations $\delta' : K(a) \rightarrow \mathbb{U}$ that extend δ and $\delta'a = b$.*

One consequence of Fact 2.1.3, using also that $(\mathbb{U}, \delta) \models DCF_0$, is that if V and $W \subseteq \tau V$ are irreducible affine algebraic varieties, such that W projects dominantly onto V , then $\nabla(V)$ is Zariski-dense in W . In fact, this latter property characterizes the models of DCF_0 ; these are the Pierce-Pillay axioms which will be discussed in Section 4.1 below.

Let $f : V \rightarrow W$ be a regular map between affine algebraic varieties, where $f = (f_1, \dots, f_s)$. Then one can define a regular map $\tau f : \tau V \rightarrow \tau W$ given by

$$\tau f(x, u) = \left(f_1(x), \dots, f_s(x), \sum_{i=1}^n \frac{\partial f_1}{\partial x_i}(x) u_i + f_1^\delta(x), \dots, \sum_{i=1}^n \frac{\partial f_s}{\partial x_i}(x) u_i + f_s^\delta(x) \right).$$

One can in fact show that τ is a (covariant) functor from the category of affine varieties to itself.

Another way to think of τV is to identify it with the $\mathbb{U}[\epsilon]/(\epsilon)^2$ -points of $V \times_K R$, where R denotes the ring of dual numbers $K[\epsilon]/(\epsilon)^2$ equipped with the K -algebra structure

$e : K \rightarrow R$ given by $e(a) = a + \delta(a)\epsilon$. One can formalize this approach by defining τV as the Weil restriction of $V \times_K R$ from $K[\epsilon]/(\epsilon)^2$ to K . Note that the base change is with respect to the K -algebra structure $e : K \rightarrow R$ while the Weil restriction is with respect to the standard K -algebra structure of $K[\epsilon]/(\epsilon)^2$. This approach is taken and generalized by Moosa and Scanlon in the series of papers [39], [38] and [40].

More generally, one can define the prolongation of an abstract algebraic variety V defined over K by patching together the prolongations of an affine covering of V . This construction yields the prolongation of V as an abstract algebraic variety τV defined over K with a canonical projection $\pi : \tau V \rightarrow V$. Then, all the results of this section also hold for abstract algebraic varieties defined over K .

When studying finite dimensional δ -algebraic varieties it is convenient to work in the following geometric category introduced by Buium [1] (see also [48]).

Definition 2.1.4. An *algebraic D-variety* is a pair (V, s) where V is an (abstract) algebraic variety and $s : V \rightarrow \tau V$ is a regular section of $\pi : \tau V \rightarrow V$. The D-variety (V, s) is said to be defined over K if V and s are defined over K . The *set of sharp points* of (V, s) is defined as $(V, s)^\# = \{a \in V : s(a) = \nabla a\}$. A *morphism of D-varieties* (V, s) and (W, t) is a regular map $f : V \rightarrow W$ such that $\tau f \circ s = t \circ f$.

Fact 2.1.5.

- (a) *The set of sharp points of a D-variety (V, s) is a finite dimensional Zariski-dense subset of V .*
- (b) *If W is a finite dimensional irreducible δ -algebraic variety defined over K , then there exists an algebraic D-variety (V, s) defined over K such that W is δ -birationally equivalent (over K) to $(V, s)^\#$.*

Example 2.1.6. In the case when a finite dimensional δ -variety $W \subseteq \mathbb{U}$ is given by a single δ -polynomial equation of the form $\delta^n x = f(x, \delta x, \dots, \delta^{n-1}x)$ where $f \in K[t_0, \dots, t_{n-1}]$, it is easy to find a D-variety (V, s) such that W is δ -birationally equivalent to $(V, s)^\#$. Let $V = \mathbb{U}^n$ and $s(x_0, \dots, x_{n-1}) = (x_0, \dots, x_{n-1}, x_1, \dots, x_{n-1}, f(x_0, \dots, x_{n-1}))$. Then (V, s) is a D-variety, and the map from W to

$$(V, s)^\# = \{(a_0, \dots, a_{n-1}) \in V : \delta^n a_0 = f(a_0, \dots, a_{n-1}) \text{ and } \delta^i a_0 = a_i, i = 1, \dots, n-1\}$$

given by $x \mapsto (x, \delta x, \dots, \delta^{n-1}x)$ is a δ -birational equivalence (in fact a δ -isomorphism in this case).

If G is an algebraic group then τG has naturally the structure of an algebraic group: if $p : G \times G \rightarrow G$ is the group operation, then the group operation on τG is given by $\tau p : \tau G \times \tau G \rightarrow \tau G$ (this uses the fact that τ commutes with products). An *algebraic D-group* is a D-variety (G, s) where G is an algebraic group and $s : G \rightarrow \tau G$ is a group homomorphism. In the presence of a group structure, part (b) of Fact 2.1.5 can be refined to show that every finite dimensional definable group in DCF_0 is definably isomorphic (rather than just δ -birationally equivalent) to $(G, s)^\#$ for some algebraic D-group (G, s) .

The rest of this chapter is devoted to extend these notions/properties to the setting of several commuting derivations. Our approach will be to relativize the prolongation with respect to a given partition of the set of derivations. Hence, the role of algebraic varieties in the above discussion will be played by differential algebraic varieties involving a subset of the set of derivations, and the role of δ will be played by the remaining derivations. In particular, an analogue of Fact 2.1.5 will thus give us a geometric category in which to study possibly infinite dimensional differential algebraic varieties. The differential algebraic machinery (i.e., differential derivations) that we will require to accomplish this will be presented in the next section.

2.2 Differential derivations

If $K < F$ is an extension of fields and δ is a derivation on K , then one can always extend δ to a derivation on F . The situation is quite different if we start with a field K equipped with several commuting derivations Δ , a Δ -field extension F , and a distinguished additional derivation D on K commuting with Δ . If we want to extend D to a derivation on F such that the extension still commutes with Δ , then we need to find solutions to a certain system of Δ -equations which might not have a solution in F . However, as we will see, this system will always have a solution in some Δ -field extension of F . This latter fact was proven by Kolchin in Chapter 0.4 of [25]. In this section we give a brief review of this result.

Let (R, Δ) be differential ring and (S, Δ) a Δ -ring extension of R .

Definition 2.2.1. Suppose $D : R \rightarrow S$ is a derivation. Then D is called a Δ -*derivation* if D commutes with Δ on R .

Fix a Δ -derivation $D : R \rightarrow S$. Let $f \in R\{x\}_\Delta$, where $x = (x_1, \dots, x_n)$. By f^D we mean the Δ -polynomial in $S\{x\}_\Delta$ obtained by applying D to the coefficients of f .

Remark 2.2.2. The map $f \mapsto f^D$ is a Δ -derivation from $R\{x\}_\Delta$ to $S\{x\}_\Delta$. Indeed, clearly the map is additive and commutes with Δ , so we only need to check that it satisfies the Leibniz rule on monomials. Recalling that Θ_Δ is the set of derivative operators (see page 10), we need to consider $f(x) = c \prod_{\theta \in \Theta_\Delta, i \leq n} (\theta x_i)^{\alpha_{\theta,i}}$ and $g(x) = d \prod_{\theta \in \Theta_\Delta, i \leq n} (\theta x_i)^{\beta_{\theta,i}}$ where $c, d \in R$ and $\alpha_{\theta,i} = \beta_{\theta,i} = 0$ except finitely many times. Then

$$\begin{aligned} (fg)^D(x) &= D(cd) \prod_{\theta \in \Theta_\Delta, i \leq n} (\theta x)^{\alpha_{\theta,i} + \beta_{\theta,i}} \\ &= (D(c)d + cD(d)) \prod_{\theta \in \Theta_\Delta, i \leq n} (\theta x)^{\alpha_{\theta,i} + \beta_{\theta,i}} \\ &= f^D(x)g(x) + f(x)g^D(x). \end{aligned}$$

Lemma 2.2.3. *Suppose $a = (a_1, \dots, a_n)$ is a tuple from S and $D' : R\{a\}_\Delta \rightarrow S$ is a Δ -derivation extending D . Then, for all $f \in R\{x\}_\Delta$, we have*

$$D'f(a) = \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a) \theta D'a_i + f^D(a).$$

Proof. By additivity it suffices to prove this for monomials $f(x) = c \prod_{\theta \in \Theta_\Delta, i \leq n} (\theta x_i)^{\alpha_{\theta,i}}$ where $c \in R$ and $\alpha_{\theta,i} = 0$ except finitely many times,

$$\begin{aligned} D'f(a) &= \sum_{\eta \in \Theta_\Delta, k \leq n} \prod_{\theta \neq \eta, i \neq k} c(\theta a_i)^{\alpha_{\theta,i}} \alpha_{\eta,k} (\eta a_k)^{\alpha_{\eta,k}-1} D'\eta a_k + D(c) \prod_{\theta \in \Theta_\Delta, i \leq n} (\theta a_i)^{\alpha_{\theta,i}} \\ &= \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a) \theta D'a_i + f^D(a) \end{aligned}$$

where in the last equality we used that $D'\theta a_i = \theta D'a_i$ for all $\theta \in \Theta_\Delta$ and $i \leq n$. □

Due to the last property, it is natural to study the following operator.

Definition 2.2.4. Let $f \in R\{x\}_\Delta$. We define the Δ -polynomial $d_{D/\Delta}f \in S\{x, u\}_\Delta$ by

$$d_{D/\Delta}f(x, u) := \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(x) \theta u_i + f^D(x).$$

where $u = (u_1, \dots, u_n)$. For a tuple a from S , we will write $d_{D/\Delta}f_a(u)$ instead of $d_{D/\Delta}f(a, u)$.

Thus, Lemma 2.2.3 says that $D'f(a) = d_{D/\Delta}f(a, D'a)$. Note that, as only finitely many algebraic indeterminates θx_i appear in f , the above displayed sum is a finite sum and so indeed $d_{D/\Delta}f \in S\{x, u\}_\Delta$.

The map $f \mapsto d_{D/\Delta}f$ from $R\{x\}_\Delta \rightarrow S\{x, u\}_\Delta$ is clearly additive. Using Remark 2.2.2, it is not hard to verify that

$$d_{D/\Delta}(fg)(x, u) = d_{D/\Delta}f(x, u)g(x) + f(x)d_{D/\Delta}g(x, u). \quad (2.2.1)$$

Hence, if we identify $R\{x\}_\Delta$ as a Δ -subring of $S\{x, u\}_\Delta$, we get that $d_{D/\Delta} : R\{x\}_\Delta \rightarrow S\{x, u\}_\Delta$ is a derivation extending D . Moreover, as $d_{D/\Delta}\theta x_i = \theta u_i = \theta d_{D/\Delta}x_i$ for all $\theta \in \Theta_\Delta$ and $i \leq n$, $d_{D/\Delta}$ is in fact a Δ -derivation.

Note that, under the assumptions of Lemma 2.2.3, for all $f \in R\{x\}_\Delta$ we have $Df(a) = d_{D/\Delta}f_a(D'a)$, and so, if $f \in \mathcal{I}(a/R)_\Delta$, we get $d_{D/\Delta}f_a(D'a) = 0$. Thus, any Δ -derivation D' from $R\{a\}_\Delta$ to S extending D yields a tuple $D'a = (D'a_1, \dots, D'a_n)$ of S at which $d_{D/\Delta}f_a$ vanishes for all $f \in \mathcal{I}(a/R)_\Delta$. The following fact is the converse of this implication and gives the criterion for when a Δ -derivation can be extended to a finitely generated Δ -ring extension.

Fact 2.2.5. *Let a be a tuple from S . Suppose there is a tuple b from S such that for every $f \in \mathcal{I}(a/R)_\Delta$,*

$$d_{D/\Delta}f_a(b) = 0.$$

Then there is a unique Δ -derivation $D' : R\{a\}_\Delta \rightarrow S$ extending D such that $D'a = b$.

Thus if we want to extend D to a Δ -derivation from $R\{a\}_\Delta$ to S , we need to find a solution to the system of Δ -equations $\{d_{D/\Delta}f_a(u) = 0 : f \in \mathcal{I}(a/R)_\Delta\}$. In the case when S is a field (of characteristic zero), this system does have a solution in some Δ -field extension of S . Indeed, the Δ -ideal

$$[d_{D/\Delta}f_a(u) : f \in \mathcal{I}(a/R)_\Delta]$$

of $S\{u\}_\Delta$ is prime. From this and Fact 2.2.5 one obtains:

Fact 2.2.6. *Suppose (K, Δ) is a differentially closed field extending (R, Δ) and $D : R \rightarrow K$ a Δ -derivation. Then there is a Δ -derivation $D' : K \rightarrow K$ extending D .*

2.3 The relative prolongation of a differential algebraic variety

For the rest of this chapter we fix a sufficiently large saturated $(\mathbb{U}, \Pi) \models DCF_{0,m}$ and a small Π -subfield K . We assume a (disjoint) partition $\Pi = \mathcal{D} \cup \Delta$ with $\mathcal{D} = \{D_1, \dots, D_r\} \neq \emptyset$.

Note that \mathcal{D} is a set of commuting Δ -derivations.

Remark 2.3.1. Suppose Π' is a basis of the space of derivations on \mathbb{U} obtained by taking linear combinations of the elements of Π over the constant field K^Π . We denote this K^Π -vector space by $\text{span}_{K^\Pi} \Pi$. Then (\mathbb{U}, Π') is again a saturated model of $DCF_{0,m}$. Indeed, (\mathbb{U}, Π') is interdefinable with (\mathbb{U}, Π) over K^Π . Thus, if $\mathcal{D}' \cup \Delta'$ is a partition of Π' , then one can apply all the definitions and results of this chapter to that partition as well.

Let $D \in \mathcal{D}$. Recall, from the previous section, that $D : K \rightarrow K$ induces a Δ -derivation $d_{D/\Delta} : K\langle x \rangle_\Delta \rightarrow K\langle x, u \rangle_\Delta$ (see Definition 2.2.4). We can extend $d_{D/\Delta}$ to all Δ -rational functions over K using the quotient rule; that is, if $f = \frac{g}{h} \in K\langle x \rangle_\Delta$ then

$$d_{D/\Delta}f(x, u) = \frac{d_{D/\Delta}g(x, u)h(x) - g(x)d_{D/\Delta}h(x, u)}{(h(x))^2}.$$

Lemma 2.2.3 extends to all Δ -rational functions over K , and so, for any tuple a from \mathbb{U} and $f \in K\langle x \rangle_\Delta$, we have

$$Df(a) = d_{D/\Delta}f(a, Da).$$

If $f = (f_1, \dots, f_s)$ is a sequence of Δ -rational functions over K by $d_{D/\Delta}f$ we mean $(d_{D/\Delta}f_1, \dots, d_{D/\Delta}f_s)$. If a is an n -tuple, we write $d_{D/\Delta}f_a(u)$ instead of $d_{D/\Delta}f(a, u)$.

Remark 2.3.2. For the first part of this section we restrict ourselves to the category of affine Δ -varieties, since the presentation in the affine setting is more intuitive and concrete. Then, in the second part, we extend the constructions and concepts to the larger category of abstract Δ -varieties.

Definition 2.3.3. Let V be an affine Δ -variety defined over K . The *relative prolongation of V (w.r.t D/Δ)*, denoted by $\tau_{D/\Delta}V$, is the affine Δ -algebraic variety defined by

$$f(x) = 0 \text{ and } d_{D/\Delta}f(x, u) = 0,$$

for all f in $\mathcal{I}(V/K)_\Delta$. The map $\pi : \tau_{D/\Delta}V \rightarrow V$ denotes the projection onto the first n coordinates and, if a is a point in V , we denote the fibre of $\tau_{D/\Delta}V$ above a by $\tau_{D/\Delta}V_a$.

Remark 2.3.4. Note that we allow the possibility of $\Delta = \emptyset$. In this case V is an affine algebraic variety and the relative prolongation $\tau_{D/\emptyset}V$ is equal to the prolongation τV of Section 2.1 with respect to the ordinary differential field structure (K, D) .

It turns out that in the definition of $\tau_{D/\Delta}V$ we can restrict the polynomials f to a set that differentially generates $\mathcal{I}(V/K)_\Delta$ (it is worth pointing out that we might not be able to find a generating set whose elements have bounded order [26]).

Lemma 2.3.5. *Let $A \subset K\{x\}_\Delta$ be such that $[A] = \mathcal{I}(V/K)$. Then $\tau_{D/\Delta}V$ is defined by $f(x) = 0$ and $d_{D/\Delta}f(x, u) = 0$ for all $f \in A$.*

Proof. Let $a \in V$ and b such that $d_{D/\Delta}f_a(b) = 0$ for all $f \in A$, we need to show that $d_{D/\Delta}g_a(b) = 0$ for all $g \in \mathcal{I}(V/K)_\Delta$. Since $d_{D/\Delta}$ is a Δ -derivation, for each $\theta \in \Theta_\Delta$, $h \in K\{x\}_\Delta$ and $f \in A$, we have that

$$d_{D/\Delta}(h\theta f)_a(b) = d_{D/\Delta}h_a(b)\theta f(a) + h(a)\theta d_{D/\Delta}f_a(b).$$

By assumption $d_{D/\Delta}f_a(b) = 0$ and since $f \in A \subseteq \mathcal{I}(V/K)_\Delta$ we get $\theta f(a) = 0$. Hence, $d_{D/\Delta}(h\theta f)_a(b) = 0$, and so it follows that $d_{D/\Delta}g_a(b) = 0$ for all $g \in [A] = \mathcal{I}(V/K)_\Delta$. \square

The previous lemma implies that the relative prolongation is independent of the Δ -field over which V is defined.

Corollary 2.3.6. *Let F be a Δ -subfield of K . If V is also defined over F , then $\tau_{D/\Delta}V$ is defined by $f(x) = 0$ and $d_{D/\Delta}f(x, u) = 0$ for all $f \in \mathcal{I}(V/F)_\Delta$ (i.e., the defining Δ -ideal of V over F).*

Proof. As V is defined over F we have that $\mathcal{I}(V/K)_\Delta = \mathcal{I}(V/F)_\Delta \cdot K\{x\}_\Delta$. In particular, $\mathcal{I}(V/K)_\Delta = [\mathcal{I}(V/F)_\Delta]$. Now apply Lemma 2.3.5 with $A = \mathcal{I}(V/F)_\Delta$. \square

Lemma 2.3.7. *If V is defined over the D -constants K^D , then $\tau_{D/\Delta}V = T_\Delta V$. Here $T_\Delta V$ is Kolchin's Δ -tangent bundle (see Definition 1.5.7).*

Proof. First note that if $f \in K^D\{x\}_\Delta$ then

$$d_{D/\Delta}f(x, u) = \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(x)\theta u_i.$$

Hence, by Corollary 2.3.6, both $\tau_{D/\Delta}V$ and $T_\Delta V$ are defined by

$$\left\{ f(x) = 0 \text{ and } \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(x)\theta u_i = 0 : f \in \mathcal{I}(V/K^D)_\Delta \right\}.$$

\square

In general, when V is not necessarily defined over K^D , we have that each fibre $T_\Delta V_a$ acts regularly on $\tau_{D/\Delta}V_a$ by translation. In fact,

Lemma 2.3.8. $\pi : \tau_{D/\Delta}V \rightarrow V$ is a torsor under $\rho : T_\Delta V \rightarrow V$ where the fibrewise action is translation.

Proof. Let $a \in V$. For each $(a, b) \in T_\Delta V_a$ and $(a, c), (a, d) \in \tau_{D/\Delta}V_a$ we show that $(a, b+c) \in \tau_{D/\Delta}V_a$ and $(a, c-d) \in T_\Delta V_a$. Let $b = (b_1, \dots, b_n)$, $c = (c_1, \dots, c_n)$, $d = (d_1, \dots, d_n)$, and $f \in \mathcal{I}(V/K)_\Delta$, then

$$\begin{aligned} d_{D/\Delta}f(a, b+c) &= \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a)\theta(b_i + c_i) + f^D(a) \\ &= \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a)\theta b_i + \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a)\theta c_i + f^D(a) \\ &= 0 + d_{D/\Delta}f(a, c) = 0. \end{aligned}$$

Thus, $(a, b+c) \in \tau_{D/\Delta}V_a$. On the other hand,

$$\begin{aligned} \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a)\theta(c_i - d_i) &= \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a)\theta c_i - \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f}{\partial(\theta x_i)}(a)\theta d_i + f^D(a) \\ &= (d_{D/\Delta}f(a, c) - f^D(a)) - (d_{D/\Delta}f(a, d) - f^D(a)) \\ &= -f^D(a) + f^D(a) = 0. \end{aligned}$$

Thus, $(a, c-d) \in T_\Delta V_a$. These facts imply that the map $T_\Delta V_a \times \tau_{D/\Delta}V_a \rightarrow \tau_{D/\Delta}V_a$ given by $((a, b), (a, c)) \mapsto (a, b+c)$ is a well defined regular action. Hence, the regular map $T_\Delta V \times_V \tau_{D/\Delta}V \rightarrow \tau_{D/\Delta}V$ given on points by $((x, u), (x, v)) \mapsto (x, u+v)$ induces a regular action of $T_\Delta V$ on $\tau_{D/\Delta}V$ over V . \square

The following lemma shows that the fibres of the relative prolongation of any K -irreducible component of V are generically the same as the fibres of the relative prolongation of V .

Lemma 2.3.9. Let $\{V_1, \dots, V_s\}$ be the K -irreducible components of V . If $a \in V_i \setminus \bigcup_{j \neq i} V_j$, then $\tau_{D/\Delta}V_a = \tau_{D/\Delta}(V_i)_a$.

Proof. Clearly $\tau_{D/\Delta}(V_i)_a \subseteq \tau_{D/\Delta}V_a$. Let $b \in \tau_{D/\Delta}V_a$ and $f \in \mathcal{I}(V_i/K)_\Delta$. Since a is not in V_j , for $j \neq i$, we can pick a $g_j \in \mathcal{I}(V_j/K)_\Delta$ such that $g_j(a) \neq 0$. Then, if $g = \prod_j g_j$, we get $fg \in \mathcal{I}(V/K)_\Delta$ and so

$$0 = d_{D/\Delta}(fg)_a(b) = d_{D/\Delta}f_a(b)g(a) + f(a)d_{D/\Delta}g_a(b) = d_{D/\Delta}f_a(b)g(a)$$

where the third equality holds because $a \in V_i$. Since $g(a) \neq 0$, we have that $d_{D/\Delta}f_a(b) = 0$, and so $b \in \tau_{D/\Delta}(V_i)_a$. \square

Remark 2.3.10. The relative prolongation can be described in terms of local Δ -derivations along the lines of Kolchin's definition of the Δ -tangent bundle (see [25], Chap. 8, §2). We give a brief explanation. Let V be an affine Δ -variety defined over K with coordinate Δ -ring

$$K\{V\}_\Delta := K\{x\}_\Delta/\mathcal{I},$$

where $x = (x_1, \dots, x_n)$. Let $a \in V(F)$, where F is a Δ -extension of K . The *local Δ -ring of V at a over K* , $\mathcal{O}_a(V/K)_\Delta$, is just the localization of $K\{V\}_\Delta$ at the prime Δ -ideal $\mathcal{P} := \{f \in K\{V\}_\Delta : f(a) = 0\}$. Note that $\mathcal{O}_a(V/K)_\Delta$ is a Δ -ring extension of K . A *local Δ -derivation at a* is an additive map $\xi : \mathcal{O}_a(V/K)_\Delta \rightarrow F$ commuting with Δ such that $\xi(fg) = \xi(f)g(a) + f(a)\xi(f)$ for all $f, g \in \mathcal{O}_a(V/K)_\Delta$. If ξ is a local Δ -derivation at a extending D , then $(\xi(\bar{x}_1), \dots, \xi(\bar{x}_n))$ is a solution in F^n to the system $\{d_{D/\Delta}f_a(u) = 0 : f \in \mathcal{I}(V/K)_\Delta\}$, where \bar{x}_i is the image in $\mathcal{O}_a(V/K)_\Delta$ of $x_i + \mathcal{I}$. Conversely, every solution b in F^n to the system $\{d_{D/\Delta}f_a(u) = 0 : f \in \mathcal{I}(V/K)_\Delta\}$ gives rise to a local Δ -derivation at a extending D defined by

$$\xi_b(f) = d_{D/\Delta}f_a(b),$$

and it satisfies $(\xi_b(\bar{x}_1), \dots, \xi_b(\bar{x}_n)) = b$. Thus, the set of local Δ -derivations at a extending D can be identified with the set of solutions in F of the system $\{d_{D/\Delta}f_a(u) = 0 : f \in \mathcal{I}(V/K)_\Delta\}$. Hence, the F -points of $\tau_{D/\Delta}V_a$ can be identified with the set of all local Δ -derivations at a extending D .

Now, let us consider the whole set $\mathcal{D} = \{D_1, \dots, D_r\}$ at once, and for each affine Δ -variety V over K define

Definition 2.3.11. The *relative prolongation of V (w.r.t. \mathcal{D}/Δ)* is the fibred product:

$$\tau_{\mathcal{D}/\Delta}V := \tau_{D_1/\Delta}V \times_V \cdots \times_V \tau_{D_r/\Delta}V.$$

If V and W are affine Δ -varieties and $f : V \rightarrow W$ is a Δ -regular map (all defined over K), then we have an induced Δ -regular map, $\tau_{\mathcal{D}/\Delta}f : \tau_{\mathcal{D}/\Delta}V \rightarrow \tau_{\mathcal{D}/\Delta}W$ defined over K , given on points by

$$(x, u_1, \dots, u_s) \mapsto (f(x), d_{D_1/\Delta}f(x, u_1), \dots, d_{D_r/\Delta}f(x, u_r)). \quad (2.3.1)$$

Proposition 2.3.12. $\tau_{\mathcal{D}/\Delta}$ is a functor from the category of affine Δ -algebraic varieties (defined over K) to itself. Moreover, $\tau_{\mathcal{D}/\Delta}$ commutes with products; that is, $\tau_{\mathcal{D}/\Delta}(V \times W)$ is naturally isomorphic to $\tau_{\mathcal{D}/\Delta}V \times \tau_{\mathcal{D}/\Delta}W$ via the map

$$(x, y, u_1, \dots, u_r, v_1, \dots, v_r) \mapsto (x, u_1, \dots, u_r, y, v_1, \dots, v_r).$$

Proof. For ease of notation we assume $\mathcal{D} = \{D\}$. We need to show that for each pair of Δ -regular maps $f : V \rightarrow W$ and $g : W \rightarrow U$ we have that $\tau_{D/\Delta}(g \circ f) = \tau_{D/\Delta}g \circ \tau_{D/\Delta}f$. By the definition of $\tau_{D/\Delta}$ it suffices to show that $d_{D/\Delta}(g \circ f)(a, b) = d_{D/\Delta}g(f(a), d_{D/\Delta}f(a, b))$ for all $(a, b) \in \tau_{D/\Delta}V$. Write $f = (f_1, \dots, f_s)$. For each $\sigma \in \Theta_\Delta$ and $j \leq s$ we have that

$$\sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial \sigma f_j}{\partial \theta x_i}(a) \theta b_i = \sum_{\theta \in \Theta_\Delta, i \leq n} \sigma \left(\frac{\partial f_j}{\partial \theta x_i}(a) \theta b_i \right),$$

and hence

$$\sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial (g \circ f)}{\partial \theta x_i}(a) \theta b_i = \sum_{\sigma \in \Theta_\Delta, j \leq s} \frac{\partial g}{\partial \sigma y_j}(f(a)) \sigma \left(\sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f_j}{\partial \theta x_i}(a) \theta b_i \right).$$

From this we get

$$\begin{aligned} (g \circ f)^D(a) &= D(g(f(a))) - \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial (g \circ f)}{\partial \theta x_i}(a) \theta D a_i \\ &= \sum_{\sigma \in \Theta_\Delta, j \leq s} \frac{\partial g}{\partial \sigma y_j}(f(a)) \sigma(D f_j(a)) + g^D(f(a)) - \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial (g \circ f)}{\partial \theta x_i}(a) \theta a_i \\ &= \sum_{\sigma \in \Theta_\Delta, j \leq s} \frac{\partial g}{\partial \sigma y_j}(f(a)) \sigma(f_j^D(a)) + g^D(f(a)) = d_{D/\Delta}g(f(a), f^D(a)). \end{aligned}$$

Finally, using these equalities we have that

$$\begin{aligned} d_{D/\Delta}(g \circ f)(a, b) &= \sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial (g \circ f)}{\partial \theta x_i}(a) \theta b_i + (g \circ f)^D(a) \\ &= \sum_{\sigma \in \Theta_\Delta, j \leq s} \frac{\partial g}{\partial \sigma y_j}(f(a)) \sigma \left(\sum_{\theta \in \Theta_\Delta, i \leq n} \frac{\partial f_j}{\partial \theta x_i}(a) \theta b_i + f_j^D(a) \right) + g^D(f(a)) \\ &= d_{D/\Delta}g(f, d_{D/\Delta}f)(a, b). \end{aligned}$$

The moreover clause follows from the fact that $\mathcal{I}(V \times W/K)_\Delta$ equals the ideal generated by $\mathcal{I}(V/K)_\Delta \subseteq K\{x\}_\Delta$ and $\mathcal{I}(W/K)_\Delta \subseteq K\{y\}_\Delta$ in $K\{x, y\}_\Delta$. \square

Let V be an affine Δ -variety defined over K . Recall that for each tuple a from \mathbb{U} , $D \in \mathcal{D}$ and $f \in K\{x\}_\Delta$ we have that $Df(a) = d_{D/\Delta}f(a, Da)$, and so for each $a \in V$ and $g \in \mathcal{I}(V/K)_\Delta$ we have

$$d_{D/\Delta}g(a, Da) = Dg(a) = 0.$$

Hence, $(a, D_1a, \dots, D_r a) \in \tau_{\mathcal{D}/\Delta}V$ for all $a \in V$. This motivates the following definition.

Definition 2.3.13. For each affine Δ -variety V defined over K , we define the map $\nabla = \nabla_{\mathcal{D}}^V$ from V to $\tau_{\mathcal{D}/\Delta}V$ by

$$\nabla a = (a, D_1 a, \dots, D_r a).$$

Clearly ∇ is a Π -regular section of $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$, i.e., a Π -regular map such that $\pi \circ \nabla = \text{Id}_V$.

Remark 2.3.14. If $f : V \rightarrow W$ is a Δ -regular map between affine Δ -algebraic varieties (all defined over K), then $\nabla_{\mathcal{D}}^W \circ f = \tau_{\mathcal{D}/\Delta}f \circ \nabla_{\mathcal{D}}^V$. In other words, the following diagram commutes

$$\begin{array}{ccc} \tau_{\mathcal{D}/\Delta}V & \xrightarrow{\tau_{\mathcal{D}/\Delta}f} & \tau_{\mathcal{D}/\Delta}W \\ \nabla_{\mathcal{D}}^V \uparrow & & \uparrow \nabla_{\mathcal{D}}^W \\ V & \xrightarrow{f} & W \end{array}$$

Indeed, since for every tuple b and $D \in \mathcal{D}$ we have that $Df(b) = d_{D/\Delta}f(b, Db)$, for every $a \in V$

$$\nabla_{\mathcal{D}}^W(f(a)) = (a, d_{D/\Delta}f(\nabla_{\mathcal{D}}^V a)) = \tau_{\mathcal{D}/\Delta}f(\nabla_{\mathcal{D}}^V a).$$

Abstract relative prolongations. While the definitions and results of this section were stated for affine Δ -varieties (for the sake of concreteness), all of them make sense and hold for abstract Δ -varieties. In the remainder of this section we discuss how to make this precise by presenting in detail the construction of the *abstract relative prolongation* and the induced map ∇ .

First, for a quasi-affine Δ -variety U defined over K , set $\tau_{\mathcal{D}/\Delta}U$ to be the preimage of U under $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$ where V is the Δ -closure (in the Δ -topology) of U over K . If $f : U \rightarrow U'$ is Δ -regular map between quasi-affine Δ -varieties (defined over K), we define $\tau_{\mathcal{D}/\Delta}f : \tau_{\mathcal{D}/\Delta}U \rightarrow \tau_{\mathcal{D}/\Delta}U'$ by specifying it on points exactly as we did for the affine case (see (2.3.1) above).

Let $(V, \{V_i\}, \{f_i\})$ be an (abstract) Δ -variety defined over K (see Definition 1.7.1), where $\{V_i\}_{i=1}^s$ is an open cover and (V_i, f_i) are the local charts. Let $U_i = f_i(V_i)$ be the affine Δ -varieties defined over K that form the affine cover of V , $U_{ij} = f_i(V_i \cap V_j)$, and $f_{ij} = f_i \circ f_j^{-1} : U_{ji} \rightarrow U_{ij}$.

Definition 2.3.15. The *relative prolongation of V (w.r.t. \mathcal{D}/Δ)*, denoted by $\tau_{\mathcal{D}/\Delta}V$, is constructed as follows. Let X be the disjoint union of the $\tau_{\mathcal{D}/\Delta}(U_i)$'s and define on X an equivalence relation E where aEb if and only if $\tau_{\mathcal{D}/\Delta}f_{ij}(a) = b$ for some i and j . Then $\tau_{\mathcal{D}/\Delta}V$ is defined as X/E with the induced topology. We define the projection $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$ by $\pi(\alpha) = f_i^{-1}(\pi_i(a))$ where $a \in \tau_{\mathcal{D}/\Delta}U_i$ is a representative of the E -class α and $\pi_i : \tau_{\mathcal{D}/\Delta}U_i \rightarrow U_i$.

Let us check that the map $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$ is well defined. Suppose aEb , then $\tau_{\mathcal{D}/\Delta}f_{ij}(a) = b$, and applying the projection π_i we get $f_i \circ f_j^{-1}(\pi_j(a)) = \pi_i(b)$. Thus, $f_j^{-1}(\pi_j(a)) = f_i^{-1}(\pi_i(b))$, as desired.

A similar construction allows us to define the Δ -tangent bundle $\rho : T_{\Delta}V \rightarrow V$ as an abstract Δ -variety. From the affine case (Lemma 2.3.8), we can deduce that in the abstract category we still have that $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$ is a torsor under $\rho : T_{\Delta}V \rightarrow V$.

We can also define the Π -regular section $\nabla = \nabla_{\mathcal{D}}^V$ of $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$ by

$$\nabla(a) = \nabla_{\mathcal{D}}^{U_i}(f_i(a))/E, \quad \text{for } a \in V_i.$$

This is well defined. Indeed, suppose $a \in V_j$. Then

$$\tau_{\mathcal{D}/\Delta}f_{ji}(\nabla_{\mathcal{D}}^{U_i}f_i(a)) = \nabla_{\mathcal{D}}^{U_j}f_{ji}(f_i(a)) = \nabla_{\mathcal{D}}^{U_j}f_j(a),$$

where the first equality follows from Remark 2.3.14. This shows that $\nabla_{\mathcal{D}}^{U_i}f_i(a) E \nabla_{\mathcal{D}}^{U_j}f_j(a)$, as desired.

Given a Δ -regular map $f : V \rightarrow W$ between abstract Δ -varieties $(V, \{V_i\}, \{f_i\})$ and $(W, \{W_i\}, \{g_i\})$. We define the Δ -regular map $\tau_{\mathcal{D}/\Delta}f : \tau_{\mathcal{D}/\Delta}V \rightarrow \tau_{\mathcal{D}/\Delta}W$ by

$$\tau_{\mathcal{D}/\Delta}f(\alpha) = \tau_{\mathcal{D}/\Delta}(g_j \circ f \circ f_i^{-1})(a)/E_W, \quad \text{for all } \alpha \in \tau_{\mathcal{D}/\Delta}V,$$

where $a \in \tau_{\mathcal{D}/\Delta}U_i$ is a representative of the E_V -class of α and $f \circ f_i^{-1}(a) \in W_j$.

Now that we have the appropriate constructions, we can deduce, from the affine case, the results of this section for abstract Δ -varieties.

From now on, unless stated otherwise, we will work in the more general category of abstract Δ -varieties, which we simply call Δ -varieties.

2.4 Relative D-varieties

In this section we present the theory of relative D-varieties. Our main goal is to introduce a convenient category for studying the geometry of those differential algebraic varieties that are not of maximal differential type (see Proposition 2.4.6).

We continue to work in our universal domain (\mathbb{U}, Π) with a partition $\Pi = \mathcal{D} \cup \Delta$ where $\mathcal{D} = \{D_1, \dots, D_r\}$, and over a base Π -field $K < \mathbb{U}$.

Let us start with the definition of relative D-variety.

Definition 2.4.1. A *relative D-variety* (w.r.t. \mathcal{D}/Δ) defined over K is a pair (V, s) where V is a Δ -variety and s is a Δ -regular section of $\pi : \tau_{\mathcal{D}/\Delta}V \rightarrow V$ (both defined over K). We require the following *integrability condition* on s : for each $a \in V$,

$$d_{D_i/\Delta}s_j(a, s_i(a)) = d_{D_j/\Delta}s_i(a, s_j(a)) \quad \text{for } i, j = 1, \dots, r, \quad (2.4.1)$$

where $(\text{Id}, s_1, \dots, s_r) : U \rightarrow \tau_{\mathcal{D}/\Delta}U$ are coordinates of s in a local chart U containing a . By a *relative D-subvariety of* (V, s) we mean a relative D-variety (W, s_W) such that W is a Δ -subvariety of V and $s_W = s|_W$. We say (V, s) is an affine relative D-variety if V is an affine Δ -variety.

Note that if $\Delta = \emptyset$ we recover the definition of algebraic D-variety, see Section 2.1.

A morphism between relative D-varieties (V, s) and (W, t) is a Δ -regular map $f : V \rightarrow W$ satisfying $\tau_{\mathcal{D}/\Delta}f \circ s = t \circ f$; that is, the following diagram commutes

$$\begin{array}{ccc} \tau_{\mathcal{D}/\Delta}V & \xrightarrow{\tau_{\mathcal{D}/\Delta}f} & \tau_{\mathcal{D}/\Delta}W \\ s \uparrow & & \uparrow t \\ V & \xrightarrow{f} & W \end{array}$$

To every relative D-variety (V, s) defined over K we can associate a Π -variety (also defined over K) given by

$$(V, s)^\sharp = \{a \in V : s(a) = \nabla(a)\}.$$

A point in $(V, s)^\sharp$ is called a *sharp point* of (V, s) .

Example 2.4.2. In this example we illustrate the fact that in the absence of the integrability condition 2.4.1 there may not be any sharp point. Consider the case of two derivations $(\mathbb{U}, \Pi = \{D_1, D_2\})$. Let $V = \mathbb{U}$, $s_1(x) = x$ and $s_2(x) = x + 1$. Then

$$s = (\text{Id}, s_1, s_2) : \mathbb{U} \rightarrow \tau_{\Pi/\emptyset}V = \mathbb{U}^3$$

is a section of $\tau_{\Pi/\emptyset}V$. Note that s does not satisfy the integrability condition. Indeed,

$$d_{D_1/\emptyset}s_2(x, s_1(x)) = x \text{ and } d_{D_2/\emptyset}s_1(x, s_2(x)) = x + 1.$$

We now show that there is no point $a \in V$ such that $s(a) = \nabla(a)$. Towards a contradiction suppose $a \in V$ is such that $s(a) = \nabla(a)$. On the one hand we have

$$\delta_2\delta_1a = \delta_2(a) = a + 1.$$

On the other hand,

$$\delta_1 \delta_2 a = \delta_1(a + 1) = \delta_1 a = a.$$

Since $\delta_2 \delta_1 a = \delta_1 \delta_2 a$, we get $1 = 0$, and so we have a contradiction.

Proposition 2.4.3. *Let (V, s) be a relative D -variety defined over K .*

1. $(V, s)^\#$ is Δ -dense in V .
2. Assume V is K -irreducible (as a Δ -variety). There exists a Δ -generic point a of V over K such that $a \in (V, s)^\#$. Moreover, any such a is a Π -generic point of $(V, s)^\#$ over K . In particular, $(V, s)^\#$ will be K -irreducible.
3. Assume V is K -irreducible (as a Δ -variety). Let ω_V be the Kolchin polynomial of V as a Δ -algebraic variety and let $\omega_{(V, s)^\#}$ be the Kolchin polynomial of $(V, s)^\#$ as a Π -algebraic variety. Let μ be the smallest positive integer such that $s_i \in K(\theta x : \theta \in \Theta_\Delta(\mu))$, for all $i = 1, \dots, r$, whenever we write s in coordinates $s = (\text{Id}, s_1, \dots, s_r)$, in a local chart. Then for sufficiently large $h \in \mathbb{N}$

$$\omega_{(V, s)^\#}(h) \leq \omega_V(\mu h).$$

In particular, $\Delta\text{-type}(V) = \Pi\text{-type}(V, s)^\#$.

Proof. It is sufficient to consider the affine case.

(1) Let $F < \mathbb{U}$ be a Π -closed field extension of K and O be any (non-empty) Δ -open subset of V defined over F . We need to find a tuple c from F such that $c \in (V, s)^\# \cap O$. Let W be an irreducible component of V such that $O \cap W \neq \emptyset$ and let a be a Δ -generic point of W over F . Clearly $a \in O$.

Now let $f \in \mathcal{I}(a/F)_\Delta = \mathcal{I}(W/F)_\Delta$. Since a is a generic point of W over F , we can find g such that $g(a) \neq 0$ and $fg \in \mathcal{I}(V/F)_\Delta$. Thus, for $i = 1, \dots, r$,

$$\begin{aligned} 0 &= d_{D_i/\Delta}(fg)_a(s_i(a)) \\ &= d_{D_i/\Delta}f_a(s_i(a))g(a) + f(a)d_{D_i/\Delta}g_a(s_i(a)) \\ &= d_{D_i/\Delta}f_a(s_i(a))g(a) \end{aligned}$$

since $f(a) = 0$. But since $g(a) \neq 0$, we get $d_{D_i/\Delta}f_a(s_i(a)) = 0$ for $i = 1, \dots, r$. Fact 2.2.5 now implies that there are Δ -derivations $D'_i : F\{a\}_\Delta \rightarrow F\{a, s_i(a)\}_\Delta = F\{a\}_\Delta$ extending $D_i : F \rightarrow F$ such that $D'_i(a) = s_i(a)$, for $i = 1, \dots, r$.

Let us check that D'_i and D'_j commute on $F\{a\}_\Delta$; that is, for $f \in F\{x\}_\Delta$, $[D'_i, D'_j]f(a) = 0$. A rather lengthy computation shows that

$$[D'_i, D'_j]f(a) = d_{[D'_i, D'_j]/\Delta}f(a, [D'_i, D'_j]a).$$

Because D_i and D_j commute on F we get $f^{[D_i, D_j]} = 0$, so we only need to check that $[D'_i, D'_j]a = 0$. Here we use the integrability condition (2.4.1). We have

$$\begin{aligned} [D'_i, D'_j]a &= D'_i s_j(a) - D'_j s_i(a) \\ &= d_{D_i/\Delta} s_j(a, D'_i a) - d_{D_j/\Delta} s_i(a, D'_j a) \\ &= d_{D_i/\Delta} s_j(a, s_i(a)) - d_{D_j/\Delta} s_i(a, s_j(a)) = 0. \end{aligned}$$

Thus, the Δ -derivations $\mathcal{D}' = \{D'_1, \dots, D'_r\}$ commute on $F\{a\}_\Delta$. So, $F\{a\}_\Delta$ together with \mathcal{D}' , is a differential ring extension of (F, Π) and it has a tuple, namely a , living in $(V, s)^\# \cap O$. Hence, there is c in F such that $c \in (V, s)^\# \cap O$ as desired.

(2) Consider the set of formulas

$$\Phi = \{x \in (V, s)^\#\} \cup \{x \notin W : W \text{ is a proper } \Delta\text{-algebraic subvariety of } V \text{ over } K\}$$

if this set were inconsistent, by compactness and K -irreducibility of V , we would have that $(V, s)^\#$ is contained in a proper Δ -algebraic subvariety of V , but this is impossible by part (1). Hence Φ is consistent and a realisation is the desired point. The rest is clear.

(3) Let a be a Δ -generic point of V over K such that $a \in (V, s)^\#$ (this is possible by (2)). Then, since $\nabla(a) = s(a)$, for any $D \in \mathcal{D}$ we can write Da as a Δ -rational function from $K(\theta x : \theta \in \Theta_\Delta(\mu))$ evaluated at a . Thus, we get that for each $h \in \mathbb{N}$,

$$K(\theta a : \theta \in \Theta_\Pi(h)) \subseteq K(\theta a : \theta \in \Theta_\Delta(\mu h)).$$

□

Next we show that morphisms between relative D-varieties are precisely those regular Δ -maps preserving sharp points.

Lemma 2.4.4. *Let (V, s) and (W, t) be relative D-varieties defined over K . A regular Δ -map $f : V \rightarrow W$ is a morphism of relative D-varieties if and only if $f((V, s)^\#) \subseteq (W, t)^\#$.*

Proof. Suppose f is a morphism of relative D-varieties and let $a \in (V, s)^\#$. We need to show that $f(a) \in (W, t)^\#$, i.e., $t(f(a)) = \nabla_{\mathcal{D}}^W(f(a))$. By Remark 2.3.14, we have

$$\nabla_{\mathcal{D}}^W \circ f(a) = \tau_{\mathcal{D}/\Delta} f \circ \nabla_{\mathcal{D}}^V(a) = \tau_{\mathcal{D}/\Delta} f \circ s(a) = t \circ f(a),$$

as desired. Conversely, assume $f((V, s)^\#) \subseteq (W, t)^\#$. For all $a \in (V, s)^\#$ we have

$$\tau_{\mathcal{D}/\Delta}f(s(a)) = \tau_{\mathcal{D}/\Delta}f(\nabla_{\mathcal{D}}^V(a)) = \nabla_{\mathcal{D}}^W(f(a)) = t(f(a)).$$

Hence, $\tau_{\mathcal{D}/\Delta}f \circ s$ and $t \circ f$ agree on all of $(V, s)^\#$, and so, by (1) of Proposition 2.4.3, they agree on a Δ -dense subset of V . Thus, $\tau_{\mathcal{D}/\Delta}f \circ s = t \circ f$. \square

The previous lemma implies that we can define a (covariant) functor $\#$ from the category of relative D-varieties defined over K to the category of Π -varieties defined over K , which is defined on objects by $(V, s)^\#$ and on morphisms by $f|_{(V, s)^\#}$.

Proposition 2.4.5. *Let (V, s) be an relative D-variety defined over K . The $\#$ -functor establishes a bijective correspondence between relative D-subvarieties of (V, s) defined over K and Π -subvarieties of $(V, s)^\#$ defined over K . The inverse is given by taking the Δ -closure (in the Δ -topology of V) over K .*

Proof. If (W, s_W) is a relative D-subvariety of (V, s) defined over K then, by (1) of Proposition 2.4.3, the Δ -closure of $(W, s_W)^\#$ is just W . On the other hand, if X is a Π -algebraic subvariety of $(V, s)^\#$ defined over K and W is its Δ -closure over K , then $(W, s_W := s|_W)$ is clearly a \mathcal{D}/Δ -subvariety of (V, s) . We need to show $X = (W, s_W)^\#$. Towards a contradiction suppose $X \neq (W, s_W)^\#$. Using the equations of the sharp points $s(x) = \nabla_{\mathcal{D}}x$, we obtain a Δ -algebraic variety U defined over K such that $X \subseteq U$ and $(W, s_W)^\# \not\subseteq U$. This would imply that $(W, s_W)^\# \not\subseteq W$, but this is impossible. \square

So far in this section and the previous one we have been using a fixed partition $\mathcal{D} \cup \Delta$ of Π . However, as we mentioned in Remark 2.3.1 at the beginning of Section 2.3, all the definitions and results make sense and hold if we replace such a partition for any two sets \mathcal{D} and Δ of linearly independent elements of the K -vector space $\text{span}_{K^\Pi} \Pi$ such that their union $\mathcal{D} \cup \Delta$ forms a basis.

We now show that, up to Π -birational equivalence, every irreducible Π -variety of Π -type less than m is the sharp points of a relative D-variety (with an appropriate choice of \mathcal{D} and Δ such that $\mathcal{D} \cup \Delta$ is a basis of $\text{span}_{K^\Pi} \Pi$). This is the analogue in several derivations of the well known characterization of finite rank ordinary differential algebraic varieties (see §4 of [51]).

Proposition 2.4.6. *Let W be a K -irreducible Π -variety defined over K with $\Pi\text{-type}(W) = \ell < m$ and $\Pi\text{-dim}(W) = d$. Then there is a basis $\mathcal{D} \cup \Delta$ of the K^Π -vector space $\text{span}_{K^\Pi} \Pi$ with $|\Delta| = \ell$, and a relative D-variety (V, s) w.r.t. \mathcal{D}/Δ defined over K with $\Delta\text{-type}(V) = \ell$ and $\Delta\text{-dim}(V) = d$, such that W is Π -birationally equivalent (over K) to $(V, s)^\#$.*

Proof. It suffices to consider the affine case. Let a be a Π -generic point of W over K . Since the Π -type of W is ℓ , by Fact 1.6.6, there exists a set Δ of ℓ linearly independent elements of $\text{span}_{K^\Pi} \Pi$ such that $K\langle a \rangle_\Pi = K\langle \alpha \rangle_\Delta$ for some tuple α of \mathbb{U} with $\Delta\text{-tp}(\alpha/K) = \ell$ and $\Delta\text{-dim}(\alpha/K) = d$. Let $\mathcal{D} = \{D_1, \dots, D_r\}$ be any set of derivations such that $\mathcal{D} \cup \Delta$ is a basis of $\text{span}_{K^\Pi} \Pi$.

For each $i = 1, \dots, r$, we have

$$D_i \alpha = \frac{f_i(\alpha)}{g_i(\alpha)}$$

where $\frac{f_i}{g_i}$ is a sequence of Δ -rational functions over K . Let

$$b := \left(\alpha, \frac{1}{g_1(\alpha)}, \dots, \frac{1}{g_r(\alpha)} \right).$$

Note that b and α are Δ -interdefinable over K , hence $\Delta\text{-type}(b/K) = \Delta\text{-type}(\alpha/K)$ and $\Delta\text{-dim}(b/K) = \Delta\text{-dim}(\alpha/K)$. Also, a and b are Π -interdefinable over K and hence it suffices to show that b is a Π -generic point of the set of sharp points of a relative D-variety w.r.t. \mathcal{D}/Δ defined over K .

Let V be the Δ -locus of b over K (then $\Delta\text{-type}(V) = r$ and $\Delta\text{-dim}(V) = d$). A standard trick gives us a sequence $s = (\text{Id}, s_1, \dots, s_r)$ of Δ -polynomials over K such that $s(b) = \nabla_{\mathcal{D}} b \in \tau_{\mathcal{D}/\Delta} V$. Since b is a Δ -generic point of V over K , we get that (V, s) is a relative D-variety defined over K . Also, $b \in (V, s)^\#$ then, by Proposition 2.4.3 (3), b is a Π -generic point of $(V, s)^\#$ over K as desired. \square

The proof of Proposition 2.4.6 actually proves:

Corollary 2.4.7. *Let W be a K -irreducible Π -variety defined over K and Δ a set of fewer than m linearly independent elements of $\text{span}_{K^\Pi} \Pi$ that bounds $\Pi\text{-type}(W)$ (see Definition 1.6.9). If we extend Δ to a basis $\mathcal{D} \cup \Delta$ of $\text{span}_{K^\Pi} \Pi$, then W is Π -birationally equivalent (over K) to $(V, s)^\#$ for some relative D-variety (V, s) w.r.t. \mathcal{D}/Δ defined over K .*

2.5 Relative D-groups

In this section we develop the basic theory of relative D-groups which we will use in Section 3.3 to understand the generalized strongly normal extensions of Section 3.2 as Galois extensions of certain differential equations.

We continue to work in our universal domain (\mathbb{U}, Π) with a partition $\Pi = \mathcal{D} \cup \Delta$, and over a base Π -field $K < \mathbb{U}$. We first study the additional properties that $\tau_{\mathcal{D}/\Delta}G$ has in the case when G is a Δ -algebraic group.

Let $\Delta' \subseteq \Delta$ and suppose V is a Δ' -variety defined over K . It is not known (at least to the author) if $\tau_{\mathcal{D}/\Delta}V = \tau_{\mathcal{D}/\Delta'}V$. However, this equality holds in the case of Δ' -algebraic groups:

Proposition 2.5.1. *Let $\Delta' \subseteq \Delta$ and let G be a Δ' -algebraic group defined over K . Then $\tau_{\mathcal{D}/\Delta}G = \tau_{\mathcal{D}/\Delta'}G$.*

Proof. It suffices to consider the case when G is affine. For ease of notation we assume $\mathcal{D} = \{D\}$. By Lemma 2.3.9 we may assume that G is a connected Δ' -algebraic group. Let Λ be a characteristic set of the prime Δ' -ideal $\mathcal{I}(G/K)_{\Delta'}$. By Chapter 7 of [13], G is also a connected Δ -algebraic group and Λ is a characteristic set of $\mathcal{I}(G/K)_{\Delta}$. Let a be a Δ -generic point of G over K . We claim that $\tau_{D/\Delta}G_a = \tau_{D/\Delta'}G_a$. Let $b \in \tau_{D/\Delta'}G_a$, then $d_{D/\Delta}f_a(b) = d_{D/\Delta'}f_a(b) = 0$ for all $f \in \Lambda$. It is easy to see now that $d_{D/\Delta}f_a(b) = 0$ for all $f \in [\Lambda]_{\Delta}$, where $[\Lambda]_{\Delta}$ denotes the Δ -ideal generated by Λ in $K\{x\}_{\Delta}$. Now let $g \in \mathcal{I}(G/K)_{\Delta}$, since Λ is a characteristic set of $\mathcal{I}(G/K)_{\Delta}$ we can find ℓ such that $H_{\Lambda}^{\ell}g \in [\Lambda]_{\Delta}$, where H_{Λ} is the product of initials and separants of the elements of Λ (see Chap. IV, §9 of [24]). Thus we have

$$\begin{aligned} 0 &= d_{D/\Delta}(H_{\Lambda}^{\ell}g)_a(b) && \text{as } H_{\Lambda}^{\ell}g \in [\Lambda]_{\Delta} \\ &= H_{\Lambda}^{\ell}(a)d_{D/\Delta}g_a(b) + g(a)d_{D/\Delta}(H_{\Lambda}^{\ell})_a(b) \\ &= H_{\Lambda}^{\ell}(a)d_{D/\Delta}g_a(b). \end{aligned}$$

Since a is a Δ -generic point of G over K , $H_{\Lambda}(a) \neq 0$, and so $d_{D/\Delta}g_a(b) = 0$. Hence, $\tau_{D/\Delta'}G_a \subseteq \tau_{D/\Delta}G_a$. The other containment is clear.

Now we show that $\tau_{D/\Delta}G_g = \tau_{D/\Delta'}G_g$ for all $g \in G$. Let $g \in G$, we can find $h \in G$ such that $g = ha$. Let $\lambda^h : G \rightarrow G$ denote left multiplication by h , then $\tau_{D/\Delta}\lambda^h(\tau_{D/\Delta}G_a) = \tau_{D/\Delta}G_g$ and $\tau_{D/\Delta'}\lambda^h(\tau_{D/\Delta'}G_a) = \tau_{D/\Delta'}G_g$. But $\tau_{D/\Delta}\lambda^h = \tau_{D/\Delta'}\lambda^h$ as G is a Δ' -algebraic group, and so, since $\tau_{D/\Delta}G_a = \tau_{D/\Delta'}G_a$, we get $\tau_{D/\Delta}G_g = \tau_{D/\Delta'}G_g$. \square

Let G be a Δ -algebraic group defined over K . Then $\tau_{\mathcal{D}/\Delta}G$ has naturally the structure of a Δ -algebraic group defined over K ; more precisely, if $p : G \times G \rightarrow G$ is the group operation on G , then $\tau_{\mathcal{D}/\Delta}p : \tau_{\mathcal{D}/\Delta}G \times \tau_{\mathcal{D}/\Delta}G \rightarrow \tau_{\mathcal{D}/\Delta}G$ is a group operation on $\tau_{\mathcal{D}/\Delta}G$. Here we are identifying $\tau_{\mathcal{D}/\Delta}(G \times G)$ with $\tau_{\mathcal{D}/\Delta}G \times \tau_{\mathcal{D}/\Delta}G$ (via the natural isomorphism given in Proposition 2.3.12).

Remark 2.5.2. By Remark 2.3.14 we have

$$\tau_{\mathcal{D}/\Delta} p \circ \nabla_{\mathcal{D}}^{G \times G} = \nabla_{\mathcal{D}}^G \circ p.$$

Hence, the section $\nabla_{\mathcal{D}}^G : G \rightarrow \tau_{\mathcal{D}/\Delta} G$ is a group homomorphism.

Let us give some useful explicit formulas for the group law of $\tau_{\mathcal{D}/\Delta} G$. For ease of notation suppose G is affine. For each $f \in K\{x\}_{\Delta}$ let

$$d_{\Delta} f_x u := \sum_{\theta \in \Theta_{\Delta}, i \leq n} \frac{\partial f}{\partial \theta x_i}(x) \theta u_i,$$

and for a tuple $f = (f_1, \dots, f_s)$ let $d_{\Delta} f_x u$ be $(d_{\Delta}(f_1)_x u, \dots, d_{\Delta}(f_s)_x u)$. Suppose λ^g and ρ^g denote left and right multiplication by $g \in G$, respectively, and p is the group operation on G , then for (g, u_1, \dots, u_r) and (h, v_1, \dots, v_r) in $\tau_{\mathcal{D}/\Delta} G$ we have

$$(g, u_1, \dots, u_r) \cdot (h, v_1, \dots, v_r) = (g \cdot h, d_{\Delta}(\lambda^g)_h v_i + d_{\Delta}(\rho^h)_g u_i + p^{D_i}(g, h) : i \leq r).$$

The inverse is given by

$$(g, u_1, \dots, u_r)^{-1} = (g^{-1}, d_{\Delta}(\lambda^{g^{-1}} \circ \rho^{g^{-1}})_g (D_i g - u_i) + D_i(g^{-1}) : i \leq r).$$

It is clear, from these formulas, that $\tau_{\mathcal{D}/\Delta} G_e$ is a normal Δ -subgroup of $\tau_{\mathcal{D}/\Delta} G$, where e is the identity of G .

Now we show that at least for Δ -algebraic groups the functor $\tau_{\mathcal{D}/\Delta}$ preserves irreducibility.

Proposition 2.5.3. *If G is a connected Δ -algebraic group then $\tau_{\mathcal{D}/\Delta} G$ is also connected.*

Proof. Let H be the connected component of $\tau_{\mathcal{D}/\Delta} G$. First note that for any $g \in G$, since $\tau_{\mathcal{D}/\Delta} G_g$ is irreducible (as it is a translate of $T_{\Delta} G_g$) and the irreducible components of G are disjoint, if $H \cap \tau_{\mathcal{D}/\Delta} G_g \neq \emptyset$ then $\tau_{\mathcal{D}/\Delta} G_g \subseteq H$. Towards a contradiction suppose that $H \neq \tau_{\mathcal{D}/\Delta} G$. Then for some $g \in G$ the intersection $H \cap \tau_{\mathcal{D}/\Delta} G_g$ is empty. Hence, $\nabla^{-1}(H)$ and $\nabla^{-1}(\tau_{\mathcal{D}/\Delta} G \setminus H)$ are proper Π -algebraic subvarieties of G whose union is all of G . This implies that G is reducible as a Π -algebraic variety, but, by Theorem 7.2.3 of [13], G will also be reducible as a Δ -variety. This contradicts our hypothesis. \square

The following relativization of algebraic D-groups is obtained by considering the group objects in the category of relative D-varieties.

Definition 2.5.4. A *relative D-group* (w.r.t. \mathcal{D}/Δ) defined over K is a relative D-variety (G, s) such that G is a Δ -algebraic group and $s : G \rightarrow \tau_{\mathcal{D}/\Delta}G$ is a group homomorphism (all defined over K). A *relative D-subgroup* of (G, s) is a relative D-subvariety (H, s_H) of (G, s) such that H is a subgroup of G .

Remark 2.5.5.

1. If (G, s) is a relative D-group then $(G, s)^\#$ is a Π -subgroup of G . Indeed, if g and $h \in (G, s)^\#$ then

$$s(g \cdot h^{-1}) = s(g) \cdot s(h)^{-1} = (\nabla g) \cdot (\nabla h)^{-1} = \nabla(g \cdot h^{-1}).$$

Hence, $g \cdot h^{-1} \in (G, s)^\#$.

2. Suppose G is a Δ -algebraic group and (G, s) is a relative D-variety. If $(G, s)^\#$ is a subgroup of G then (G, s) is a relative D-group. Indeed, since $\nabla_{\mathcal{D}}^G$ is a group homomorphism, the restriction of s to $(G, s)^\#$ is a group homomorphism, and so, since $(G, s)^\#$ is Δ -dense in G and s is a regular Δ -map, s is a group homomorphism on all of G .

The correspondence of Proposition 2.4.5 specializes to a natural correspondence between relative D-subgroups of a relative D-group and Π -subgroups of its set of sharp points:

Lemma 2.5.6. *Suppose (G, s) is a relative D-group defined over K such that $(G, s)^\#(\bar{K}) = (G, s)^\#(K)$, for some Π -closure \bar{K} of K . Then the \sharp -functor establishes a bijective correspondence between relative D-subgroups of (G, s) defined over K and Π -algebraic subgroups of $(G, s)^\#$ defined over K . The inverse is given by taking the Δ -closure (in the Δ -topology of G) over K .*

Proof. This is an immediate consequence of Proposition 2.4.5 and the discussion above, except for the fact that the Δ -closure over K of a Π -algebraic subgroup of $(G, s)^\#$ is a subgroup of G . Let us prove this. Let H be a Π -algebraic subgroup of $(G, s)^\#$ and \bar{H} its closure over K . Let

$$X := \{a \in \bar{H} : b \cdot a \in \bar{H} \text{ and } b \cdot a^{-1} \in \bar{H} \text{ for all } b \in \bar{H}\} \subseteq \bar{H},$$

then X is a Δ -algebraic subgroup of G defined over K . We claim that $\bar{H} = X$. By the definition of \bar{H} , it suffices to show that $H(\bar{K}) \subseteq X(\bar{K})$. Let $a \in H(\bar{K}) = H(K)$, and consider $Y_a = \{x \in \bar{H} : x \cdot a \in \bar{H} \text{ and } x \cdot a^{-1} \in \bar{H}\}$. Then $H \subset Y_a$, but Y_a is Δ -closed and defined over K , so $\bar{H} \subset Y_a$. Thus $a \in X$, as desired. \square

We now show that every definable group G in (\mathbb{U}, Π) of differential type less than m is, after possibly replacing Π by some independent linear combination, definably isomorphic to a relative D-group w.r.t. \mathcal{D}/Δ with $|\Delta| = \Pi\text{-type}(G)$. By Remark 1.7.4 we can find a set Δ of linearly independent elements of $\text{span}_{K^\Pi} \Pi$ which witnesses the Π -type of G , and hence it suffices to prove the following.

Theorem 2.5.7. *Let G be a connected K -definable group. Suppose Δ is a set of fewer than m linearly independent elements of $\text{span}_{K^\Pi} \Pi$ that bounds the Π -type of G . If we extend Δ to a basis $\mathcal{D} \cup \Delta$ of $\text{span}_{K^\Pi} \Pi$, then G is K -definably isomorphic to $(H, s)^\#$ for some relative D-group (H, s) w.r.t. \mathcal{D}/Δ defined over K .*

Proof. Recall that any K -definable group is K -definably isomorphic to a Π -algebraic group defined over K (to the author's knowledge the proof of this fact, for the partial case, does not appear anywhere; however, it is well known that the proof for the ordinary case [50] extends with little modification). Thus, we assume that G is a Π -algebraic group defined over K . By Corollary 2.4.7, there is a Π -rational map α and a relative D-variety (V, s) w.r.t. \mathcal{D}/Δ , both defined over K , such that α yields a Π -birational equivalence between G and $(V, s)^\#$. Let p be the Π -generic type of G over K and q be the Δ -generic type of V over K .

We first show that there is a generically defined Δ -group structure on q . Note that α maps realisations of p to elements of $(V, s)^\#$ realising q . Let g and h be Π -independent realisations of p then

$$K\langle g, h \rangle_\Pi = K\langle \alpha(g), \alpha(h) \rangle_\Pi = K\langle \alpha(g), \alpha(h) \rangle_\Delta,$$

and so $\alpha(g \cdot h) \in K\langle \alpha(g), \alpha(h) \rangle_\Delta$. Thus, we can find a Δ -rational map ρ defined over K such that $\alpha(g \cdot h) = \rho(\alpha(g), \alpha(h))$. Hence, for any g, h Π -independent realisations of p

$$\alpha(g \cdot h) = \rho(\alpha(g), \alpha(h)),$$

and so $\rho(\alpha(g), \alpha(h))$ realises q . Now, since $\alpha(g)$ and $\alpha(h)$ are Δ -independent realisations of q , we get that for any x, y Δ -independent realisations of q , $\rho(x, y)$ realises q . Moreover, if g, h, l are Π -independent realisations of p then one easily checks that

$$\rho(\alpha(g), \rho(\alpha(h), \alpha(l))) = \rho(\rho(\alpha(g), \alpha(h)), \alpha(l)),$$

but $\alpha(g)$, $\alpha(h)$ and $\alpha(l)$ are Δ -independent realisations of q , so for any x, y, z Δ -independent realisations of q we get $\rho(x, \rho(y, z)) = \rho(\rho(x, y), z)$.

We thus have a stationary type q (in the language of Δ -rings) and a Δ -rational map ρ satisfying the conditions of Hrushovski's theorem on groups given generically (see [20] or

[53]), and so there is a connected K -definable group H and a K -definable injection β (both in the language of Δ -rings) such that β maps the realisations of q onto the realisations of the generic type of H over K and $\beta(\rho(x, y)) = \beta(x) \cdot \beta(y)$ for all x, y Δ -independent realisations of q . We assume, without loss of generality, that H is a Δ -algebraic group defined over K . We have a (partial) definable map $\gamma := \beta \circ \alpha : G \rightarrow H$ such that if g, h are Π -independent realisations of p then $\gamma(g \cdot h) = \gamma(g) \cdot \gamma(h)$. It follows that there is a K -definable group embedding $\bar{\gamma} : G \rightarrow H$ extending γ . Indeed, let

$$U := \{x \in G : \gamma(xy) = \gamma(x)\gamma(y) \text{ and } \gamma(yx) = \gamma(y)\gamma(x) \text{ for all } y \models p \text{ with } x \underset{K}{\perp} y\},$$

then U is a K -definable subset of G (by definability of types in $DCF_{0,m}$). If $g \models p$ then $g \in U$, and so every element of G is a product of elements of U (see §7.2 of [33]). Let $g \in G$ and let $u, v \in U$ be such that $g = u \cdot v$, then $\bar{\gamma}$ is defined by

$$\bar{\gamma}(g) = \gamma(u) \cdot \gamma(v).$$

It is well known that this construction yields a group embedding (see for example §3 of Marker's survey [31]). We also have a (partial) definable map $t := \tau_{\mathcal{D}/\Delta} \beta \circ s \circ \beta^{-1} : H \rightarrow \tau_{\mathcal{D}/\Delta} H$ such that for every $g \models p$, as $\alpha(g) \in (V, s)^\#$, we have $t(\gamma(g)) = \nabla_{\mathcal{D}}^H(\gamma(g))$. Thus, for g, h Π -independent realisations of p

$$t(\gamma(g) \cdot \gamma(h)) = \nabla_{\mathcal{D}}^H(\gamma(g) \cdot \gamma(h)) = (\nabla_{\mathcal{D}}^H \gamma(g)) \cdot (\nabla_{\mathcal{D}}^H \gamma(h)) = t(\gamma(g)) \cdot t(\gamma(h)).$$

Thus, for x, y Δ -independent realisations of the generic type of H over K we have $t(x \cdot y) = t(x) \cdot t(y)$, and so we can extend t to a K -definable group homomorphism $\bar{t} : H \rightarrow \tau_{\mathcal{D}/\Delta} H$. Therefore, (H, \bar{t}) is a relative D-group w.r.t. \mathcal{D}/Δ . Clearly, if g is a Π -generic of G over K , then $\gamma(g)$ is a Π -generic (over K) of both $\bar{\gamma}(G)$ and $(H, \bar{t})^\#$. Hence, $\bar{\gamma}(G) = (H, \bar{t})^\#$, as desired. \square

Chapter 3

Generalized Galois theory for partial differential fields

In this chapter we develop the fundamental results of generalized strongly normal extensions for partial differential fields in the possibly infinite dimensional setting. The connection to logarithmic differential equations on relative D-groups is also established. This theory generalizes simultaneously the parametrized Picard-Vessiot theory of Cassidy and Singer [7] and the finite dimensional differential Galois theory of Pillay [44]. We finish with two examples of non-linear Galois groups associated to logarithmic differential equations.

3.1 A review of Galois theory for ordinary differential fields

In this section we give a brief review of the several differential Galois theories for ordinary differential fields. We also mention some of the extensions of these theories for partial differential fields. We work in a sufficiently saturated $(\mathbb{U}, \delta) \models DCF_0$ and a base δ -subfield K of \mathbb{U} . Recall that we denote by \mathbb{U}^δ and K^δ the field of constants of \mathbb{U} and K , respectively.

Recall that, in usual Galois theory, the Galois extension associated to a polynomial is precisely the splitting field of the polynomial in question, and the Galois group is the group

The results in this chapter form the basis for a paper entitled “Relative D-groups and differential Galois theory in several derivations” that is to appear in *Transactions of the AMS*.

of automorphisms of the extension that leave the base field unchanged. It is then shown that the structure of the Galois group reflects the algebraic properties of the roots of the polynomial via the correspondence between its subgroups and the intermediate fields of the extension and the base field. One can proceed in an analogous fashion when dealing with homogeneous linear ordinary differential equations over K ; that is, equations of the form

$$a_n \delta^n x + a_{n-1} \delta^{n-1} x + \cdots + a_1 \delta x + a_0 x = 0, \quad a_i \in K.$$

One can easily rewrite each such equation in vector form

$$\delta y = My$$

where y is an $n \times 1$ column of indeterminates and M is an $n \times n$ matrix over K .

Definition 3.1.1. A *Picard-Vessiot extension* of K for the homogeneous linear differential equation $\delta y = My$ is a δ -field extension L of K such that

1. $L^\delta = K^\delta$ is algebraically closed.
2. L is generated by the entries of a matrix A in $\text{GL}_n(L)$ such that $\delta A = MA$.

The matrix A is called a *fundamental system of solutions* of the system $\delta y = My$. We let $\text{Aut}_\delta(L/K)$ denote the group of δ -automorphisms of L over K .

Let us mention that most authors do not require, in condition (1) of the definition of Picard-Vessiot extension, for K^δ to be algebraically closed, only that no new constants are added as we pass from K to L . We restrict ourselves to this case for the sake of a smoother exposition of the fundamental results of the classical theory.

Let us summarize the basic properties of Picard-Vessiot extensions:

- (i) Picard-Vessiot extensions always exist and are unique, up to δ -isomorphism over K .
- (ii) Let L be a Picard-Vessiot extension. Then, $\text{Aut}_\delta(L\langle \mathbb{U}^\delta \rangle_\delta / K\langle \mathbb{U}^\delta \rangle_\delta)$ is isomorphic to the \mathbb{U}^δ -points of a linear algebraic group \mathcal{G} defined over K^δ .
- (iii) There is a natural Galois correspondence between the algebraic subgroups of $\mathcal{G}(\mathbb{U}^\delta)$ defined over K^δ and the intermediate δ -fields of K and L .

For proofs and more properties of Picard-Vessiot extensions see [60] or [23].

One more fundamental property of a Picard-Vessiot extension L is that for any δ -isomorphism σ from L into \mathbb{U} over K , it is always the case that $\sigma(L) \subset L\langle\mathbb{U}^\delta\rangle_\delta$. Kolchin noticed that this implicit property (i.e., it does not refer to a differential equation) of Picard-Vessiot extensions is what yields a good Galois theory. Based on this observation, he was able to generalize the theory to a much larger class of extensions, including some that deal with extensions that arise from certain non-linear differential equations.

Definition 3.1.2. A finitely generated δ -field extension L of K is said to be *strongly normal* if:

1. $L^\delta = K^\delta$ is algebraically closed
2. For any δ -isomorphism σ from L into \mathbb{U} over K we have $\sigma(L) \subset L\langle\mathbb{U}^\delta\rangle_\delta$

In [24], Kolchin shows, amongst other things, that the group of automorphisms

$$\text{Aut}_\delta(L\langle\mathbb{U}^\delta\rangle_\delta/K\langle\mathbb{U}^\delta\rangle_\delta)$$

is isomorphic to the \mathbb{U}^δ -points of an algebraic (not necessarily linear) group defined over K^δ , and that there is a Galois correspondence analogous to (iii) above. It is worth mentioning that, in [54], Poizat observed that condition (2) in the definition of strongly normal can be rephrased in terms of the model theoretic notion of internality, this lead him to direct proofs of Kolchin's results using the machinery of ω -stable theories. Indeed, this is how model theorists became interested in differential Galois theory.

In the spirit of Poizat's observation, Pillay took differential Galois theory one step further. He successfully replaced the role of the constants \mathbb{U}^δ in the definition of strongly normal, by an arbitrary K -definable set X . He defined *X -strongly normal extensions* as those finitely generated δ -field extensions L of K such that $X(K) = X(\bar{L})$, for some δ -closure \bar{L} of L , and for any δ -isomorphism σ from L into \mathbb{U} over K we have $\sigma(L) \subset L\langle X\rangle_\delta$. Then, he proved, using model theoretic techniques, that the automorphism group

$$\text{Aut}_\delta(L\langle X\rangle_\delta/K\langle X\rangle_\delta)$$

is isomorphic to a finite dimensional δ -algebraic group defined over K (not necessarily in the constants), and that there is a Galois correspondence [45].

In [44], Pillay takes a different approach and reformulates his Galois theory in terms of logarithmic differential equations. His idea is to replace the role of GL_n in the definition of

Picard-Vessiot extensions by an arbitrary algebraic D-group (G, s) . In this situation one gets a surjective crossed homomorphism $\ell_s : G \rightarrow TG_e$, where TG_e is the Lie algebra of G (i.e., the tangent space at the identity e). The analogue of a linear equation is now an equation of the form $\ell_s x = \alpha$, where α is a K -point in the Lie algebra of G . Under some additional technical conditions, Pillay defines the notion of a differential Galois extension for the equation $\ell_s x = \alpha$, proves its existence and uniqueness, and observes that they are $(G, s)^\#$ -strongly normal extensions. In fact, he shows, under the additional assumption that K is algebraically closed, that every generalized strongly normal extension is of this form.

We will extend Pillay's theory to the partial differential, infinite dimensional, setting. Let us therefore mention some of the extensions to the partial setting that have already been considered in the literature. Suppose now that we are in a sufficiently saturated $(\mathbb{U}, \Pi) \models DCF_{0,m}$ and we have a partition $\mathcal{D} \cup \Delta$ of Π with $\mathcal{D} = \{D_1, \dots, D_r\}$. In [7], Cassidy and Singer extend the Picard-Vessiot theory to a parametrized Picard-Vessiot theory. For equations of the form

$$D_1 y = M_1 y, \dots, D_r x = M_r x,$$

where M_i are $n \times n$ matrices over a Π -subfield K satisfying the *integrability condition*

$$D_i M_j - D_j M_i = [M_i, M_j],$$

they define PPV-extensions as a Π -field extension L of K such that $L^\mathcal{D} = K^\mathcal{D}$ is Δ -closed and L is generated as a Π -field by the entries of a matrix A in $\text{GL}_n(L)$ such that $D_i A = M_i A$ for $i = 1, \dots, r$. They proved that the results from the Picard-Vessiot theory have natural extensions to this parametrized setting.

In the same spirit of PPV-extensions, Landesman parametrizes Kolchin's strongly normal extensions. He defines a Δ -strongly normal extension L as a finitely generated Π -field extension of K such that $L^\mathcal{D} = K^\mathcal{D}$ is Δ -closed and for any Π -isomorphism σ from L into \mathbb{U} over K we have $\sigma(L) \subset L\langle \mathbb{U}^\mathcal{D} \rangle_\Pi$. Then, he shows that the results for strongly normal extensions have natural extensions to this parametrized setting [30].

The goal of this chapter is to extend the differential Galois theories of Pillay (generalized strongly normal and logarithmic differential) to the partial setting. In fact both of his theories have more or less immediate extensions to the partial case as long as the definable sets in question are finite dimensional. Thus, our work is to extend Pillay's theories from finite dimensional to possibly infinite dimensional definable sets. As we will see, the PPV extensions of Cassidy and Singer and the Δ -strongly normal extensions of Landesman are special cases of what we present here.

3.2 Generalized strongly normal extensions

In this section we extend Pillay’s strongly normal extensions [45] to the possibly infinite dimensional partial differential setting. In fact most of the arguments here are extensions of Pillay’s arguments from [45], but for the sake of review, as well as completeness, we will give the details.

Fix a sufficiently large saturated $(\mathbb{U}, \Pi) \models DCF_{0,m}$ and a base Π -subfield K of \mathbb{U} . All definable sets will be identified with their \mathbb{U} -points. Unless otherwise specified, the notation and terminology of this section will be with respect to the derivations Π ; for example, $K\langle x \rangle$ means $K\langle x \rangle_\Pi$, generated means Π -generated, isomorphism means Π -isomorphism, etc.

Definition 3.2.1. Let X be a K -definable set. A finitely generated Π -field extension L of K is said to be *X -strongly normal* if:

1. $X(K) = X(\bar{L})$ for some (equivalently any) differential closure \bar{L} of L .
2. For any isomorphism σ from L into \mathbb{U} over K , $\sigma(L) \subseteq L\langle X \rangle$.

A Π -field extension L of K is called a *generalized strongly normal extension* of K if it is an X -strongly normal extension for some K -definable X .

Remark 3.2.2.

- (i) Suppose $\mathcal{D} \cup \Delta$ is a partition of Π with $\mathcal{D} \neq \emptyset$. If $X = \mathbb{U}^{\mathcal{D}}$, then the X -strongly normal extensions are exactly the “ Δ -strongly normal extensions” of Landesman [30]. Indeed, one need only observe that in this case (1) is equivalent to $K^{\mathcal{D}} = L^{\mathcal{D}}$ is Δ -closed (see for example Lemma 9.3 of [7]).
- (ii) If $L = K\langle a \rangle$ is an X -strongly normal extension and b realizes $tp(a/K)$, then $L' = K\langle b \rangle$ is also an X -strongly normal extension. Moreover, if $b \in \bar{L}$ then $L = L'$. Indeed, if σ is an automorphism of \mathbb{U} over K such that $\sigma(a) = b \in \bar{L}$, then, by (2), $L' \subseteq L\langle X \rangle$. But as $L' \subseteq \bar{L} \models DCF_{0,m}$, we have that $L' \subseteq L\langle X(\bar{L}) \rangle = L\langle X(K) \rangle = L$. A similar argument shows $L \subseteq L'$.
- (iii) When X is finite, X -strongly normal extensions are just the usual Galois (i.e., finite and normal) extensions of K . In the other extreme, if $\Pi\text{-type}(X) = m$ then X -strongly normal extensions are trivial. Indeed, suppose that $\Pi\text{-type}(X) = m$ and $X(K) = X(\bar{L})$. Then, since X having Π -type equal to m implies that X projects Π -dominantly onto some coordinate, we get a Π -dense K -definable subset of \mathbb{U} whose

\bar{L} -points are contained in K . This implies that $K = \bar{L}$, and hence X -strongly normal extensions are trivial. Therefore, this notion is mainly of interest when X is infinite and of Π -type strictly less than m .

For the rest of this section we fix a K -definable set X .

Proposition 3.2.3. *For any L finitely generated Π -field extension of K , condition (2) of Definition 3.2.1 is equivalent to L being generated by a fundamental system of solutions of an X -internal type over K .*

Proof. If a type $p \in S_n(K)$ is X -internal, then there is a sequence (a_1, \dots, a_ℓ) of realisations of p , called a fundamental system of solutions, such that for every $b \models p$ there is a tuple c from X such that $b \in \text{dcl}(K, a_1, \dots, a_\ell, c)$ (see §7.4 of [46]).

Suppose $L = K\langle a \rangle$ satisfies condition (2). We claim that $tp(a/K)$ is X -internal and a is itself a fundamental system of solutions. If $b \models tp(a/K)$ there is an automorphism σ of \mathbb{U} over K such that $b = \sigma(a) \in L\langle X \rangle$. Hence, $b \in \text{dcl}(K, a, c)$ for some tuple c from X . Conversely, suppose $L = K\langle a \rangle$ and a is a fundamental system of solutions of an X -internal type over K . Note that $tp(a/K)$ is also X -internal with fundamental system of solutions a itself. So if σ is an isomorphism from L into \mathbb{U} over K , then $\sigma(a) \models tp(a/K)$ and hence there is a tuple c from X such that $\sigma(a) \in \text{dcl}(K, a, c) = L\langle c \rangle$. That is $\sigma(L) \subseteq L\langle X \rangle$, as desired. \square

If $L = K\langle a \rangle$ is an X -strongly normal extension then, by Lemma 1.6.11 and Proposition 3.2.3,

$$\Pi\text{-type}(a/K) \leq \Pi\text{-type}(X).$$

Moreover, if $\Delta \subseteq \text{span}_{K^\Pi} \Pi$ bounds the Π -type of X then Δ also bounds the Π -type of a over K (see 1.6.9 for the definition of “bounds the Π -type”). Specializing to the case when $\mathcal{D} \cup \Delta$ is a partition of Π and $X = \mathbb{U}^\mathcal{D}$, we recover Theorem 1.24 of [30] that every $\mathbb{U}^\mathcal{D}$ -strongly normal extension of K is finitely Δ -generated over K .

We now start presenting the main properties of X -strongly normal extensions.

Lemma 3.2.4. *Suppose L is an X -strongly normal extension of K . Then L is contained in some differential closure of K .*

Proof. Let $L = K\langle a \rangle$ and $p = tp(a/K)$. Then for any realisation b of p , we have $b \in K\langle a, X \rangle$. By compactness, we get a K -definable function f and $\phi \in p$ such that for any b_1, b_2 realisations of ϕ there is a tuple c of X such that $b_1 = f(b_2, c)$. Let \bar{L} be a differential closure of L and \bar{K} a differential closure K contained in \bar{L} . Let $b \in \bar{K}$ satisfy ϕ . Then there is a tuple c from $X(\bar{L}) = X(K)$ such that $a = f(b, c)$. Thus, $a \in \bar{K}$ and $L < \bar{K}$. \square

The previous lemma implies that condition (1) of Definition 3.2.1 can be replaced by

$$(1') \quad X(K) = X(\bar{K}) \text{ and } L \subseteq \bar{K} \text{ for some differential closure } \bar{K} \text{ of } K.$$

The following result is fundamental for the analysis of the group of automorphisms of generalized strongly normal extensions.

Lemma 3.2.5. *Let L be an X -strongly normal extension of K . Let σ an isomorphism from L into \mathbb{U} over K . Then σ extends to a unique automorphism of $L\langle X \rangle$ over $K\langle X \rangle$.*

Proof. Let $L = K\langle a \rangle$ and ϕ a formula isolating $tp(a/K)$ (the existence of such a formula is given by Lemma 3.2.4). We show that ϕ also isolates $tp(a/K\langle X \rangle)$. If not, there is $\psi(x, y)$ (over K) and c a tuple from X such that $\psi(x, c) \in tp(a/K\langle X \rangle)$ and also a b realising $\phi(x)$ but not $\psi(x, c)$. Then we can find such b and c in \bar{L} , and so $c \in X(\bar{L}) = X(K)$ contradicting the fact that ϕ isolates $tp(a/K)$. This implies that $tp(a/K\langle X \rangle) = tp(\sigma(a)/K\langle X \rangle)$, and so there is a unique isomorphism from $L\langle X \rangle$ onto $\sigma(L)\langle X \rangle$ over $K\langle X \rangle$ extending σ . However, it follows from condition (2) of Definition 3.2.1 that $L\langle X \rangle = \sigma(L)\langle X \rangle$. \square

Now, let $L = K\langle a \rangle$ be an X -strongly normal extension of K . Following Pillay (and Kolchin), we define $\text{Gal}_X(L/K) := \text{Aut}_\Pi(L\langle X \rangle/K\langle X \rangle)$; that is, $\text{Gal}_X(L/K)$ is the group of automorphisms of $L\langle X \rangle$ over $K\langle X \rangle$. Note that $tp(a/K)^\mathbb{U}$ is in $\text{dcl}(L \cup X)$, and thus $\text{Gal}_X(L/K)$ acts naturally on $tp(a/K)^\mathbb{U}$. It follows, from Lemma 3.2.5, that this action is regular. We also let $\text{gal}(L/K) := \text{Aut}_\Pi(L/K)$. By Lemma 3.2.5, every element of $\text{gal}(L/K)$ extends uniquely to an element of $\text{Gal}_X(L/K)$. Via this map, we identify $\text{gal}(L/K)$ with a subgroup of $\text{Gal}_X(L/K)$.

Lemma 3.2.6. *$\text{Gal}_X(L/K)$ and its action on $tp(a/K)^\mathbb{U}$ do not depend (up to isomorphism) on the choice of X nor on the choice of generator a of L over K . We therefore simply write $\text{Gal}(L/K)$ for $\text{Gal}_X(L/K)$.*

Proof. If L is also an X' -strongly normal extension of K then, by Lemma 3.2.5, there is a (unique) group isomorphism

$$\psi : \text{Gal}_X(L/K) \rightarrow \text{Gal}_{X'}(L/K)$$

such that $\sigma|_{L\langle X \rangle \cap L\langle X' \rangle} = \psi(\sigma)|_{L\langle X \rangle \cap L\langle X' \rangle}$. Indeed, if $\sigma \in \text{Gal}_X(L/K)$, then the restriction of σ to L gives rise to an isomorphism from L into \mathbb{U} over K and so $\psi(\sigma)$ is defined as the unique extension of $\sigma|_L$ to an automorphism of $L\langle X' \rangle$ over $K\langle X' \rangle$. It is clear that this

yields the desired group isomorphism. Furthermore, if a' is another generator of L over K , then

$$tp(a/K)^\cup \cup tp(a'/K)^\cup \text{ is in } \text{dcl}(L \cup X \cup X')$$

and so any K -definable bijection between $tp(a/K)^\cup$ and $tp(a'/K)^\cup$, together with ψ , gives rise to an isomorphism between the natural actions of $\text{Gal}_X(L/K)$ on $tp(a/K)^\cup$ and $\text{Gal}_{X'}(L/K)$ on $tp(a'/K)^\cup$. \square

The following theorem asserts the existence of the *definable* Galois group of a generalized strongly normal extension and summarises some of its basic properties.

Theorem 3.2.7. *Suppose $L = K\langle a \rangle$ is an X -strongly normal extension of K .*

1. *There is a K -definable group \mathcal{G} in $\text{dcl}(K \cup X)$ with an L -definable regular action on $tp(a/K)^\cup$ such that \mathcal{G} together with its action on $tp(a/K)^\cup$ is (abstractly) isomorphic to $\text{Gal}(L/K)$ together with its natural action on $tp(a/K)^\cup$. We call \mathcal{G} the Galois group of L over K .*
2. *$\Pi\text{-type}(\mathcal{G}) = \Pi\text{-type}(a/K)$ and $\Pi\text{-dim}(\mathcal{G}) = \Pi\text{-dim}(a/K)$. Moreover, $\Delta \subseteq \text{span}_{K^\Pi} \Pi$ bounds the Π -type of \mathcal{G} if and only if Δ bounds the Π -type of a over K .*
3. *L is a \mathcal{G} -strongly normal extension.*
4. *The action of \mathcal{G} on $tp(a/K)^\cup$ is K -definable if and only if \mathcal{G} is abelian.*
5. *If $\mu : \text{Gal}(L/K) \rightarrow \mathcal{G}$ is the isomorphism from (1), then $\mu(\text{gal}(L/K)) = \mathcal{G}(\bar{K}) = \mathcal{G}(K)$.*

Proof.

(1) The construction is exactly as in [45], but we recall it briefly. Let $Z = tp(a/K)^\cup$, by Lemma 3.2.4, $tp(a/K)$ is isolated and so Z is K -definable. If $b \in Z$ then b is a tuple from $L\langle X \rangle$, so by compactness we can find a K -definable function $f_0(x, y)$ such that for any $b \in Z$ there is a tuple c from X with $b = f_0(a, c)$. Let $Y_0 = \{c \in X : f_0(a, c) \in Z\}$, then Y_0 is a K -definable set of tuples from X . Consider the equivalence relation on Y_0 given by $E(c_1, c_2)$ if and only if $f_0(a, c_1) = f_0(a, c_2)$, then E is K -definable. By elimination of imaginaries we can find a K -definable set Y in $\text{dcl}(K \cup X)$ which we can identify with Y_0/E . Now define $f : Z \times Y \rightarrow Z$ by $f(b, d) = f_0(b, c)$ where c is such that $d = c/E$. Now note that for any $b_1, b_2 \in Z$ there is a unique $d \in Y$ such that $b_2 = f(b_1, d)$, and so we can write $d = h(b_1, b_2)$ for some K -definable function h .

Define $\mu : \text{Gal}(L/K) \rightarrow Y$ by $\mu(\sigma) = h(a, \sigma a)$. Then μ is a bijection. Let \mathcal{G} denote the group with underlying set Y and with the group structure induced by μ . Then \mathcal{G} is a K -definable group and is in $\text{dcl}(K \cup X)$. Consider the action of \mathcal{G} on Z induced (via μ) from the action of $\text{Gal}(L/K)$ on Z , i.e., for each $g \in \mathcal{G}$ and $b \in Z$ let $g.b := \mu^{-1}(g)(b)$. This action is indeed L -definable since

$$\mu^{-1}(g)(b) = f(\mu^{-1}(g)(a), h(a, b)) = f(f(a, g), h(a, b)). \quad (3.2.1)$$

(2) Since \mathcal{G} acts regularly and definably on $tp(a/K)^\cup$, the map $g \mapsto g.a$ is a definable bijection between these two K -definable sets. Part (2) follows since Π -type, Π -dim and the property “ Δ bounds the Π -type” for definable sets are all preserved under extension of the set of parameters, as well as under definable bijection.

(3) Since \mathcal{G} is in $\text{dcl}(K \cup X)$ and $X(\bar{K}) = X(K)$, we have that $\mathcal{G}(\bar{K}) = \mathcal{G}(K)$. If σ is an isomorphism from $L = K\langle a \rangle$ into \mathbb{U} over K , then, by Lemma 3.2.5, we can extend it uniquely to an element of $\text{Gal}_X(L/K)$. Then, by (1), we get that $\sigma(a) = \mu(\sigma).a$ is in $\text{dcl}(L \cup \mathcal{G})$ and so $\sigma(L) \subseteq L\langle \mathcal{G} \rangle$.

(4) If the action is K -definable, for any $g_1, g_2 \in \mathcal{G}$ we get

$$(g_1 g_2).a = g_1.(\mu^{-1}(g_2)(a)) = \mu^{-1}(g_2)(g_1.a) = (g_2 g_1).a$$

so $g_1 g_2 = g_2 g_1$. Conversely, suppose \mathcal{G} is abelian and let $\sigma \in \text{Gal}(L/K)$ be such that $\sigma(a) = b$, then

$$\mu^{-1}(g)(b) = \mu^{-1}(g)(\sigma(a)) = \sigma(\mu^{-1}(g)(a)) = \sigma(f(a, g)) = f(b, g).$$

(5) Note that if $g \in \mathcal{G}(\bar{K})$ then $g = p(c)$ where $p \in K\langle x \rangle$ and c is a tuple from $X(\bar{K}) = X(K)$, and so $g \in \mathcal{G}(K)$. Thus $\mathcal{G}(\bar{K}) = \mathcal{G}(K) = \mathcal{G}(L)$. Now if $\sigma \in \text{gal}(L/K)$ then $\sigma(a)$ is a tuple from L , and $\mu(\sigma) = h(a, \sigma(a)) \in \mathcal{G}(L)$. Conversely, let $g \in \mathcal{G}(L)$ with $g = \mu(\sigma)$ and $\sigma \in \text{Gal}(L/K)$. Then $\sigma(a) = f(a, g)$ is in L and thus the restriction of σ to L determines an automorphism. Hence, $\sigma \in \text{gal}(L/K)$. \square

Remark 3.2.8. Using the notation of the proof of Theorem 3.2.7 (1), one could consider the map $Z \times \mathcal{G} \rightarrow Z$ given by $(b, g) \mapsto f(b, g) =: g * b$ and ask if this is a well-defined action. In general it is not, since it is not necessarily the case that $g_1 * (g_2 * b) = (g_1 g_2) * b$; however, it is the case that $g_1 * (g_2 * b) = (g_2 g_1) * b$. Hence, f induces a regular K -definable action of \mathcal{G}_{op} on Z , where \mathcal{G}_{op} is the group with universe \mathcal{G} and product $g \cdot h := h g$. Moreover, \mathcal{G} will be abelian if and only if $g * b$ defines an action of \mathcal{G} on Z , and in this case $g * b$ will coincide with the action of Theorem 3.2.7 (1).

Corollary 3.2.9. *A differential field extension L of K is a generalized strongly normal extension if and only if there are a K -definable group \mathcal{H} satisfying $\mathcal{H}(\bar{L}) = \mathcal{H}(K)$ and a K -definable principal homogeneous \mathcal{H} -space Z such that L is generated by an element of Z .*

Proof. Suppose $L = K\langle a \rangle$ is a generalized strongly normal extension. Let $Z = tp(a/K)^\cup$ and $\mathcal{H} = \mathcal{G}_{\text{op}}$ where \mathcal{G} is the K -definable group from Theorem 3.2.7. In Remark 3.2.8 we saw that the f from the proof of Theorem 3.2.7 (1) makes Z into a K -definable principal homogeneous \mathcal{H} -space. The converse is clear, since L will be an \mathcal{H} -strongly normal extension of K . \square

Now we explain why the construction of the definable group \mathcal{G} in the proof of Theorem 3.2.7 is independent of X, a as well as the choice of Y, f .

Lemma 3.2.10. *Let $L = K\langle a \rangle$ be an X -strongly normal extension and $\mu : \text{Gal}(L/K) \rightarrow \mathcal{G}$ as in Theorem 3.2.7. If \mathcal{G}' is a K -definable group satisfying the following:*

1. $\mathcal{G}'(\bar{K}) = \mathcal{G}'(K)$
2. *For some a' such that $L = K\langle a' \rangle$, there is an L -definable action of \mathcal{G}' on $Z' := tp(a'/K)^\cup$.*
3. *There is a group isomorphism $\mu' : \text{Gal}(L/K) \rightarrow \mathcal{G}'$ preserving the action on Z' .*

Then $\mu' \circ \mu^{-1} : \mathcal{G} \rightarrow \mathcal{G}'$ is the unique K -definable group isomorphism such that: any K -definable bijection between $Z := tp(a/K)^\cup$ and Z' , together with $\mu' \circ \mu^{-1}$, gives an isomorphism between the action of \mathcal{G} on Z and the action of \mathcal{G}' on Z' .

Proof. There is a K -definable function $q(x, y, z)$ such that $q(a', y, z) : Z' \times \mathcal{G}' \rightarrow Z'$ defines the action of \mathcal{G}' on Z' . Let $f'(y, z) := q(y, y, z)$. For any $b_1, b_2 \in Z'$ there is a unique $g' \in \mathcal{G}'$ such that $f'(b_1, g') = b_2$ and so we can write $g' = h'(b_1, b_2)$ for some K -definable function h' . We have that for any $\sigma \in \text{Gal}(L/K)$

$$\sigma(a') = q(a', a', \mu'(\sigma)) = f'(a', \mu'(\sigma)).$$

Thus, $\mu'(\sigma) = h'(a', \sigma a')$. Let $\rho : K\langle a \rangle \rightarrow K\langle a' \rangle$ be a K -definable bijection such that $\rho(a) = a'$. For every $g \in \mathcal{G}$ we get

$$\mu' \circ \mu^{-1}(g) = h'(a', \mu^{-1}(g)(a')) = h'(\rho(a), \rho(f(a, g))),$$

where f is the K -definable function from the proof of Theorem 3.2.7. For every $b \in Z$, we can find σ an automorphism of \mathbb{U} over K such that $\sigma(a) = b$ and σ fixes both \mathcal{G} and \mathcal{G}' pointwise (since L is a \mathcal{G} as well as a \mathcal{G}' -strongly normal extension). Then

$$h'(\rho(b), \rho(f(b, g))) = \sigma(h'(\rho(a), \rho(f(a, g)))) = h'(\rho(a), \rho(f(a, g))).$$

Hence, $\mu' \circ \mu^{-1}$ is K -definable. The rest is clear. \square

Even though the construction of \mathcal{G} in Theorem 3.2.7 depends on the choice of (X, a, Y, f) , if we choose different data (X', a', Y', f') , for the same extension L/K , and construct the corresponding \mathcal{G}' and μ' , Lemma 3.2.10 shows that \mathcal{G} and \mathcal{G}' will be K -definably isomorphic via $\mu' \circ \mu^{-1}$. Moreover, any K -definable bijection between $tp(a/K)^\mathbb{U}$ and $tp(a'/K)^\mathbb{U}$, together with $\mu' \circ \mu^{-1}$, gives rise to an isomorphism between the actions of \mathcal{G} on $tp(a/K)^\mathbb{U}$ and \mathcal{G}' on $tp(a'/K)^\mathbb{U}$. On the other hand, if \mathcal{G}' satisfies the conditions of Lemma 3.2.10 then it comes from such a construction (since in this case L will be a \mathcal{G}' -strongly normal extension and there is a natural choice of data which will give rise to \mathcal{G}'). Thus, to any generalized strongly normal extension we can associate a unique (up to K -definable isomorphism) K -definable group \mathcal{G} equipped with an L -definable regular action on $tp(a/K)^\mathbb{U}$ for each generator a of L over K . We call \mathcal{G} *the Galois group* of L over K .

Example 3.2.11. Suppose we have a partition $\mathcal{D} \cup \Delta$ of Π , $X = \mathbb{U}^\mathcal{D}$, L is an X -strongly normal extension of K and \mathcal{G} is its Galois group. Then \mathcal{G} is in $\text{del}(K \cup \mathbb{U}^\mathcal{D})$. Hence, \mathcal{G} is a Δ -algebraic group over $K^\mathcal{D}$ in the \mathcal{D} -constants. Thus, (1) and (2) of Theorem 3.2.7 recover Theorem 1.24 of [30].

We have the following Galois correspondence.

Theorem 3.2.12. *Let $L = K\langle a \rangle$ be an X -strongly normal extension of K with Galois group \mathcal{G} , and let $\mu : \text{Gal}(L/K) \rightarrow \mathcal{G}$ be the isomorphism from Theorem 3.2.7.*

1. *If $K \leq F \leq L$ is an intermediate Π -field, then L is an X -strongly normal extension of F . Moreover,*

$$\mathcal{G}_F := \mu(\text{Gal}(L/F))$$

is a K -definable subgroup of \mathcal{G} and is the Galois group of L over F . The map $F \mapsto \mathcal{G}_F$ establishes a 1-1 correspondence between the intermediate differential fields and K -definable subgroups of \mathcal{G} .

2. *F is an X -strongly normal extension of K if and only if \mathcal{G}_F is a normal subgroup of \mathcal{G} , in which case $\mathcal{G}/\mathcal{G}_F$ is the Galois group of F over K .*

Proof.

(1) Clearly L is an X -strongly normal extension of F . Also, F is finitely generated over K , in fact $F = K\langle b \rangle$ where b generates the minimal Π -field of definition of the set of realisations of $tp(a/F)$ (recall this type is isolated). Let $b = p(a)$ for $p \in K\langle x \rangle$. Note that if $g \in \mathcal{G}$ then $g \in \mathcal{G}_F$ if and only if $\mu^{-1}(g)(c) = c$ for all $c \in F$ which is equivalent to $\mu^{-1}(g)(p(a)) = p(a)$. Thus, $g \in \mathcal{G}_F$ if and only if $g \in \mathcal{G}$ and $p(f(a, g)) = p(a)$, which is a K -definable condition since $tp(a/K)$ is isolated. Thus \mathcal{G}_F is K -definable. It is clear, from the definition, that \mathcal{G}_F is the Galois group of L over F .

Now we prove the Galois correspondence. Let F_1 and F_2 be distinct intermediate differential fields, we show that if $b \in F_2 \setminus F_1$ there is $\sigma \in \text{Gal}(L/F_1)$ such that $\sigma(b) \neq b$. As $F_1 = \text{dcl}(F_1)$, there is an automorphism σ' of \mathbb{U} over F_1 such that $\sigma'(b) \neq b$. The restriction of σ' to $L\langle X \rangle$ yields the desired element of $\text{Gal}(L/F_1)$. Thus, $\mathcal{G}_{F_1} \neq \mathcal{G}_{F_2}$ and so the map $F \mapsto \mathcal{G}_F$ is injective.

For surjectivity let H be a K -definable subgroup of \mathcal{G} . Consider $Y = \{\mu^{-1}(h)(a) \in L\langle X \rangle : h \in H\}$ and let b be a tuple that generates the minimal Π -field of definition of Y . Then $b \in L$ and let $F = K\langle b \rangle$. We show $H = \mathcal{G}_F$. Let $h \in H$ then clearly $\mu^{-1}(h)(Y) = Y$, so $\mu^{-1}(h)(b) = b$ and then $h \in \mathcal{G}_F$. Conversely, if $g \in \mathcal{G}_F$ then, since $a \in Y$, $\mu^{-1}(g)(a) \in Y$. Thus, $\mu^{-1}(g) = \mu^{-1}(h)$ for some $h \in H$ and so $g = h \in H$. This establishes the Galois correspondence.

(2) Suppose F is an X -strongly normal extension of K , we need to show $\text{Gal}(L/F)$ is normal in $\text{Gal}(L/K)$. This is immediate since for all $\sigma \in \text{Gal}(L/K)$ we get that $\sigma(F) \subset F\langle X \rangle$. Now suppose \mathcal{G}_F is normal in \mathcal{G} , in order to show that F is an X -strongly normal extension of K we only need to check that for every isomorphism ϕ from F into \mathbb{U} over K , $\phi(F) \subset F\langle X \rangle$. For contradiction suppose there is $b \in F$ such that $\phi(b) \notin F\langle X \rangle$. We can extend ϕ to $\psi \in \text{Gal}(L/K)$ and thus by part (1) we can find $\sigma \in \text{Gal}(L/F)$ such that $\sigma \circ \psi(b) \neq \psi(b)$. Hence, $\psi^{-1} \circ \sigma \circ \psi(b) \neq b$ and this contradicts normality. Finally, the group isomorphism $\eta : \mathcal{G}/\mathcal{G}_F \rightarrow \text{Gal}(F/K)$ given by $\eta(g\mathcal{G}_F) = \mu^{-1}(g)|_{F\langle X \rangle}$ shows that $\mathcal{G}/\mathcal{G}_F$ is the Galois group of F over K . \square

Corollary 3.2.13. *Let L be a generalized strongly normal extension of K with Galois group \mathcal{G} . Then \mathcal{G} is connected if and only if K is relatively algebraically closed in L .*

Proof. Suppose $b \in L$ is algebraic over K but $b \notin K$. Let $F \leq L$ be the (algebraic) normal closure of $K(b)$, then F is an X -strongly normal extension of K . Thus, $\text{Gal}(F/K)$ is a nontrivial finite group and from (2) of Theorem 3.2.12 we get that $\text{Gal}(L/F)$ is a proper normal definable subgroup of $\text{Gal}(L/K)$ of finite index. This implies \mathcal{G} is not connected. Conversely, if \mathcal{G} is not connected, we can find a proper normal definable subgroup of finite

index. This subgroup will be of the form \mathcal{G}_F for some X -strongly normal extension F of K contained in L . But $\text{Gal}(F/K)$ will be finite and nontrivial, so F will be a proper algebraic extension of K in L . \square

Finally, as in the ordinary case, we have a positive answer to the “baby” inverse Galois problem:

Proposition 3.2.14. *Let G be a connected definable group with $\Pi\text{-type}(G) < m$. Then there is a Π -field K , over which G is defined, and a G -strongly normal extension L of K such that G is the Galois group of L over K .*

Proof. Let K_0 be a Π -closed field over which G is defined. By the embeddability property of differential algebraic groups into algebraic groups, there is a connected algebraic group H defined over K_0 and a K_0 -definable group embedding of G into H (to the author’s knowledge the proof of this fact, for the partial case, does not appear anywhere; however, it is well known that the proof for the ordinary case [50] extends with little modification). We therefore assume that G is a subgroup of the algebraic group H .

Let H/G be set of left cosets of G in H and $\nu : H \rightarrow H/G$ be the canonical projection. By elimination of imaginaries for $DCF_{0,m}$, we may assume H/G and ν are a definable set and a definable function, respectively, both over K_0 . Let a be a Π -generic point of H over K_0 . Then, since H is an algebraic group, $U(a/K_0) = \omega^m \cdot d$ where d is the (algebraic-geometric) dimension of H . Let $\alpha = \nu(a)$, $K = K_0\langle\alpha\rangle_\Pi$ and $L = K\langle a\rangle_\Pi$. Note that $L = K_0\langle a\rangle_\Pi$.

We first check that $G(\bar{L}) = G(K_0)$. Since K_0 is Π -closed, if $b \in \bar{L}$ and $b \downarrow_{K_0} L$ then $b \in K_0$ (this follows from definability of types in $DCF_{0,m}$). Thus it suffices to show that $g \downarrow_{K_0} a$ for all $g \in G$. Let $g \in G$, then $U(g/K_0) < \omega^m$ and by the Lascar inequalities

$$\omega^m \cdot d \leq U(a, g/K_0) \leq U(a/K_0, g) \oplus U(g/K_0).$$

Hence, $U(a/K_0, g) = \omega^m \cdot d$ and so $g \downarrow_{K_0} a$. Now suppose σ is an isomorphism from L into \mathbb{U} over K . Then $\nu(a) = \nu(\sigma a)$, and thus $a^{-1} \cdot \sigma a \in G$. Hence, $\sigma a \in L\langle G\rangle$, which implies that $\sigma(L) \subset L\langle G\rangle$. This shows that L is a G -strongly normal extension.

Now we check that the Galois group is G . It follows from the proof of Theorem 3.2.7 (i) that the Galois group is given by

$$\{g \in G : a \cdot g = \sigma(a) \text{ for some } \sigma \in \text{Gal}(L/K)\},$$

and thus it suffices to show that for each $g \in G$ we have that $tp(a \cdot g/K) = tp(a/K)$. Let $g \in G$, then we have that $\nu(a \cdot g) = \nu(a) = \alpha$. Since $\{x \in H : \nu x = \alpha\}$ is in definable bijection with G and the latter has Lascar rank less than ω^m , then $U(a/K) < \omega^m$ and $U(a \cdot g/K) < \omega^m$. Using Lascar inequalities again we get that $U(\alpha/K_0) = \omega^m \cdot d$ and also that $U(a \cdot g/K_0) = \omega^m \cdot d$. Then $a \cdot g$ is a generic point of H over K_0 , and so $tp(a \cdot g/K_0) = tp(a/K_0)$. But $\nu(a \cdot g) = \nu(a) = \alpha$, thus $tp(a \cdot g/K) = tp(a/K)$ as desired. \square

3.3 Relative logarithmic differential equations and their Galois extensions

In ordinary differential Galois theory, Picard-Vessiot extensions are Galois extensions corresponding to certain differential equations on linear algebraic groups in the constants (i.e., linear ODE's). Pillay's generalized strongly normal extensions correspond to certain differential equations on algebraic D-groups (under the assumption that the base field is algebraically closed). In this section we introduce the logarithmic derivative on a relative D-group and the Galois extensions associated to (relative) logarithmic differential equations. We then show, under appropriate assumptions, that these are precisely the generalized strongly normal extensions of Section 3.2. The theory we develop here extends Pillay's theory from [44].

We still continue to work in our universal domain $(\mathbb{U}, \Pi) \models DCF_{0,m}$ and a base Π -field $K < \mathbb{U}$. We assume a partition $\Pi = \mathcal{D} \cup \Delta$ with $\mathcal{D} = \{D_1, \dots, D_r\} \neq \emptyset$.

Definition 3.3.1. Let (G, s) be a relative D-group w.r.t. \mathcal{D}/Δ defined over K (c.f. Section 2.5) and e be its identity. We say that $\alpha \in \tau_{\mathcal{D}/\Delta}G_e$ is an *integrable point* of (G, s) if $(G, \alpha s)$ is a relative D-variety (but not necessarily a relative D-group). Here $\alpha s : G \rightarrow \tau_{\mathcal{D}/\Delta}G$ is the Δ -section given by $(\alpha s)(x) = \alpha \cdot s(x)$, where the latter product occurs in $\tau_{\mathcal{D}/\Delta}G$. Clearly the identity of $\tau_{\mathcal{D}/\Delta}G$ is integrable.

Example 3.3.2. Suppose G is a linear Δ -algebraic group defined over $K^{\mathcal{D}}$ (that is, a Δ -algebraic subgroup of GL_n for some n). Then, by Lemma 2.3.7, $\tau_{\mathcal{D}/\Delta}G$ is equal to $(T_{\Delta}G)^r$, the r -th fibred product of the Δ -tangent bundle of G , and so there is a zero section $s_0 : G \rightarrow (T_{\Delta}G)^r$. Then (G, s_0) is a relative D-group. In this case a point $(\text{Id}, A_1, \dots, A_r) \in \tau_{\mathcal{D}/\Delta}G_{\text{Id}} = (\mathcal{L}_{\Delta}(G))^r$, where $\mathcal{L}_{\Delta}(G) = T_{\Delta}G_e$ is the Δ -Lie algebra of G , is integrable if and only if it satisfies

$$D_i A_j - D_j A_i = [A_i, A_j], \quad \text{for } i, j = 1, \dots, r. \quad (3.3.1)$$

Indeed, by definition, the point $(\text{Id}, A_1, \dots, A_r)$ is integrable if and only if the section

$$(\text{Id}, A_1, \dots, A_r) \cdot (x, 0, \dots, 0) = (x, A_1x, \dots, A_rx)$$

satisfies the integrability condition (2.4.1). This means that for each i, j

$$d_{D_i/\Delta}(A_jx)(x, A_ix) = d_{D_j/\Delta}(A_ix)(x, A_jx),$$

which is equivalent to

$$A_jA_ix + D_i(A_j)x = A_iA_jx + D_j(A_i)x.$$

From this we get the desired equations. These equations (3.3.1) are the integrability conditions on A_1, \dots, A_r that one finds in [7] or [60].

Lemma 3.3.3. *A point $\alpha \in \tau_{\mathcal{D}/\Delta}G_e$ is integrable if and only if the system of Π -equations*

$$\nabla_{\mathcal{D}}x = \alpha s(x)$$

has a solution in G .

Proof. If α is integrable then, by Proposition 2.4.3 (1), $\{g \in G : \nabla_{\mathcal{D}}g = \alpha s(g)\}$ is Δ -dense in G . Conversely, suppose there is $g \in G$ such that $\nabla_{\mathcal{D}}g = \alpha s(g)$. To prove α is integrable it suffices to show that $(G, \alpha s)^\sharp$ is Δ -dense in G . This follows from $(G, \alpha s)^\sharp = g(G, s)^\sharp$ and the fact that $(G, s)^\sharp$ is Δ -dense in G . \square

Let (G, s) be a relative D-group defined over K . The *logarithmic derivative associated to (G, s)* is defined by

$$\begin{aligned} \ell_s : G &\rightarrow \tau_{\mathcal{D}/\Delta}G_e \\ g &\mapsto (\nabla_{\mathcal{D}}g) \cdot (s(g))^{-1} \end{aligned}$$

where the product and inverse occur in $\tau_{\mathcal{D}/\Delta}G$. Note that since ∇ and s are sections of $\pi : \tau_{\mathcal{D}/\Delta}G \rightarrow G$, the product $(\nabla_{\mathcal{D}}g) \cdot (s(g))^{-1}$ lies indeed in $\tau_{\mathcal{D}/\Delta}G_e$ (see the explicit formulas for the group law of $\tau_{\mathcal{D}/\Delta}G$ in Section 2.5).

We now list some properties of ℓ_s .

Lemma 3.3.4.

1. ℓ_s is a crossed-homomorphism, i.e., $\ell_s(gh) = (\ell_s g)(g * \ell_s h)$. Here $*$ is the adjoint action of G on $\tau_{\mathcal{D}/\Delta} G_e$, that is, $g * \alpha := \tau_{\mathcal{D}/\Delta} C_g(\alpha)$ for each $\alpha \in \tau_{\mathcal{D}/\Delta} G_e$ where C_g denotes conjugation by g .
2. The kernel of ℓ_s is $(G, s)^\#$.
3. The image of ℓ_s is exactly the set of integrable points of (G, s) .
4. For all $a \in G$, $tp(a/K \ell_s a)$ is $(G, s)^\#$ -internal and has a as a fundamental system of solutions.

Proof.

(1) This can be shown as in Pillay [44]. We include a sketch of the proof. An easy computation shows that $g * \alpha = u\alpha u^{-1}$ for any $u \in \tau_{\mathcal{D}/\Delta} G_g$. Thus, the adjoint action of G on $\tau_{\mathcal{D}/\Delta} G_e$ is by automorphisms and

$$\begin{aligned} \ell_s(gh) &= (\nabla_{\mathcal{D}} g)(\nabla_{\mathcal{D}} h)(s(h))^{-1}(s(g))^{-1} \\ &= (\ell_s g)(s(g))(\ell_s h)(s(g))^{-1} = (\ell_s g)(g * \ell_s h). \end{aligned}$$

(2) By definition of ℓ_s .

(3) Follows from Lemma 3.3.3, since if $\alpha = \ell_s(g)$ for some $g \in G$ then $\nabla_{\mathcal{D}} g = \alpha s(g)$.

(4) By basic properties of crossed-homomorphisms $\ell_s^{-1}(\ell_s a) = a \ker(\ell_s) = a(G, s)^\#$ for all $a \in G$. Thus, if $b \models tp(a/K \ell_s a)$ then $\ell_s b = \ell_s a$, and so $b \in \ell_s^{-1}(\ell_s a) = a(G, s)^\#$. \square

Extending the work of Pillay in [44], we point out that these relative logarithmic differential equations give rise to generalized strongly normal extensions.

Let (G, s) be a relative D-group defined over K with $(G, s)^\#(K) = (G, s)^\#(\bar{K})$, for some (equivalently any) Π -closure \bar{K} of K , and α be an integrable K -point of (G, s) . By Lemma 3.3.4 (3), the set of solutions in G to $\ell_s x = \alpha$ is nonempty. Hence there is a maximal Π -ideal $\mathcal{M} \subset K\{x\}_\Pi$ containing $\{\ell_s x - \alpha\} \cup \mathcal{I}(G/K)$. It is a well known fact that every maximal Π -ideal of $K\{x\}_\Pi$ is a prime ideal (see for example [21]). Let a be a tuple of \mathbb{U} such that $\mathcal{M} = \mathcal{I}(a/K)_\Pi$. Note that $tp(a/K)$ is therefore isolated (by the formula which sets the radical differential generators of $\mathcal{I}(a/K)_\Pi$ to zero) and so $K\langle a \rangle_\Pi$ is contained in a Π -closure of K . Moreover, Lemma 3.3.4 (4) tells us that, $tp(a/K)$ is $(G, s)^\#$ -internal and a is a fundamental system of solutions. Hence, by Proposition 3.2.3, $K\langle a \rangle_\Pi$ is a $(G, s)^\#$ -strongly normal extension of K .

Definition 3.3.5. We call the above $K\langle a \rangle_\Pi$ the *Galois extension associated to $\ell_s x = \alpha$* . The Galois group associated to this generalized strongly normal extension is called the *Galois group associated to $\ell_s x = \alpha$* .

Let us point out that the above construction does not depend on the choice of a (up to isomorphism over K). Indeed, if b is another solution such that $\mathcal{I}(b/K)_\Pi$ is a maximal Π -ideal, then both $tp(a/K)$ and $tp(b/K)$ are isolated and so we can find a Π -closure \bar{K} of K containing b and an embedding $\phi : K\{a\}_\Pi \rightarrow \bar{K}$ over K . Since α is a K -point, $\ell_s \phi(a) = \alpha$, thus, by Lemma 3.3.4 (1), $b^{-1}\phi(a) \in (G, s)^\#(\bar{K}) = (G, s)^\#(K)$. Hence, $\phi(a)$ and b are interdefinable over K and so $K\langle a \rangle_\Pi$ is isomorphic to $K\langle b \rangle_\Pi$ over K . This argument actually shows that there is exactly one such extension in each Π -closure of K .

Remark 3.3.6 (On the condition $(G, s)^\#(K) = (G, s)^\#(\bar{K})$).

- (i) Let G be a Δ -group defined over $K^\mathcal{D}$ and $s_0 : G \rightarrow \tau_{\mathcal{D}/\Delta}G = (T_\Delta G)^r$ its zero section. If $K^\mathcal{D}$ is Δ -closed then $(G, s_0)^\#(K) = (G, s_0)^\#(\bar{K})$. On the other hand, if $(G, s_0)^\#(K) = (G, s_0)^\#(\bar{K})$ and $\Delta\text{-type}(G) = |\Delta|$ then $K^\mathcal{D}$ is Δ -closed.
- (ii) Let (G, s) be a relative D-group defined over K . If $(G, s)^\#$ is the Galois group of a generalized strongly normal extension of K then $(G, s)^\#(K) = (G, s)^\#(\bar{K})$ (see (5) of Theorem 3.2.7).

Example 3.3.7 (The linear case). Suppose $G = GL_n$ and $s_0 : G \rightarrow \tau_{\mathcal{D}/\Delta}G = (T_\Delta G)^r$ is the zero section. By Proposition 2.5.1, $T_\Delta G = TG$ the (algebraic) tangent bundle of G , and so $\tau_{\mathcal{D}/\Delta}G_{\text{Id}} = \{\text{Id}\} \times (Mat_n)^r$. If $\alpha = (\text{Id}, A_1, \dots, A_r) \in \tau_{\mathcal{D}/\Delta}G_{\text{Id}}$, then the logarithmic differential equation $\ell_{s_0} x = \alpha$ reduces to the system of linear differential equations

$$D_1 x = A_1 x, \dots, D_r x = A_r x.$$

As we already pointed out in Example 3.3.2, α will an integrable point if and only if

$$D_i A_j - D_j A_i = [A_i, A_j] \text{ for } i, j = 1, \dots, r.$$

Also, in this case, $(G, s_0)^\#(K) = (G, s_0)^\#(\bar{K})$ if and only if $K^\mathcal{D}$ is Δ -closed. Thus, the Galois extensions of K associated to logarithmic differential equations of (GL_n, s_0) are precisely the parametrized Picard-Vessiot extensions considered by Cassidy and Singer in [7].

Proposition 3.3.8. *Let (G, s) be a relative D -group defined over K with $(G, s)^\sharp(K) = (G, s)^\sharp(\bar{K})$ and α be an integrable K -point of (G, s) . Let L be a Π -field extension of K generated by a solution to $\ell_s x = \alpha$. Then, L is the Galois extension associated to $\ell_s x = \alpha$ if and only if $(G, s)^\sharp(K) = (G, s)^\sharp(\bar{L})$ for some (any) Π -closure \bar{L} of L .*

Proof. By the above discussion, the Galois extension associated to $\ell_s x = \alpha$ is a $(G, s)^\sharp$ -strongly normal. Thus, $(G, s)^\sharp(K) = (G, s)^\sharp(\bar{L})$. For the converse, suppose $L = K\langle b \rangle_\Pi$ where b is a solution to $\ell_s x = \alpha$. Then, since $tp(b/K)$ is $(G, s)^\sharp$ -internal and b is a fundamental system of solutions, L is a $(G, s)^\sharp$ -strongly normal extension and so L is contained in a Π -closure \bar{K} of K . Let a be a tuple from \bar{K} such that $\mathcal{I}(a/K)_\Pi$ is a maximal Π -ideal. Then $K\langle a \rangle_\Pi$ is the Galois extension associated to $\ell_s x = \alpha$. But, by Lemma 3.3.4, $b^{-1}a \in (G, s)^\sharp(\bar{K}) = (G, s)^\sharp(K)$, and hence $L = K\langle a \rangle_\Pi$. \square

The proof of Lemma 3.9 of [44] extends directly to the partial case and yields the following proposition.

Proposition 3.3.9. *Let (G, s) and α be as in Proposition 3.3.8. Then the Galois group associated to $\ell_s x = \alpha$ is of the form $(H, s_H)^\sharp$ for some relative D -subgroup (H, s_H) of (G, s) defined over K . Moreover, if a is a solution to $\ell_s x = \alpha$ such that $\mathcal{I}(a/K)_\Pi$ is a maximal Π -ideal, then the action of $(H, s_H)^\sharp$ on $tp(a/K)^\cup$ is given by $h.b = (aha^{-1})b$.*

Proof. We give a slightly more direct argument than what is found in [44]. Let $L = K\langle a \rangle_\Pi$ be the Galois extension associated to $\ell_s x = \alpha$, where a is a solution to $\ell_s x = \alpha$ and $\mathcal{I}(a/K)_\Pi$ is a maximal Π -ideal. Let Z be the Π -locus of a over K , note that $Z = tp(a/K)^\cup$ and that Z is a Π -algebraic subvariety of $(G, \alpha s)^\sharp$. Let f be the multiplication on G and $Y = \{g \in (G, s)^\sharp : Zg = Z\}$. Following the construction of Theorem 3.2.7 (1) with the data $((G, s)^\sharp, a, Y, f)$, we get a bijection $\mu : \text{Gal}(L/K)_\Pi \rightarrow Y$ defined by $\mu(\sigma) = a^{-1}\sigma(a)$. It follows that μ is in fact a group isomorphism, where Y is viewed as a subgroup of $(G, s)^\sharp$. Let H be the Δ -closure of Y over K . Since Y is a Π -algebraic subgroup of $(G, s)^\sharp$ defined over K , H equipped with $s_H := s|_H$ is a \mathcal{D}/Δ -subgroup of (G, s) defined over K and $(H, s_H)^\sharp = Y$ (see Lemma 2.5.6).

The moreover clause follows by (3.2.1) in the proof of Theorem 3.2.7 (1). \square

The Galois correspondence given by Theorem 3.2.12 specializes to this context and, composed with the bijective correspondence between relative D -subgroups of a given relative D -group and the Π -algebraic subgroups of the sharp points given in Lemma 2.5.6, yields the following correspondence.

Corollary 3.3.10. *Let (G, s) and α be as in Proposition 3.3.8. Let L be the Galois extension associated to $\ell_s x = \alpha$ with Galois group $(H, s_H)^\#$. Then there is a Galois correspondence between the intermediate Π -fields (of K and L) and the relative D -subgroups of (H, s_H) defined over K .*

Now, let (G, s) and α be as in Proposition 3.3.8, and L be the Galois extension associated to $\ell_s x = \alpha$ with Galois group $(H, s_H)^\#$. Suppose that there is $\Delta' \subseteq \Delta$ such that G is a Δ' -algebraic group and s is a Δ' -section (recall that in this case $\tau_{\mathcal{D}/\Delta} G = \tau_{\mathcal{D}/\Delta'} G$). Let $\Pi' = \mathcal{D} \cup \Delta'$. We can consider the Galois extension L' and Galois group $(H', s_{H'})^\#$ associated to $\ell_s x = \alpha$ when the latter is viewed as a logarithmic Π' -equation (note that α is also an integrable point when (G, s) is viewed as a relative D -group w.r.t. \mathcal{D}/Δ'). In other words, L' is a Π' -field extension of K of the form $K\langle a \rangle_{\Pi'}$ where $\ell_s a = \alpha$ and $\mathcal{I}(a/K)_{\Pi'}$ is a maximal Π' -ideal of $K\{x\}_{\Pi'}$, and $(H', s_{H'})$ is a relative D -subgroup of (G, s) w.r.t. \mathcal{D}/Δ' such that $(H', s_{H'})^\#$ is (abstractly) isomorphic to the group of Π' -automorphisms $\text{Gal}(L'/K)_{\Pi'}$. We have the following relation between the Galois extensions L and L' , and the groups H and H' :

Proposition 3.3.11. *Let L, L', H and H' be as above. If $L = K\langle a \rangle_{\Pi}$ then $L' = K\langle a \rangle_{\Pi'}$, and H' equals the Δ' -closure (in the Δ' -Zariski topology) of H over K .*

Proof. Since $(G, s)^\#(K) = (G, s)^\#(\bar{L})$ for some Π -closure \bar{L} of L , then $(G, s)^\#(K) = (G, s)^\#(\bar{K}\langle a \rangle_{\Pi'})$ for some Π' -closure $\bar{K}\langle a \rangle_{\Pi'}$ of $K\langle a \rangle_{\Pi'}$. Now Proposition 3.3.8 implies that $K\langle a \rangle_{\Pi'}$ is the Galois extension associated to $\ell_s x = \alpha$ when viewed as a logarithmic Π' -equation, and so $L' = K\langle a \rangle_{\Pi'}$. Now, to show that H' is the Δ' -closure of H over K it suffices to show that $(H', s_{H'})^\#$ is the Π' -closure of $(H, s_H)^\#$ over K . First we check that $(H, s_H)^\# \subseteq (H', s_{H'})^\#$. Let $h \in (H, s_H)^\#$, then there is $\sigma \in \text{Gal}(L/K)_{\Pi}$ such that $h = a^{-1}\sigma(a)$. Since σ restricts to an element of $\text{Gal}(L'/K)_{\Pi'}$, $a^{-1}\sigma(a) \in (H', s_{H'})^\#$, showing the desired containment. Let Y be the Π' -closure of $(H, s_H)^\#$ over K , then adapting the proof of Lemma 2.5.6 one can see that Y is a Π' -algebraic subgroup of $(H', s_{H'})^\#$. The fixed field of $(H, s_H)^\#$ is K , then the fixed field of Y (as a Π' -subgroup of $(H', s_{H'})^\#$) is also K . Hence, by the Galois correspondence, $Y = (H', s_{H'})^\#$. \square

In the case when $(G, s) = (GL, s_0)$ and $\Delta' = \emptyset$, Proposition 3.3.11 specializes to Proposition 3.6 of [7].

We finish this section by showing that, under some natural assumptions on the differential field K , every generalized strongly normal extension of K is the Galois extension of a (relative) logarithmic differential equation. This is in analogy with Remark 3.8 of [44].

Theorem 3.3.12. *Let X be a K -definable set and L an X -strongly normal extension of K . Suppose Δ is a set of fewer than m linearly independent elements of $\text{span}_{K^\Pi} \Pi$ that bounds the Π -type of X , and such that K is Δ -closed. If we extend Δ to a basis $\mathcal{D} \cup \Delta$ of $\text{span}_{K^\Pi} \Pi$, then there is a connected relative D -group (H, s) w.r.t. \mathcal{D}/Δ defined over K and α an integrable K -point of (H, s) such that L is the Galois extension associated to $\ell_s x = \alpha$.*

Proof. We just need to check that the argument given in Proposition 3.4 (ii) of [45] for the finite-dimensional case extends to this setting. Let \mathcal{G} be the Galois group of L over K . Note that, by Theorem 3.2.7 (5), \mathcal{G} is connected and that, by Lemma 1.6.11 (2) and Theorem 3.2.7 (2), Δ also bounds the Π -type of \mathcal{G} . Thus, Theorem 2.5.7 implies that \mathcal{G} is of the form $(H, s)^\#$ for some relative D -group (H, s) w.r.t. \mathcal{D}/Δ defined over K .

Now, let b be a tuple such that $L = K\langle b \rangle_\Pi$ and let \bar{K} be a Π -closure of K that contains L . Let μ be the canonical isomorphism from $\text{Gal}(L/K)$ to $(H, s)^\#$. We know there is some K -definable function h such that $\mu(\sigma) = h(b, \sigma(b))$ for all $\sigma \in \text{Gal}(L/K)$. Consider the map $\nu : \text{Aut}_\Pi(\bar{K}/K) \rightarrow H(\bar{K})$ defined by $\nu(\sigma) = h(b, \sigma(b))$. Let $\sigma_i \in \text{Aut}_\Pi(\bar{K}/K)$ for $i = 1, 2$, and denote by σ'_i the unique elements of $\text{Gal}(L/K)$ such that $\sigma'_i(b) = \sigma_i(b)$. We have

$$\begin{aligned} \nu(\sigma_1 \circ \sigma_2) &= h(b, \sigma_1 \circ \sigma_2(b)) = h(b, \sigma'_1 \circ \sigma'_2(b)) \\ &= \mu(\sigma'_1 \circ \sigma'_2) = \mu(\sigma'_1) \mu(\sigma'_2) \\ &= h(b, \sigma'_1(b)) h(b, \sigma'_2(b)) = h(b, \sigma_1(b)) h(b, \sigma_2(b)) \\ &= \nu(\sigma_1) \nu(\sigma_2) = \nu(\sigma_1) \sigma_1(\nu(\sigma_2)) \end{aligned}$$

where the last equality follows from $(H, s)^\#(\bar{K}) = (H, s)^\#(K)$. In the terminology of ([25], Chap. 7) or [49] we say that ν is a definable cocycle from $\text{Aut}_\Pi(\bar{K}/K)$ to H . Using the fact that K is Δ -closed, the argument from Proposition 3.2 of [45] extends to show that every definable cocycle is cohomologous to the trivial cocycle. In other words, we get a tuple $a \in H(\bar{K})$ such that $\nu(\sigma) = a^{-1} \sigma(a)$ for all $\sigma \in \text{Aut}_\Pi(\bar{K}/K)$.

Claim. $K\langle a \rangle_\Pi = K\langle b \rangle_\Pi$.

Towards a contradiction suppose $a \notin K\langle b \rangle_\Pi$. Since $a \in \bar{K}$ and \bar{K} is also a Π -closure of $K\langle b \rangle_\Pi$ (see [47], Chap. 8), we get that $tp(a/K, b)$ is isolated. Thus we can find $c \in \bar{K}$ realising $tp(a/K, b)$ such that $c \neq a$. Then there is $\sigma \in \text{Aut}_\Pi(\bar{K}/K\langle b \rangle_\Pi)$ such that $\sigma(a) = c$ (see [47], Chap. 8), but this is impossible since σ fixes b iff $\nu(\sigma) = e$ (where e is the identity of H) iff $a^{-1} \sigma(a) = e$ iff σ fixes a . The other containment is analogous. This proves the claim.

By Proposition 3.3.8, all that is left to show is that $\alpha := l_s(a)$ is a K -point. Let $\sigma \in \text{Aut}_\Pi(\mathbb{U}/K)$, then $a^{-1}\sigma(a) \in (H, s)^\#$. Thus, $\nabla_{\mathcal{D}}(a^{-1}\sigma(a)) = s(a^{-1}\sigma(a))$ and so

$$\begin{aligned}\sigma(\alpha) &= \sigma((\nabla_{\mathcal{D}}a)(s(a))^{-1}) \\ &= (\nabla_{\mathcal{D}}\sigma(a))(s(\sigma(a)))^{-1} \\ &= (\nabla_{\mathcal{D}}a)(s(a))^{-1} \\ &= \alpha,\end{aligned}$$

as desired. □

3.4 Two examples

In this section we give two non-linear examples of Galois groups associated to logarithmic differential equations. Our examples are modeled after Pillay's non-linear example given in [44].

First we exhibit a finite-dimensional non-linear Galois group in two derivations:

Example 3.4.1. Let $\Pi = \{\delta_t, \delta_w\}$. Let $G = \mathbb{G}_m \times \mathbb{G}_a$ and $s : G \rightarrow \tau_{\Pi/\emptyset}G = (TG)^2$ be the (algebraic) section defined by $s(x, y) = (x, y, xy, 0, xy, 0)$. Here \mathbb{G}_m and \mathbb{G}_a denote the multiplicative and additive groups, respectively. Then (G, s) is a relative D-group w.r.t. Π/\emptyset , and the logarithmic derivative $\ell_s : G \rightarrow (TG)_{(1,0)}^2$ is given by

$$\ell_s(x, y) = \left(1, 0, \frac{\delta_t x}{x} - y, \delta_t y, \frac{\delta_w x}{x} - y, \delta_w y\right).$$

Thus,

$$(G, s)^\# = \{(x, y) \in G : \delta_t y = \delta_w y = 0 \text{ and } \delta_t x = \delta_w x = yx\}.$$

We take the ground Π -field to be $K := \mathbb{C}(t, e^{ct}, e^{cw} : c \in \mathbb{C})$, where we regard t and w as two complex variables, and the Π -field extension $L := K(w, e^{2wt+w^2})$. Then L is contained in a Π -closure $\bar{K} = \bar{\mathbb{C}}$ of K and \mathbb{C} .

We now show that $(G, s)^\#(\bar{K}) = (G, s)^\#(K)$. Let $(a, b) \in (G, s)^\#(\bar{K})$, then $b \in \bar{K}^\Pi = \mathbb{C}$ and

$$\delta_t \left(\frac{a}{e^{b(t+w)}} \right) = \delta_w \left(\frac{a}{e^{b(t+w)}} \right) = 0.$$

Thus $a = ce^{b(t+w)}$ for some $c \in \mathbb{C}$, and so $a \in K$.

Now, as L is generated by $(e^{2wt+w^2}, 2w)$ and this pair is a solution to

$$\begin{cases} \delta_t x = yx \\ \delta_t y = 0 \\ \delta_w x = (y + 2t)x \\ \delta_w y = 2 \end{cases} ,$$

L is the Galois extension associated to $\ell_s(x, y) = (1, 0, 0, 0, 2t, 2)$. Also, since the transcendence degree of L over K is 2 and $(G, s)^\#$ is a connected Π -algebraic group whose Kolchin polynomial is constant equal to 2, then the Galois group associated to $\ell_s(x, y) = (1, 0, 0, 0, 2t, 2)$ is $(G, s)^\#$.

Suppose $\mathcal{D} \cup \Delta$ is a partition of Π . Note that while Proposition 3.2.14 shows that for every connected relative D-group (G, s) the subgroup $(G, s)^\#$ is the Galois group of a generalized strongly normal extension, it is not known if $(G, s)^\#$ is the Galois group of a logarithmic differential equation. The following proposition gives a sufficient condition on G that allows a construction of a Galois extension of a logarithmic differential equation on (G, s) with Galois group $(G, s)^\#$.

Proposition 3.4.2. *Let (G, s) be a relative D-group and suppose G is a connected algebraic group. Then there is Π -field K and an integrable K -point α of (G, s) such that the Galois group associated to $\ell_s x = \alpha$ is $(G, s)^\#$.*

Proof. We follow the construction given in the proof of Proposition 3.2.14. Let K_0 be a Π -closed field over which the \mathcal{D}/Δ -group (G, s) is defined and let a be a Π -generic point of G over K_0 . Let $\alpha = \ell_s a$, $K = K_0 \langle \alpha \rangle_\Pi$ and $L = K \langle a \rangle_\Pi$. Using now the same arguments as in the proof of Proposition 3.2.14, we get that L is a $(G, s)^\#$ -strongly normal extension of K and $(G, s)^\#$ is the Galois group of L over K . Hence, by Proposition 3.3.8, L is the Galois extension of K associated to $\ell_s x = \alpha$ and $(G, s)^\#$ is the associated Galois group. \square

We finish with an example of an infinite-dimensional Galois group associated to a non-linear logarithmic differential equation.

Example 3.4.3. Let $\Pi = \{\delta_1, \delta_2\}$. Let $G = \mathbb{G}_m \times \mathbb{G}_a$ and $s : G \rightarrow \tau_{\delta_2/\delta_1} G = TG$ be the δ_1 -section given by

$$s(x, y) = (x, y, xy, \delta_1 y).$$

Then (G, s) is a relative D-group w.r.t. δ_2/δ_1 . The logarithmic derivative $\ell_s : G \rightarrow TG_e$ is given by

$$\ell_s(x, y) = \left(1, 0, \frac{\delta_2 x}{x} - y, \delta_2 y - \delta_1 y\right).$$

Thus the sharp points are given by

$$(G, s)^\# = \{(x, y) \in G : \delta_2 x = xy \text{ and } \delta_2 y = \delta_1 y\}.$$

Note that $\Pi\text{-type}(G, s)^\# = 1$ and $\Pi\text{-dim}(G, s)^\# = 2$. By Proposition 3.4.2, there is (α_1, α_2) such that $(G, s)^\#$ is the Galois group associated to the non-linear logarithmic differential equation

$$\begin{cases} \delta_2 x = x(y + \alpha_1) \\ \delta_2 y = \delta_1 y + \alpha_2 \end{cases} .$$

Chapter 4

Geometric axioms for $DCF_{0,m}$

In this chapter we establish three characterizations of the existentially closed differential fields with $\ell + 1$ commuting derivations in terms of the geometry of $DCF_{0,\ell}$. We treat the $(\ell + 1)^{\text{th}}$ -derivation as a differential derivation with respect to the first ℓ , and use relative prolongations to produce solutions to certain systems of differential equations. We observe that one of these characterizations is first order.

4.1 A review of the Pierce-Pillay axioms

In this section we briefly review the case of a single derivation and describe the geometric axiomatization given by Pierce and Pillay. Then we present an example, constructed by Hrushovski, showing that the naive extension of their axioms to the partial case is false.

Recall that in ordinary differential fields the prolongation functor is simply written as τ as opposed to $\tau_{\delta/\emptyset}$, and it goes from the category of algebraic varieties to itself.

Fact 4.1.1 (Pierce-Pillay axioms [43]). $(K, \delta) \models DCF_0$ if and only if

1. $K \models ACF_0$
2. Let V and $W \subseteq \tau V$ be irreducible affine algebraic varieties defined over K such that W projects dominantly onto V . Let O_V and O_W be nonempty Zariski-open subsets

The results in this chapter form the basis for a paper entitled “Geometric axioms for differentially closed fields with several commuting derivations” and its corrigendum that appear in the *Journal of Algebra*, Vol. 362 and Vol. 382, respectively

of V and W , respectively, defined over K . Then there is a K -point $a \in O_V$ such that $(a, \delta a) \in O_W$.

This geometric characterization of DCF_0 is indeed expressible in a first order way. First, one needs to check that irreducibility of a Zariski-closed set is a definable condition on the parameters. This means that if $f_1(u, x), \dots, f_s(u, x)$ is a set of polynomials over \mathbb{Q} in the variables $u = (u_1, \dots, u_r)$ and $x = (x_1, \dots, x_n)$, then there exists a formula $\phi(u)$ (in the language of rings) such that for any algebraically closed field K of characteristic zero and $a \in K^r$ we have that $K \models \phi(a)$ if and only if the Zariski-closed set

$$\{b \in K^n : f_1(a, b) = \dots = f_s(a, b) = 0\}$$

is irreducible. The latter is equivalent to asking for the ideal $(f_1(a, x), \dots, f_s(a, x))$ to be a prime ideal of $K[x]$. Hence, the existence of such a ϕ is a consequence of the well known bounds to check primality of ideals in polynomial rings in finitely many variables (see for example [59]).

Second, one needs to check that dominant projections of Zariski-closed sets onto irreducible Zariski-closed sets is a definable condition on the parameters. More precisely, if $f_{u,v} : W_v \rightarrow V_u$ is an algebraic family of polynomial maps between Zariski-closed sets over \mathbb{Q} , then there is a formula (in the language of rings) such that for any algebraically closed field of characteristic zero, K , and tuple (a, b) from K we have that $K \models \phi(a, b)$ if and only if V_a is irreducible and $f_{a,b}(W_b)$ is a Zariski-dense in V_a . Once we know that V_a is irreducible, the latter condition is equivalent to $f_{a,b}(W_b)$ having the same algebraic-geometric dimension as V_a . Hence, the existence of such a ϕ is a consequence of the fact that dimension is definable in ACF .

Third, one needs to know that the family of prolongations of a definable family of irreducible algebraic varieties is again a definable family. This means that if V_u is an algebraic family of Zariski-closed sets over \mathbb{Q} , then there is a δ -algebraic family W_u of Zariski closed sets over \mathbb{Q} such that for any algebraically closed differential field K (of characteristic zero) and tuple a from K we have that if V_a is irreducible then $W_a = \tau(V_a)$. To see this, suppose V_u is given by $f_1(u, x) = 0, \dots, f_s(u, x) = 0$, and let W_u be given by

$$f_i(u, x) = 0 \text{ and } \sum_{j=1}^n \frac{\partial f_i}{\partial x_j}(u, x)y_j + f_i(\delta u, x) = 0, \quad i = 1, \dots, s.$$

For arbitrary a , $W_a \neq \tau(V_a)$; however, if we knew that $(f_1(a, x), \dots, f_s(a, x)) = \mathcal{I}(V/K)$ then we would get equality. Hence, for those a such that $(f_1(a, x), \dots, f_s(a, x))$ is prime we get $W_a = \tau(V_a)$. But we know primality is a definable condition on the parameters.

Now, if one tries to extend these axioms to the partial case $\Delta = \{\delta_1, \dots, \delta_m\}$ by simply replacing in condition (2) the prolongation functor τ for $\tau_\Delta := \tau_{\Delta/\emptyset}$ and $(a, \delta a)$ for $(a, \delta_1 a, \dots, \delta_m a)$, then the corresponding statement will not be true for the existentially closed models as the following example of Hrushovski's shows (one can also see this using Example 2.4.2):

Example 4.1.2 ([41], Counterexample 3.2). Let $\Delta = \{\delta_1, \delta_2\}$. Let $V = \mathbb{A}^3$ and $W \subset \tau_\Delta V = \mathbb{A}^9$ be defined by

$$W = \{(x_0, y_0, z_0, x_1, y_1, z_1) : x_1 = z_0^2, x_2 = z_0, y_1 = 2x_0, y_2 = z_0\}.$$

Note that W projects dominantly onto V . Let $O_V \subset V$ be the Zariski-open set defined by $z_0 \neq 0$. We claim that there is no Δ -field K such that there exists a K -point $(a, b, c) \in O_V$ with

$$(a, b, c, \delta_1 a, \delta_1 b, \delta_1 c, \delta_2 a, \delta_2 b, \delta_2 c) \in W.$$

Suppose there is such a Δ -field K and K -point (a, b, c) . Then we compute,

$$\begin{aligned} 2c\delta_2 c &= \delta_2(c^2) \\ &= \delta_2\delta_1 a && \text{as } \delta_1 a = c^2 \text{ is implied by } W \\ &= \delta_1\delta_2 a \\ &= \delta_1 c && \text{as } \delta_2 a = c \text{ is implied by } W \\ &= \delta_1\delta_2 b && \text{as } \delta_2 b = c \text{ is implied by } W \\ &= \delta_2\delta_1 b \\ &= \delta_2(2a) && \text{as } \delta_1 b = 2a \text{ is implied by } W \\ &= 2\delta_2 a \\ &= 2c. \end{aligned}$$

Since $c \in O_V$, $c \neq 0$, we have that $\delta_2 c = 1$. Hence,

$$\begin{aligned} 2 &= 2\delta_2 c \\ &= \delta_2(2c) \\ &= \delta_2\delta_1 c && \text{as } \delta_1 c = 2c \\ &= \delta_1\delta_2 c \\ &= \delta_1 1 \\ &= 0, \end{aligned}$$

which is the desired contradiction.

In the following section we take a different approach to the problem of geometric axioms for partial differentially closed fields. We establish a characterization (and ultimately an axiomatization) of $DCF_{0,\ell+1}$ which is geometric relative to the theory $DCF_{0,\ell}$. The characterization we present is very much in the spirit of the Pierce-Pillay axioms, and can be used essentially in the same way. It is worth pointing out that first order axiomatizability is a complication that arises in our setting (as explained above this issue does not appear in the ordinary case). Characteristic sets of prime differential ideals are behind our solution to this problem.

It is worth mentioning that Pierce-Pillay type axiomatizations appear in various other contexts: difference fields [8], ordinary differential-difference fields [3], fields with commuting Hasse-Schmidt derivations in positive characteristic [28], fields with free operators [40], and theories having a “geometric notion of genericity” [18]. However, none of the techniques used in these works seem to extend to our context.

4.2 Geometric characterizations and an axiomatization

In this section we work in the language $\mathcal{L}_{\ell+1}$ with derivations $\{\delta_1, \dots, \delta_{\ell+1}\}$. Fix a sufficiently saturated $(\mathbb{U}, \Pi) \models DCF_{0,\ell+1}$ and a small Π -subfield K of \mathbb{U} .

In Chapter 2 we have already presented/developed all the notions and machinery needed to state and prove the following characterization.

Theorem 4.2.1. *Let $\Delta = \{\delta_1, \dots, \delta_\ell\}$ and $D = \delta_{\ell+1}$. Then, $(K, \Pi) \models DCF_{0,\ell+1}$ if and only if*

1. $(K, \Delta) \models DCF_{0,\ell}$
2. *Suppose V and W are irreducible affine Δ -algebraic varieties defined over K such that $W \subseteq \tau_{D/\Delta}V$ and W projects Δ -dominantly onto V . If O_V and O_W are nonempty Δ -open subsets of V and W respectively, defined over K , then there exists a K -point $a \in O_V$ such that $\nabla a \in O_W$.*

Proof. Suppose $(K, \Pi) \models DCF_{0,\ell+1}$, and V, W, O_V and O_W are as in condition (2). Let (a, b) be a Δ -generic point of W over K ; that is, $\mathcal{I}(a, b/K)_\Delta = \mathcal{I}(W/K)_\Delta$. Then $(a, b) \in O_W$. Since $(a, b) \in \tau_{D/\Delta}V$ we have that $d_{D/\Delta}f_a(b) = 0$ for all $f \in \mathcal{I}(V/K)_\Delta$. The fact that W projects Δ -dominantly onto V implies that a is a Δ -generic point of V over K ,

so $a \in O_V$ and $\mathcal{I}(a/K)_\Delta = \mathcal{I}(V/K)_\Delta$. Hence, $d_{D/\Delta} f_a(b) = 0$ for all $f \in \mathcal{I}(a/K)_\Delta$. By Fact 2.2.5, there is a unique Δ -derivation $D' : K\{a\} \rightarrow \mathbb{U}$ extending D such that $D'a = b$. By Fact 2.2.6, we can extend D' to all of \mathbb{U} , call it D'' . Hence, $(\mathbb{U}, \Delta \cup \{D''\})$ is a differential field extending (K, Π) . Since $a \in O_V$, $(a, b) \in O_W$ and $D''a = b$, we get a point (a', b') in K such that $a' \in O_V$, $(a', b') \in O_W$ and $Da' = b'$.

The converse is essentially as in [43]. For the sake of completeness we give the details. Let $\phi(x)$ be a conjunction of atomic $\mathcal{L}_{\ell+1}$ -formulas over K with a realisation a in \mathbb{U} . Let

$$\phi(x) = \psi(x, \delta_{\ell+1}x, \dots, \delta_{\ell+1}^r x)$$

where ψ is a conjunction of atomic \mathcal{L}_ℓ -formulas over K and $r > 0$. Let $c = (a, Da, \dots, D^{r-1}a)$ and V be the Δ -locus of c over K . Let W be the Δ -locus of ∇c over K . Let

$$\chi(x_0, \dots, x_{r-1}, y_0, \dots, y_{r-1}) := \psi(x_0, \dots, x_{r-1}, y_{r-1}) \wedge (\bigwedge_{i=1}^{r-1} x_i = y_{i-1})$$

then χ is realised by $(c, Dc) = \nabla c$. Since $(K, \Delta) \models DCF_{0,\ell}$, and c and ∇c are Δ -generic points of V and W respectively, over K , we have that W projects Δ -dominantly onto V . Also, since $\nabla c \in \tau_{D/\Delta} X$, we have $W \subseteq \tau_{D/\Delta} V$. Applying condition (2) with $V = O_V$ and $W = O_W$, there is a K -point d in V such that $\nabla d \in W$. Let $d = (d_0, \dots, d_{r-1})$ then $(d_0, \dots, d_{r-1}, Dd_0, \dots, Dd_{r-1})$ realises χ . Thus, $(d_0, Dd_0, \dots, D^r d_0)$ realises ψ . Hence, d_0 is a tuple of K realising ϕ . This proves that $(K, \Pi) \models DCF_{0,\ell+1}$. \square

Remark 4.2.2. As can be seen from the proof, it would have been equivalent in condition (2) to take $O_V = V$ and $O_W = W$. Also note that, under the convention that $DCF_{0,0}$ is the theory of algebraically closed fields of characteristic zero ACF_0 , when $\ell = 0$ this is exactly the Pierce-Pillay axioms.

It remains open as to whether condition (2) of Theorem 4.2.1 is expressible in a first order way for $m > 0$. One of the problems lies in determining the definability of differential irreducibility and differential dominance in definable families of differential algebraic varieties (this problem seems to be related to the generalized Ritt problem, see [14] and [16]). Another problem is to determine whether the family of prolongations of a definable family of differential algebraic varieties is again a definable family.

We get around these problems by injecting some differential algebra into the above characterisation, modifying it to make it first order. That is, we use characteristic sets of prime differential ideals (see Section 1.2 for the definition of characteristic sets). We will use the following notation.

Definition 4.2.3. Given a characteristic set Λ of a prime Δ -ideal, by $\mathcal{V}^*(\Lambda)$ we mean $\mathcal{V}(\Lambda) \setminus \mathcal{V}(H_\Lambda)$ where H_Λ is the product of the initials and separants of the elements of Λ (see Section 1.2).

Recall that the notation $\mathcal{V}^*(\Lambda)$ has already been used in Proposition 1.5.9.

Theorem 4.2.4. *Let $\Delta = \{\delta_1, \dots, \delta_\ell\}$ and $D = \delta_{\ell+1}$. Then, $(K, \Pi) \models DCF_{0, \ell+1}$ if and only if*

(i) $(K, \Delta) \models DCF_{0, \ell}$

(ii) *Suppose Λ and Γ are characteristic sets of prime Δ -ideals of $K\{x\}_\Delta$ and $K\{x, y\}_\Delta$ respectively, such that*

$$\mathcal{V}^*(\Gamma) \subseteq \mathcal{V}(f, d_{D/\Delta}f : f \in \Lambda).$$

Suppose O is a nonempty Δ -open subset of $\mathcal{V}^(\Lambda)$ defined over K such that the projection of $\mathcal{V}^*(\Gamma)$ to $\mathcal{V}(\Lambda)$ contains O . Then there is a K -point $a \in \mathcal{V}^*(\Lambda)$ such that $\nabla_D a \in \mathcal{V}^*(\Gamma)$.*

Proof. Suppose $(K, \Pi) \models DCF_{0, \ell+1}$, and we are given Λ, Γ and O satisfying the hypotheses of (ii). Then Λ is a characteristic set of the prime differential ideal $\mathcal{P} = [\Lambda] : H_\Lambda^\infty$ and, by Proposition 1.5.9, $\mathcal{V}^*(\Lambda) = \mathcal{V}(\mathcal{P}) \setminus \mathcal{V}(H_\Lambda)$. So $O_V := \mathcal{V}^*(\Lambda)$ is a nonempty Δ -open subset of the irreducible affine Δ -algebraic variety $V := \mathcal{V}(\mathcal{P})$. Similarly, $O_W := \mathcal{V}^*(\Gamma)$ is a nonempty Δ -open subset of the irreducible affine Δ -algebraic variety $W := \mathcal{V}([\Gamma] : H_\Gamma^\infty)$.

Next we show that $\tau_{D/\Delta}V|_O = \mathcal{V}(f, d_{D/\Delta}f : f \in \Lambda)|_O$. Recall that, by definition, $\tau_{D/\Delta}V$ is $\mathcal{V}(f, d_{D/\Delta}f : f \in \mathcal{I}(V/K)_\Delta)$. It is easy to see that $\mathcal{V}(f, d_{D/\Delta}f : f \in \Lambda) = \mathcal{V}(f, d_{D/\Delta}f : f \in [\Lambda])$. So, supposing that (a, b) is a root of f and $d_{D/\Delta}f$ for all $f \in [\Lambda]$, and $a \in O$, we need to show that (a, b) is a root of $d_{D/\Delta}g$ for all $g \in \mathcal{I}(V/K)_\Delta$. But $\mathcal{I}(V/K)_\Delta = [\Lambda] : H_\Lambda^\infty$, so $H_\Lambda^\ell g \in [\Lambda]$ for some ℓ . We get

$$\begin{aligned} 0 &= d_{D/\Delta}(H_\Lambda^\ell g)(a, b) && \text{as } H_\Lambda^\ell g \in [\Lambda] \\ &= H_\Lambda^\ell(a)d_{D/\Delta}g(a, b) + g(a)d_{D/\Delta}H_\Lambda^\ell(a, b) \\ &= H_\Lambda^\ell(a)\tau g(a, b). \end{aligned}$$

Since O is disjoint from $\mathcal{V}(H_\Lambda)$ we have that $\tau g(a, b) = 0$, as desired.

It follows that a nonempty Δ -open subset of W is contained in $\tau_{D/\Delta}V$, and hence, by irreducibility, $W \subseteq \tau_{D/\Delta}V$. Since O is contained in the projection of W , W projects Δ -dominantly onto V . Applying (2) of Theorem 4.2.1 to V, W, O_V, O_W , we get a K -point $a \in O_V$ such that $\nabla a \in O_W$, as desired.

For the converse we assume that (ii) holds and aim to prove condition (2) of Theorem 4.2.1. In fact, it suffices to prove this statement in the case when $O_V = V$ and $O_W = W$ (see Remark 4.2.2). We thus have irreducible affine Δ -algebraic varieties V

and $W \subseteq \tau_{D/\Delta}V$ such that W projects Δ -dominantly onto V , and we show that there is $a \in V$ such that $\nabla a \in W$. Let Λ and Γ be characteristic sets of $\mathcal{I}(V/K)_\Delta$ and $\mathcal{I}(W/K)_\Delta$, respectively. Then, by Proposition 1.5.9

$$\mathcal{V}^*(\Gamma) = W \setminus \mathcal{V}(H_\Gamma) \subseteq \tau_{D/\Delta}V \subseteq \mathcal{V}(f, d_{D/\Delta}f : f \in \Lambda).$$

Since W projects Δ -dominantly onto V , $\mathcal{V}^*(\Gamma)$ projects Δ -dominantly onto V . Thus, by quantifier elimination for $DCF_{0,\ell}$ and the assumption $(K, \Delta) \models DCF_{0,\ell}$, there is O a nonempty Δ -open subset of $\mathcal{V}^*(\Lambda)$ defined over K such that the projection of $\mathcal{V}^*(\Gamma)$ to $\mathcal{V}(\Lambda)$ contains O . We are in the situation of condition (ii), and there is $a \in \mathcal{V}^*(\Lambda) \subseteq V$ such that $\nabla a \in \mathcal{V}^*(\Gamma) \subseteq W$. \square

Remark 4.2.5. What the above proof shows is that, in a Δ -closed field, each instance of condition (2) of Theorem 4.2.1 is equivalent to an instance of condition (ii) of Theorem 4.2.4. This is accomplished by passing from prime differential ideals to their characteristic sets, and from Δ -dominant projections to containment of a nonempty Δ -open set.

By Fact 1.3.1, the condition that “ $\Lambda = \{f_1, \dots, f_s\}$ is a characteristic set of a prime Δ -ideal of $K\{\bar{x}\}_\Delta$ ” is a definable property on the coefficients of f_1, \dots, f_s . That is, there is an \mathcal{L}_m -formula which specifies for which parameters of the f_i 's, over any model of $DF_{0,m}$, the set Λ is characteristic set of a prime Δ -ideal. Thus, condition (ii) of Theorem 4.2.4 is first order expressible in the language of differential rings. Indeed, the scheme of axioms will consist of $DCF_{0,\ell}$ and an axiom for each choice of the shape of Λ , Γ , O and Q together with a formula specifying for which parameters these finite sets of polynomials satisfy condition (ii). We therefore obtain a *geometric first order axiomatization* for the theory $DCF_{0,\ell+1}$ that refers to the geometry of $DCF_{0,\ell}$.

We conclude by giving another geometric characterization in terms of relative D-varieties that is not on the face of it first order, but may be of independent interest. We show that the models of $DCF_{0,\ell+1}$ are precisely those differential fields that have sharp points on every relative D-variety (where we allow K^Π -linearly independent transformations of Π).

Theorem 4.2.6. $(K, \Pi) \models DCF_{0,\ell+1}$ if and only if

(\dagger) For every basis $\Delta \cup \{D\}$ of $\text{span}_{K^\Pi} \Pi$, if (V, s) is an affine relative D-variety w.r.t. D/Δ defined over K then $(V, s)^\#$ has a K -point.

Proof. Suppose $(K, \Pi) \models DCF_{0,\ell+1}$. By (1) of Theorem 2.4.3, $(V, s)^\#$ is Π -dense in V and hence nonempty. As (K, Π) is existentially closed, $(V, s)^\#$ has a K -point.

For the converse, we assume condition (\dagger) and prove that $(K, \Pi) \models DCF_{0, \ell+1}$. Suppose $\phi(x)$ is a conjunction of atomic $\mathcal{L}_{\ell+1}$ -formulas over K , with a realisation $a = (a_1, \dots, a_n)$ in \mathbb{U} . We need to find a realisation of ϕ in K . We may assume that a is in a Π -closure of K and hence each a_i is Π -algebraic over K .

By Fact 1.6.5, we can find a natural number $k > 0$ and a basis $\Delta \cup \{D\}$ of $\text{span}_{K^\Pi} \Pi$ such that $K\langle a \rangle_\Pi = K\langle a, \dots, D^k a \rangle_\Delta$. Hence,

$$D^{k+1}a = \frac{f(a, \dots, D^k a)}{g(a, \dots, D^k a)} \quad (4.2.1)$$

for some sequence $\frac{f}{g}$ of Δ -rational functions over K . Let

$$c = \left(a, Da, \dots, D^k a, \frac{1}{g(a, Da, \dots, D^k a)} \right).$$

Let V be the Δ -locus of c over K and W the Δ -locus of ∇c over K . A standard trick yields a sequence $s = (\text{Id}, s')$ of Δ -polynomials over K such that $s(c) = \nabla c$. Since c is a generic point of V over K , we get that (V, s) is a relative D-variety w.r.t. D/Δ defined over K .

Claim. $W = s(V)$.

Let $b \in V$. If h is a Δ -polynomial over K vanishing at (c, Dc) , then $h(\cdot, s(\cdot))$ vanishes at c and hence on all of V . So $(b, s(b))$ is in the Δ -locus of (c, Dc) over K . That is, $(b, s(b)) \in Y$. The other containment is clear since $\nabla c = s(c) \in s(V)$.

By condition (\dagger) , $(V, s)^\#$ has a K -point say d . Hence, $d \in V$ and $\nabla d = s(d)$. The previous claim yields that $(d, Dd) \in W$.

Now, let $\rho(x)$ be the $\mathcal{L}_{\ell+1}$ -formula over K obtained from ϕ by replacing each δ_i , $i = 1, \dots, \ell + 1$, for $e_{i,1}\delta_1 + \dots + e_{i,\ell+1}\delta_{\ell+1}$, where $(e_{i,j}) \in \text{GL}_{\ell+1}(K^\Pi)$ is the transition matrix taking the basis $\Delta \cup \{D\}$ to the basis Π . By construction, $\phi^{(K, \Pi)} = \rho^{(K, \Delta \cup \{D\})}$. Thus it suffices to find a realisation of ρ in $(K, \Delta \cup \{D\})$. We may assume that the k of (4.2.1) is large enough so that we can write

$$\rho(x) = \psi(x, \delta_{\ell+1}x, \dots, \delta_{\ell+1}^k x)$$

where ψ is a conjunction of atomic \mathcal{L}_ℓ -formulas over K . Let

$$\chi(x_0, \dots, x_{k+1}, y_0, \dots, y_{k+1}) := \psi(x_0, \dots, x_k) \wedge (\wedge_{i=1}^k x_i = y_{i-1}).$$

Then $(\mathbb{U}, \Delta) \models \chi(c, Dc)$, and so, as $(d, Dd) \in W$, we have that $(\mathbb{U}, \Delta) \models \chi(d, Dd)$. But since d is a K -point, we get $(K, \Delta) \models \chi(d, Dd)$. Writing the tuple d as (d_0, \dots, d_{k+1}) , we see that d_0 is a realisation of ρ in $(K, \Delta \cup \{D\})$. This completes the proof. \square

Notice that, in contrast with condition (2) of Theorem 4.2.1, condition (\dagger) does not involve differential irreducibility or differential dominance; however, it does involve the condition $s : V \rightarrow \tau_{D/\Delta}V$. Hence, since it is not known if the family of prolongations of a definable family of differential algebraic varieties is definable, it also remains open as to whether condition (\dagger) is expressible in a first order way for $m > 0$.

Chapter 5

The model theory of partial differential fields with an automorphism

It has been known for over fifteen years that the theory of ordinary differential fields of characteristic zero with an automorphism admits a model companion. In other words, the class of existentially closed ordinary differential fields of characteristic zero equipped with an automorphism is an axiomatizable class. This theory has been studied extensively by Bustamante [3], [4], [2]. However, the techniques used there to prove the existence of the model companion do not extend to partial differential fields with an automorphism. In this chapter we prove that the latter theory indeed has a model companion. After establishing some of its basic model theoretic properties, we will prove the canonical base property and the Zilber dichotomy for finite dimensional types.

Our approach is similar to the one from Section 4.2. We observe that the existentially closed partial differential fields with an automorphism are characterized by a certain geometric condition in terms of irreducible Δ -varieties and Δ -dominant projections (this already appears in [15]), very much in the spirit of the geometric axioms for *ACFA*. Of course, we run into the same difficulties as we did for Theorem 4.2.1 when trying to express this geometric condition in a first order way: it is not known if differential irreducibility and differential dominance are definable conditions in definable families of differential algebraic varieties. As we did in Section 4.2, we bypass these issues by applying the differential algebraic machinery of characteristic sets of prime differential ideals.

5.1 A review of *ACFA*

Fields equipped with an automorphism (i.e., difference fields) enjoy some formal similarities to differential fields. While our eventual purpose is to consider a combination of differential and difference structure, we begin by reviewing in this section the model theory of difference fields. For a complete exposition of difference algebra the reader may consult [10], and for the model theoretic results [8].

We work in the language $\mathcal{L}_\sigma = \{0, 1, +, -, \times, \sigma\}$ of rings expanded by a unary function symbol σ .

Definition 5.1.1. By a *difference field*, or σ -*field*, we mean a field K equipped with an automorphism $\sigma : K \rightarrow K$ (often one only asks for σ to be injective, what we call here difference field is sometimes referred as *inversive difference field*). The *fixed field* of a difference field (K, σ) is by definition $K^\sigma = \{a \in K : \sigma a = a\}$.

Let (K, σ) be a difference field. If V is an affine algebraic variety defined over K then V^σ is the affine algebraic variety defined by $\{f^\sigma = 0 : f \in \mathcal{I}(V/K)\}$, where f^σ is the polynomial obtained by applying σ to the coefficients of f .

Fact 5.1.2 ([8], §1). *The class of existentially closed difference fields is axiomatizable. This theory, denoted by *ACFA*, is axiomatized by the scheme of axioms expressing:*

- (i) $K \models \text{ACF}$
- (ii) *Suppose V and W are irreducible affine algebraic varieties defined over K such that $W \subseteq V \times V^\sigma$ and W projects dominantly onto both V and V^σ . Then there is a K -point $a \in V$ such that $(a, \sigma a) \in W$.*

In [8], Chatzidakis and Hrushovski carry out a thorough analysis of the theory *ACFA*. In particular, they show that:

- The fixed field of $(K\sigma) \models \text{ACFA}$ is a pseudofinite field; that is, K^σ is perfect, K^σ has a unique extension of degree n for all $n < \omega$, and K^σ is pseudo-algebraically closed (i.e., every absolutely irreducible variety defined over K has a K -point).
- If $A \subseteq K$, then $\text{acl}(A)$ is the field theoretic algebraic closure of the σ -field generated by A .

- The completions of $ACFA$ are determined by specifying the characteristic and the action of σ on the algebraic closure of the prime field. Each completion of $ACFA$ admits elimination of imaginaries.

Let us restrict our attention to characteristic zero. Fix a sufficiently saturated model $(\mathbb{U}, \sigma) \models ACFA_0$. Given subsets A, B, C of \mathbb{U} define A to be *independent from B over C* , denoted by $A \downarrow_C B$, if $\mathbb{Q}\langle A \cup C \rangle_\sigma$ is algebraically disjoint from $\mathbb{Q}\langle B \cup C \rangle_\sigma$ over $\mathbb{Q}\langle C \rangle_\sigma$. Here the notation $\mathbb{Q}\langle A \rangle_\sigma$ means the σ -field generated by A .

Then, as in the case of differentially closed fields, this captures Shelah's nonforking: if $C \subseteq B$ and a is a tuple, $tp(a/B)$ is a nonforking extension of $tp(a/C)$ if and only if $a \downarrow_C B$.

In [8], Chatzidakis and Hrushovski proved that $ACFA_0$ is *supersimple*; that is, they prove the following properties of independence:

1. (*Invariance*) If ϕ is an automorphism of \mathbb{U} and $A \downarrow_C B$, then $\phi(A) \downarrow_{\phi(C)} \phi(B)$.
2. (*Local character*) There is a finite subset $B_0 \subseteq B$ such that $A \downarrow_{B_0} B$.
3. (*Extension*) There is a tuple b such that $tp(a/C) = tp(b/C)$ and $b \downarrow_C B$.
4. (*Symmetry*) $A \downarrow_C B$ if and only if $B \downarrow_C A$.
5. (*Transitivity*) Suppose $C \subseteq B \subseteq D$, then $A \downarrow_C D$ if and only if $A \downarrow_C B$ and $A \downarrow_B D$.
6. (*Independence theorem*) Suppose C is an algebraically closed σ -field, and
 - (i) A and B are supersets of C with $A \downarrow_C B$, and
 - (ii) a and b are tuples such that $tp(a/C) = tp(b/C)$ and $a \downarrow_C A$ and $b \downarrow_C B$.

Then there is a common nonforking extension $p \in S_n(A \cup B)$ of $tp(a/A)$ and $tp(b/B)$.

Note that unlike the case of differentially closed fields we do not have stationarity (see Fact 1.4.1). The independence theorem takes its place. This corresponds to the fact that $ACFA_0$ is not a *stable* theory.

Now, exactly as in the case of differentially closed field, we make use of the foundation rank associated to the partial ordering on types induced by forking. For historical reasons,

this is called SU -rank in this case, rather than U -rank, but the definition is exactly as in Definition 1.4.2, and it also enjoys the appropriate analogues of the properties listed in Fact 1.4.3.

A very consequential property of $ACFA_0$ proved in [8] is

Fact 5.1.3 (Zilber dichotomy). *Suppose K is an algebraically closed σ -field and $p = tp(a/K)$ is of SU -rank 1. Then p is either one-based or almost \mathbb{U}^σ -internal.*

Here $p = tp(a/K)$ is *almost \mathbb{U}^σ -internal* means that there is a σ -field L extension of K with $a \downarrow_K L$ such that $a \in \text{acl}(L, c) = L\langle a \rangle_\sigma^{\text{alg}}$ for some tuple c from \mathbb{U}^σ , and *one-based* means that for every algebraically closed σ -field L extension of K we have that $Cb(a/L) \subseteq \text{acl}(K, a)$ where $Cb(a/L)$ is the *canonical base* of $tp(a/L)$ which we explain below.

The Zilber dichotomy can be seen as a consequence of the following result proved by Pillay and Ziegler in [52]. First, given a complete type $tp(a/K)$, where K is an algebraically σ -field, the *canonical base* of p , denoted by $Cb(p)$, is a (finite) tuple from K that comes from the study of supersimple theories in general and was introduced in [17]. We will not give the precise definition here, but content ourselves with pointing out that in $ACFA_0$, up to interalgebraicity, it is the same thing as the minimal σ -field of definition of the σ -locus of a over K . That is, if V is the intersection of all solution sets of systems of difference polynomials over K that contain a , which is again given by a (finite) system of such polynomials [10], and F is the minimal σ -field of definition of V , which exists by model theoretic considerations (by elimination of imaginaries, see for example Chapter 1.1 of [46]), then $\text{acl}(Cb(p)) = F^{\text{alg}}$.

Fact 5.1.4 (Canonical base property, [52]). *Suppose a is a tuple such that $SU(a/K) < \omega$ and L is an algebraically closed σ -field extension of K . Then $tp(Cb(a/L)/a, K)$ is almost \mathbb{U}^σ -internal.*

Our goal in this chapter will be to extend the results discussed here to the case of (partial) differential-difference fields. In the case of a single derivation, this was done by Bustamante in [3] and [2]. That is, consider the class of ordinary differential-difference fields (K, δ, σ) where K is a field of characteristic zero, δ is a derivation, and σ is an automorphism of (K, δ) (i.e., an automorphism of K such that $\delta\sigma = \sigma\delta$). Then the class of existentially closed ordinary differential-difference fields is axiomatizable, and the theory is denoted by DCF_0A . The axioms given in [3] do not precisely generalize those of Fact 5.1.2 above, and do not extend to give an axiomatization of existentially closed partial differential-difference fields (K, Δ, σ) , where Δ is a set of $m > 1$ commuting derivations on K and

σ is an automorphism of (K, Δ) . This will be accomplished in the next section, thereby introducing $DCF_{0,m}A$. Nevertheless, Bustamante does show that DCF_0A exists, that its completions are described by the σ -field structure on \mathbb{Q}^{alg} , that each completion eliminates imaginaries and is supersimple (i.e., for the appropriate definition of \perp properties (1)-(6) stated above for $ACFA$, hold for DCF_0A), and satisfies the canonical base property (and hence also the Zilber dichotomy). We will prove all these properties for $DCF_{0,m}A$, except that for the canonical base property (and Zilber dichotomy) we will restrict ourselves to *finite dimensional* types.

5.2 The model companion $DCF_{0,m}A$

We work in the language $\mathcal{L}_{m,\sigma}$ of differential rings with m derivations expanded by a unary function symbol. We denote by $DF_{0,m,\sigma}$ the $\mathcal{L}_{m,\sigma}$ -theory of differential fields of characteristic zero with m commuting derivations and an automorphism that commutes with the derivations. We say that (K, Δ, σ) is a differential-difference field if $(K, \Delta, \sigma) \models DF_{0,m,\sigma}$. We are interested in axiomatizing the existentially closed differential-difference fields, and we will begin with a known geometric characterization of them.

First some notation. If V is an affine Δ -variety defined over K then V^σ is the affine Δ -variety defined by $\{f^\sigma = 0 : f \in \mathcal{I}(V/K)_\Delta\}$, where f^σ is the differential polynomial obtained by applying σ to the coefficients of f . Here, when speaking of Δ -varieties, we are implicitly working in a sufficiently large saturated $(\mathbb{U}, \Delta) \models DCF_{0,m}$, where our affine Δ -varieties live in.

Fact 5.2.1 ([15], §2). *Let (K, Δ, σ) be a differential-difference field. Then, (K, Δ, σ) is existentially closed if and only if*

(i) $(K, \Delta) \models DCF_{0,m}$

(ii) *Suppose V and W are irreducible affine Δ -varieties defined over K such that $W \subseteq V \times V^\sigma$ and W projects Δ -dominantly onto both V and V^σ . If O_V and O_W are nonempty Δ -open sets of V and W respectively, defined over K , then there is a K -point $a \in O_V$ such that $(a, \sigma a) \in O_W$.*

Proof. Formally this is done in [15] for the case when $O_V = V$ and $O_W = W$. The argument there can easily be adapted to the more general statement presented above. Nevertheless, we give some details for the sake of completeness.

Suppose (K, Δ, σ) is existentially closed and V, W, O_V and O_W are as in condition (ii). Let (c, d) be a Δ -generic point of W over K ; that is, $I(c, d/K)_\Delta = I(W/K)_\Delta$. Clearly $(c, d) \in O_W$ and, since W projects Δ -dominantly onto V and V^σ , c and d are generic points of V and V^σ , respectively, over K . Thus, $c \in O_V$. Because $DCF_{0,m}$ has quantifier elimination we have that $tp_\Delta(d/K) = \sigma(tp_\Delta(c/K))$, where $tp_\Delta(a/K)$ denotes the type of a over K in the language of differential rings. Hence, there is an automorphism σ' of \mathbb{U} extending σ such that $\sigma'(c) = d$. Since (K, Δ, σ) is existentially closed we can find a point in K with the desired properties.

Now suppose conditions (i) and (ii) are satisfied. Let $\phi(x)$ be a conjunction of atomic $\mathcal{L}_{m,\sigma}$ -formulas over K . Suppose ϕ has a realisation a in some differential-difference field (F, Δ, σ) extending (K, Δ, s) . We may assume, and we do, that (F, Δ) is a Δ -subfield of our ambient differentially closed field (\mathbb{U}, Δ) . Let

$$\phi(x) = \psi(x, \sigma x, \dots, \sigma^r x)$$

where ψ is a conjunction of atomic \mathcal{L}_m -formulas over K and $r > 0$. Let $b = (a, \sigma a, \dots, \sigma^{r-1}a)$ and V be the Δ -locus of b over K . Let W be the Δ -locus of $(b, \sigma b)$ over K . Let

$$\chi(x_0, \dots, x_{r-1}, y_0, \dots, y_{r-1}) := \psi(x_0, \dots, x_{r-1}, y_{r-1}) \wedge (\bigwedge_{i=1}^{r-1} x_i = y_{i-1})$$

then χ is realised by $(b, \sigma b)$. Using the assumption that $(K, \Delta) \models DCF_{0,m}$, and since $(b, \sigma b)$ is a Δ -generic point of W , b is a Δ -generic point of V and σb is a Δ -generic point of V^σ , over K , we have that W projects Δ -dominantly onto both V and V^σ . Applying (2) with $V = O_V$ and $W = O_W$, there is c in V such that $(c, \sigma c) \in W$. Let $c = (c_0, \dots, c_{r-1})$, then $(c_0, \dots, c_{r-1}, \sigma c_0, \dots, \sigma c_{r-1})$ realises χ . Thus, $(c_0, \sigma c_0, \dots, \sigma^r c_0)$ realises ψ . Hence, c_0 is a tuple from K realising ϕ . This proves (K, Δ, σ) is existentially closed. \square

Remark 5.2.2. As can be seen from the proof, it would have been equivalent in condition (2) to take $O_V = V$ and $O_W = W$. Also note that, under the convention that $DCF_{0,0}$ is ACF_0 , when $m = 0$ this is exactly the geometric axiomatization of $ACFA_0$.

Condition (ii) is not, on the face of it, first order as it is not known if differential irreducibility and differential dominance are definable properties. We have seen this difficulty before, in Chapter 4, and we remedy the situation similarly; by working with characteristic sets of prime differential ideals. Recall that for Λ a characteristic set of a prime Δ -ideal, $\mathcal{V}^*(\Lambda) := \mathcal{V}(\Lambda) \setminus \mathcal{V}(H_\Lambda)$ (see definition 4.2.3).

Theorem 5.2.3. *Let (K, Δ, σ) be a differential-difference field. Then, (K, Δ, σ) is existentially closed if and only if*

1. $(K, \Delta) \models DCF_{0,m}$
2. Suppose Λ and Γ are characteristic sets of prime differential ideals of $K\{x\}$ and $K\{x, y\}$, respectively, such that

$$\mathcal{V}^*(\Gamma) \subseteq \mathcal{V}(\Lambda) \times \mathcal{V}(\Lambda^\sigma).$$

Suppose O and Q are nonempty Δ -open subsets of $\mathcal{V}^*(\Lambda)$ and $\mathcal{V}^*(\Lambda^\sigma)$ respectively, defined over K , such that $O \subseteq \pi_x(\mathcal{V}^*(\Gamma))$ and $Q \subseteq \pi_y(\mathcal{V}^*(\Gamma))$. Then there is a K -point $a \in \mathcal{V}^*(\Lambda)$ such that $(a, \sigma a) \in \mathcal{V}^*(\Gamma)$.

Here, if $\Lambda = \{f_1, \dots, f_s\}$, then $\Lambda^\sigma = \{f_1^\sigma, \dots, f_s^\sigma\}$, and π_x and π_y denote the canonical projections from $\mathcal{V}(\Lambda) \times \mathcal{V}(\Lambda^\sigma)$ to $\mathcal{V}(\Lambda)$ and $\mathcal{V}(\Lambda^\sigma)$, respectively.

Proof. Suppose (K, Δ, σ) is existentially closed, and let Λ , Γ , O and Q be as in condition (2). Then Λ is a characteristic set of the prime differential ideal $\mathcal{P} = [\Lambda] : H_\Lambda^\infty$ and, by Proposition 1.5.9, $\mathcal{V}^*(\Lambda) = \mathcal{V}(\mathcal{P}) \setminus \mathcal{V}(H_\Lambda)$. So $O_V := \mathcal{V}^*(\Lambda)$ is a nonempty Δ -open subset of the irreducible affine Δ -variety $V := \mathcal{V}(\mathcal{P})$. Similarly, $O_W := \mathcal{V}^*(\Gamma)$ is a nonempty Δ -open subset of the irreducible affine Δ -variety $W := \mathcal{V}([\Gamma] : H_\Gamma^\infty)$. Note that both O_V and O_W are defined over K .

Next we show that $W \subseteq V \times V^\sigma$. By Proposition 1.5.9

$$W \setminus \mathcal{V}(H_\Gamma) = \mathcal{V}^*(\Gamma) \subseteq \mathcal{V}(\Lambda) \times \mathcal{V}(\Lambda^\sigma),$$

so that by taking Δ -closures $W \subseteq \mathcal{V}(\Lambda) \times \mathcal{V}(\Lambda^\sigma)$. Since W is irreducible it must be contained in some $X \times Y$ where X and Y are irreducible components of $\mathcal{V}(\Lambda)$ and $\mathcal{V}(\Lambda^\sigma)$, respectively. Hence $O \subseteq \pi_x(W) \subseteq X$. On the other hand, $O \subseteq O_V \subset V$, and so $V = X$. Similarly, since $V^\sigma = \mathcal{V}([\Lambda^\sigma] : H_{\Lambda^\sigma}^\infty)$ and working with Q rather than O , we get $Y = V^\sigma$. Therefore, $W \subseteq V \times V^\sigma$. Since $O \subseteq \pi_x(W)$ and $Q \subseteq \pi_y(W)$, W projects Δ -dominantly onto both V and V^σ . Applying (ii) of Fact 5.2.1 to V , W , O_V , O_W , we get a K -point $a \in O_V$ such that $(a, \sigma a) \in O_W$, as desired.

For the converse we assume that (2) holds and aim to prove condition (ii) of Fact 5.2.1. In fact, it suffices to prove this statement in the case when $O_V = V$ and $O_W = W$ (see Remark 5.2.2). We thus have irreducible affine Δ -varieties V and $W \subseteq V \times V^\sigma$ such that W projects Δ -dominantly onto both V and V^σ , and we show that there is a K -point $a \in V$ such that $(a, \sigma a) \in W$. Let Λ and Γ be characteristic sets of $\mathcal{I}(V/K)_\Delta$ and $\mathcal{I}(W/K)_\Delta$, respectively. Then, by Proposition 1.5.9

$$\mathcal{V}^*(\Gamma) = W \setminus \mathcal{V}(H_\Gamma) \subseteq V \times V^\sigma \subseteq \mathcal{V}(\Lambda) \times \mathcal{V}(\Lambda^\sigma).$$

Since W projects Δ -dominantly onto V and V^σ , $\mathcal{V}^*(\Gamma)$ projects Δ -dominantly onto both V and V^σ . Thus, by quantifier elimination for $DCF_{0,m}$ and the assumption $(K, \Delta) \models DCF_{0,m}$, there are nonempty Δ -open sets O and Q of $\mathcal{V}^*(\Lambda)$ and $\mathcal{V}^*(\Lambda^\sigma)$ respectively, defined over K , such that $O \subseteq \pi_x(\mathcal{V}^*(\Gamma))$ and $Q \subseteq \pi_y(\mathcal{V}^*(\Gamma))$. We are in the situation of condition (2), and there is a K -point $a \in \mathcal{V}^*(\Lambda) \subseteq V$ such that $(a, \sigma a) \in \mathcal{V}^*(\Gamma) \subseteq W$. \square

Because being a characteristic set of a prime Δ -ideal is a definable property (see Fact1.3.1), statement (2) of Theorem 5.2.3 is first-order expressible. We thus have a first order axiomatization of the existentially closed differential-difference fields. That is,

Corollary 5.2.4. *The theory $DF_{0,m,\sigma}$ has a model companion.*

Henceforth we denote this model companion by $DCF_{0,m}A$.

5.3 Basic model theory of $DCF_{0,m}A$

In this section we present some of the model theoretic properties of the theory $DCF_{0,m}A$. Many of these results are consequences of the work of Chatzidakis and Pillay in [9] or more or less immediate adaptations of the arguments from [3] or §5 of [40].

Let (K, Δ, σ) be a differential-difference field and $A \subseteq K$. The differential-difference field generated by A , denoted by $\mathbb{Q}\langle A \rangle_{\Delta, \sigma}$, is the smallest differential-difference subfield of K containing A . Note that $\mathbb{Q}\langle A \rangle_{\Delta, \sigma}$ is simply the subfield of K generated by

$$\{\delta_m^{e_m} \cdots \delta_1^{e_1} \sigma^k a : a \in A, e_i < \omega, k \in \mathbb{Z}\}.$$

If k is a differential-difference subfield of K , we write $k\langle B \rangle_{\Delta, \sigma}$ instead of $\mathbb{Q}\langle k \cup B \rangle_{\Delta, \sigma}$.

Proposition 5.3.1. *Let (K, Δ, σ) and (L, Δ', σ') be models of $DCF_{0,m}A$.*

- (i) *Suppose K and L have a common algebraically closed differential-difference subfield F , then $(K, \Delta, \sigma) \equiv_F (L, \Delta', \sigma')$. In particular, the completions of $DCF_{0,m}A$ are determined by the difference field structure on \mathbb{Q}^{alg} .*
- (ii) *If $A \subseteq K$, then $\text{acl}(A) = \mathbb{Q}\langle A \rangle_{\Delta, \sigma}^{alg}$.*
- (iii) *Suppose F is a differential-difference subfield of K and a, b are tuples from K . Then $tp(a/F) = tp(b/F)$ if and only if there is an F -isomorphism from $(F\langle a \rangle_{\Delta, \sigma}^{alg}, \Delta, \sigma)$ to $(F\langle b \rangle_{\Delta, \sigma}^{alg}, \Delta, \sigma)$ sending a to b .*

Proof. In [9], Chatzidakis and Pillay prove the following general results. Suppose T is a stable theory with elimination of imaginaries and of quantifiers, and T_0 is the theory whose models are structures (\mathcal{M}, σ) where $\mathcal{M} \models T$ and $\sigma \in \text{Aut}(\mathcal{M})$. Assuming that T_0 has a model companion TA , they show that:

- If $(\mathcal{M}_1, \sigma_1)$ and $(\mathcal{M}_2, \sigma_2)$ are models of TA , containing a common substructure (A, σ) such that $A = \text{acl}_T(A)$ (i.e., A is algebraically closed in the sense of T), then $(\mathcal{M}_1, \sigma_1) \equiv_A (\mathcal{M}_2, \sigma_2)$.
- If $(\mathcal{M}, \sigma) \models TA$ and $A \subseteq M$, then

$$\text{acl}_{TA}(A) = \text{acl}_T(\sigma^k(A) : k \in \mathbb{Z}).$$

Hence, (i) and (ii) follow from this general results by specializing T to $DCF_{0,m}$ and the fact that we have already shown that the model companion $TA = DCF_{0,m}A$ exists (Corollary 5.2.4).

We now prove (iii). Suppose $tp(a/F) = tp(b/F)$, working in a sufficiently saturated $(\mathbb{U}, \Delta, \sigma) \models DCF_{0,m}A$ extension of (K, Δ, σ) , we get $\phi \in \text{Aut}(\mathbb{U}/F)$ such that $\phi(a) = b$. Then, the restriction $\phi|_{F\langle a \rangle_{\Delta, \sigma}^{alg}}$ yields the desired isomorphism. Conversely, suppose f is an F -isomorphism from $(F\langle a \rangle_{\Delta, \sigma}^{alg}, \Delta, \sigma)$ to $(F\langle b \rangle_{\Delta, \sigma}^{alg}, \Delta, \sigma)$ such that $f(a) = b$. Let $(E, \Delta) \models DCF_{0,m}$ be a sufficiently large saturated extension of (K, Δ) . Then there is a Δ -automorphism ρ of E extending f . Let $L = \rho(K)$. Then ρ induces a difference structure σ' on L such that ρ is now an isomorphism between (K, Δ, σ) and (L, Δ, σ') . Note that $F\langle b \rangle_{\Delta, \sigma}^{alg}$ is a common substructure of (K, Δ, σ) and (L, Δ, σ') . If $\phi(x)$ is an $\mathcal{L}_{m, \sigma}$ -formula over F , we get

$$K \models \phi(a) \iff L \models \phi(b) \iff K \models \phi(b)$$

where the second equivalence follows from part (i). We have shown that $tp(a/F) = tp(b/F)$, as desired. \square

Let (K, Δ, σ) be a differential-difference field. We denote by K^σ the *fixed field* of K , that is $K^\sigma = \{a \in K : \sigma a = a\}$, and by K^Δ we denote the *field of (total) constants* of K , that is $K^\Delta = \{a \in K : \delta a = 0 \text{ for all } \delta \in \Delta\}$. We let \mathcal{C}_K be the field $K^\sigma \cap K^\Delta$.

More generally, if Δ' is a set of linearly independent elements of the \mathcal{C}_K -vector space $\text{span}_{\mathcal{C}_K} \Delta$, we let $K^{\Delta'}$ be the field of Δ' -constants of K , that is $K^{\Delta'} = \{a \in K : \delta a = 0 \text{ for all } \delta \in \Delta'\}$. In particular, $K^\emptyset = K$ and if Δ' is a basis of $\text{span}_{\mathcal{C}_K} \Delta$ then $K^{\Delta'} = K^\Delta$. Note that both K^σ and $K^{\Delta'}$ are differential-difference subfields of (K, Δ, σ) . Also, (K, Δ', σ) is itself a differential-difference field.

In the following proposition $DCF_{0,0}A$ stands for $ACFA_0$.

Proposition 5.3.2. *Let $(K, \Delta, \sigma) \models DCF_{0,m}A$, Δ_1 and Δ_2 disjoint sets such that $\Delta' := \Delta_1 \cup \Delta_2$ forms a basis of $\text{span}_{\mathcal{C}_K} \Delta$, and $r = |\Delta_1|$. Then*

1. $(K, \Delta', \sigma) \models DCF_{0,m}A$
2. $(K, \Delta_1, \sigma) \models DCF_{0,r}A$
3. $(K^{\Delta_2}, \Delta_1, \sigma) \models DCF_{0,r}A$
4. $K^{\Delta_2} \cap K^\sigma$ is a pseudofinite field.

Proof.

(1) It is easy to see that a set $V \subseteq K^n$ is Δ -closed if and only if it is Δ' -closed. Hence, irreducibility in the Δ -topology is equivalent to irreducibility in the Δ' -topology. Similarly a projection (onto any set of coordinates) is Δ -dominant if and only if it is Δ' -dominant. Therefore, each instance of the axioms of $DCF_{0,m}A$ (or rather of the characterization given by Fact 5.2.1) that needs to be checked for Δ' is true, as it is true for Δ .

(2) By (1) we may assume that $\Delta_1 \subseteq \Delta$. One needs to show that every instance of the axioms of $DCF_{0,r}A$ that needs to be checked for Δ_1 is an instance of the axioms of $DCF_{0,m}A$ which we know is true for Δ . For this, one needs to know that every Δ_1 -closed set irreducible in the Δ_1 -topology is also irreducible in the Δ -topology. This is not obvious, rather it is a general version of Kolchin's Irreducibility Theorem. A proof of this can be found in Chap. 7 of [13].

(3) By (1) we may assume that $\Delta_1 \cup \Delta_2 = \Delta$. We show that $(K^{\Delta_2}, \Delta_1, \sigma)$ is existentially closed. Let $\phi(x)$ be a quantifier free $\mathcal{L}_{\Delta_1, \sigma}$ -formula over K^{Δ_2} with a realisation a in some differential-difference field (F, Ω_1, γ) extending $(K^{\Delta_2}, \Delta_1, \sigma)$. Let $\Omega := \Omega_1 \cup \{\rho_1, \dots, \rho_{m-r}\}$ where each ρ_i is the trivial derivation on F . Hence, (F, Ω, γ) is a differential-difference field extending $(K^{\Delta_2}, \Delta, \sigma)$. Let (L, Ω, γ) be a model of $DCF_{0,m}A$ extending (F, Ω, γ) . Since K^{Δ_2} is a common algebraically closed differential-difference subfield of K and L , and $L \models \phi(a)$ and $\rho_i a = 0$ for $i = 1, \dots, m-r$, by (i) of Proposition 5.3.1 K has a realisation b of ϕ such that $b \in K^{\Delta_2}$. Thus, since ϕ is quantifier free, $K^{\Delta_2} \models \phi(b)$.

(4) By (3), we have that $(K^{\Delta_2}, \Delta_1, \sigma) \models DCF_{0,r}A$, and so, by (2), $(K^{\Delta_2}, \sigma) \models ACF_{A_0}$. Hence, the fixed field of (K^{Δ_2}, σ) , which is $K^{\Delta_2} \cap K^\sigma$, is pseudofinite. \square

We now proceed to show that $DCF_{0,m}A$ is supersimple. Fix a sufficiently saturated $(U, \Delta, \sigma) \models DCF_{0,m}A$ and a differential-difference subfield K of \mathbb{U} . Given subsets A, B, C of \mathbb{U} define A to be *independent from B over C* , denoted by $A \downarrow_C B$, if $\mathbb{Q}\langle A \cup C \rangle_{\Delta, \sigma}$ is algebraically disjoint from $\mathbb{Q}\langle B \cup C \rangle_{\Delta, \sigma}$ from $\mathbb{Q}\langle C \rangle_{\Delta, \sigma}$.

Theorem 5.3.3. *Every completion of $DCF_{0,m}A$ is supersimple; that is, independence satisfies the following properties:*

1. (Invariance) *If ϕ is an automorphism of \mathbb{U} and $A \downarrow_C B$, then $\phi(A) \downarrow_{\phi(C)} \phi(B)$.*
2. (Local character) *There is a finite subset $B_0 \subseteq B$ such that $A \downarrow_{B_0} B$.*
3. (Extension) *There is a tuple b such that $tp(a/C) = tp(b/C)$ and $b \downarrow_C B$.*
4. (Symmetry) *$A \downarrow_C B$ if and only if $B \downarrow_C A$.*
5. (Transitivity) *Suppose $C \subseteq B \subseteq D$, then $A \downarrow_C D$ if and only if $A \downarrow_C B$ and $A \downarrow_B D$.*
6. (Independence theorem) *Suppose C is an algebraically closed differential-difference field, and*
 - (i) *A and B are supersets of C with $A \downarrow_C B$, and*
 - (ii) *a and b are tuples such that $tp(a/C) = tp(b/C)$ and $a \downarrow_C A$ and $b \downarrow_C B$.*

Then there is $d \downarrow_C A \cup B$ with $tp(d/A) = tp(a/A)$ and $tp(d/B) = tp(b/B)$.

In particular, by Theorem 4.2 of [22], our notion of independence captures Shelah's nonforking. That is, if $C \subseteq B$ and a is a tuple, $tp(a/B)$ is a nonforking extension of $tp(a/C)$ if and only if $a \downarrow_C B$.

Proof. As in the proof of Proposition 5.3.1, let us recall some general results of Chatzidakis and Pillay from [9]. Suppose T is a stable theory with elimination of imaginaries and of quantifiers, and T_0 is the theory whose models are structures (\mathcal{M}, σ) where $\mathcal{M} \models T$ and $\sigma \in \text{Aut}(\mathcal{M})$. Assuming that T_0 has a model companion TA and letting (\mathcal{M}, σ) be a sufficiently saturated model of TA , they define a notion of independence in TA as follows: if A, B, C are subsets of M , then $A \downarrow_C B$ iff $\text{acl}_T(\sigma^k(A \cup C) : k \in \mathbb{Z})$ is independent from $\text{acl}_T(\sigma^k(B \cup C) : k \in \mathbb{Z})$ over $\text{acl}_T(\sigma^k(C) : k \in \mathbb{Z})$ in the sense of T , here acl_T means algebraic closure in the sense of T . Then they show that this notion of independence in TA satisfies (1)-(5) above and, under the assumption that $(C, \sigma) \models T_0$, also (6).

Hence, by specializing T to $DCF_{0,m}$ and using the fact that in this case the model companion $TA = DCF_{0,m}A$ exists (Corollary 5.2.4), we get (1)-(5) for our notion of independence in $DCF_{0,m}A$. Note that the results of Chatzidakis and Pillay imply (6) under

the additional assumption that $(C, \Delta) \models DCF_{0,m}$. However, as we now show, in the case of $DCF_{0,m}A$ this assumption is unnecessary.

We follow very closely the strategy used by Moosa and Scanlon in Theorem 5.9 of [40] where they prove the analogous result for fields with free operators (in fact their strategy is an extension of the one used by Chatzidakis and Hrushovski in [8]).

Let $c \models tp(a/C)$. It suffices to find A', B' such that

- (i) $c \perp_C A' \cup B'$
- (ii) $A', c \models tp(A, a/C)$
- (iii) $B', c \models tp(B, b/C)$
- (iv) $A', B' \models tp(A, B/C)$.

Indeed, by (iv) there would be $\phi \in Aut(\mathbb{U}/C)$ such that $\phi(A' \cup B') = A \cup B$, then if we set $d = \phi(c)$ we would get: from (i) that $d \perp_C A \cup B$, from (ii) that $tp(d/A) = tp(a/A)$, and from (iii) that $tp(d/B) = tp(b/B)$.

Since $tp(c/C) = tp(a/C) = tp(b/C)$, there are A', B' satisfying (ii) and (iii). We may assume, by extensionality, that $A' \perp_{C,c} B'$, and, by transitivity, this implies (i). We are only missing (iv).

Let $K_0 = \text{acl}(A') \cdot \text{acl}(B')$ and $K_1 = \text{acl}(A', c) \cdot \text{acl}(B', c)$ where \cdot means compositum in the sense of fields. Also, let $K_2 = \text{acl}(A', B')$. Then, K_1 and K_2 are field extensions of K_0 . We will induce on K_2 a differential-difference field structure $(\Delta' = \{\delta'_1, \dots, \delta'_m\}, \sigma')$ such that there is an isomorphism over C between (K_2, Δ', σ') and $(\text{acl}(A, B), \Delta, \sigma)$. Let $\alpha, \beta \in Aut(\mathbb{U}, C)$ be such that α takes A to A' and β takes B to B' . Since $A \perp_C B$ and $A' \perp_C B'$ and C is algebraically closed, we get that $\text{acl}(A)$ and $\text{acl}(A')$ are linearly disjoint from $\text{acl}(B)$ and $\text{acl}(B')$ over C , respectively. Hence, $\text{acl}(A) \cdot \text{acl}(B)$ is the field of fractions of $\text{acl}(A) \otimes_C \text{acl}(B)$, and similarly for A' and B' . This implies that $\alpha \otimes \beta$ induces a field isomorphism over C between $(\text{acl}(A) \cdot \text{acl}(B))^{alg} = \text{acl}(A, B)$ and $(\text{acl}(A') \cdot \text{acl}(B'))^{alg} = \text{acl}(A', B') = K_2$. This field isomorphism induces on K_2 the desired differential-difference field structure (Δ', σ') , and it follows that (K_0, Δ, σ) is a substructure of (K_2, Δ', σ') .

We now use the following result proved by Chatzidakis and Hrushovski (see the remark following the generalized independence theorem in [8]) to obtain that K_1 and K_2 are linearly disjoint over K_0 : If F_1, F_2, F_3 are algebraically closed fields extending an algebraically closed field F , with F_3 algebraically independent from $F_1 \cdot F_2$ over F , then $(F_1 \cdot F_3)^{alg} \cdot (F_2 \cdot F_3)^{alg}$ is linearly disjoint from $(F_1 \cdot F_2)^{alg}$ over $F_1 \cdot F_2$.

Now, since K_1 and K_2 are linearly disjoint over K_0 , their compositum is the field of fractions of $K_1 \otimes_{K_0} K_2$ and hence we can induce a differential-difference field structure $(\Delta'' = \{\delta''_1, \dots, \delta''_m\}, \sigma'')$ on $K_1 \cdot K_2$ by letting:

$$\sigma''(a \otimes b) = \sigma(a) \otimes \sigma'(b) \quad \text{and} \quad \delta''_i(a \otimes b) = \delta_i(a) \otimes b + a \otimes \delta'_i(b) \quad \text{for } i = 1, \dots, m.$$

It follows that (K_1, Δ, σ) and (K_2, Δ', σ') are substructures of $(K_1 \cdot K_2, \Delta'', \sigma'')$. We can embed the latter in some $(L, \Delta^*, \sigma^*) \models DCF_{0,m}A$, and note that it shares with $(\mathbb{U}, \Delta, \sigma)$ the algebraically closed substructure C . Therefore, by Proposition 5.3.1 (i), $L \equiv_C \mathbb{U}$ and so, by saturation of U , (L, Δ^*, σ^*) is embeddable in $(\mathbb{U}, \Delta, \sigma)$ over C . Replace A' and B' with their images under such an embedding (this images still satisfy conditions (i)-(iii)). Then it follows, from Proposition 5.3.1 (iii) and the isomorphism between (K_2, Δ', σ') and $(\text{acl}(A, B), \Delta, \sigma)$, that A' and B' satisfy condition (iv). \square

We now aim towards proving elimination of imaginaries for $DCF_{0,m}A$. We will follow very closely the strategy used by Moosa and Scanlon in [40] to prove elimination of imaginaries for fields with free operators in characteristic zero (in fact their strategy is an extension of the one used by Chatzidakis and Hrushovski in [8] to prove elimination of imaginaries for $ACFA$).

We will make use of the following notion of *dimension*:

Definition 5.3.4. Let a be a tuple from \mathbb{U} . We define

$$\dim_K a = (\text{trdeg}_K(\Theta_r a) : r < \omega)$$

where $\Theta_r a = \{\delta_m^{e_m} \dots \delta_1^{e_1} \sigma^k a : e_i, k \in \omega \text{ and } e_1 + \dots + e_m + k \leq r\}$. We view $\dim_K a$ as an element of ω^ω equipped with the lexicographical order.

This dimension can be considered as an analogue of the Kolchin polynomial (see Section 1.6). It is not preserved under interdefinability; however, it is a good measure for forking:

Lemma 5.3.5. *Suppose a is a tuple from \mathbb{U} and $L > K$ is a differential-difference field. Then $a \downarrow_K L$ if and only if $\dim_K a = \dim_L a$.*

It is worth pointing out here that our differential-difference fields do not fit into the formalism of Moosa and Scanlon from [40] so that the results of this chapter are not a consequence of [40].

Proof. It follows from (v) of Proposition 5.3.1 that $a \downarrow_K L$ if and only if $K\langle\sigma^k a : k \geq 0\rangle_\Delta$ is algebraically disjoint from L over K . Hence, we have the following

$$\begin{aligned}
a \downarrow_K L &\iff K\langle\sigma^k a : k \geq 0\rangle_\Delta \text{ is algebraically disjoint from } L \text{ over } K \\
&\iff K(\Theta_r a) \text{ is algebraically disjoint from } L \text{ over } K \text{ for all } r < \omega \\
&\iff \text{trdeg}_K(\Theta_r a) = \text{trdeg}_L(\Theta_r a) \text{ for all } r < \omega \\
&\iff \dim_K a = \dim_L a.
\end{aligned}$$

□

Proposition 5.3.6. *Every completion of $DCF_{0,m}A$ eliminates imaginaries.*

Proof. We will work in the saturated multi-sorted structure \mathbb{U}^{eq} . The unfamiliar reader is referred to Chapter 1.1 of [46] for definitions and basic results. Let $e \in \mathbb{U}^{eq}$. It suffices to show that there is a tuple d from \mathbb{U} such that $\text{dcl}^{eq}(e) = \text{dcl}^{eq}(d)$. Let $E = \text{acl}^{eq}(e) \cap \mathbb{U}$.

Claim. $e \in \text{dcl}^{eq}(E)$

Proof of Claim. Let f be a definable function (without parameters) and a a tuple from \mathbb{U} such that $e = f(a)$. We first show that there exists a tuple c from \mathbb{U} such that $f(c) = f(a)$ and $c \downarrow_E a$. By Lemma 1.4 of [11], there exists $b \models tp(a/E, e)$ such that

$$\text{acl}^{eq}(E, a) \cap \text{acl}^{eq}(E, b) \cap \mathbb{U} = E.$$

Let us prove that b can be chosen such that $\dim_{E\langle a \rangle_{\Delta, \sigma}} b$ is maximal (in the lexicographic order). For each r , choose b_r so that the above holds and $\text{trdeg}_{E\langle a \rangle_{\Delta, \sigma}}(\Theta_i b_r : i \leq r)$ is maximal possible. Let $n_r = \text{trdeg}_{E\langle a \rangle_{\Delta, \sigma}}(\Theta_r b_r)$ and let Φ be the set of formulas (over $E\langle a \rangle_{\Delta, \sigma}$) saying that

$$x \models tp(a/E, e), \quad \text{acl}^{eq}(E, a) \cap \text{acl}^{eq}(E, x) \cap \mathbb{U} = E,$$

and

$$\text{trdeg}_{E\langle a \rangle_{\Delta, \sigma}}(\Theta_r x) \geq n_r \text{ for each } r < \omega.$$

The b_r 's witness that Φ is finitely satisfiable, and so, by saturation, Φ is satisfiable. Any realisation b of Φ has the desired properties, i.e., $b \models tp(a/E, e)$, $\text{acl}^{eq}(E, a) \cap \text{acl}^{eq}(E, b) \cap \mathbb{U} = E$ and $\dim_{E\langle a \rangle_{\Delta, \sigma}} b$ is maximal.

Let $c \models tp(b/E, a)$ with $c \downarrow_{E, a} b$. Then $c \models tp(b/E, e) = tp(a/E, e)$ and so $f(c) = f(a)$. It remains to show that $c \downarrow_E a$. Since $c \downarrow_{E, a} b$, we have that

$$\text{acl}^{eq}(E, c) \cap \text{acl}^{eq}(E, b) \cap \mathbb{U} \subseteq \text{acl}^{eq}(E, a) \cap \text{acl}^{eq}(E, b) \cap \mathbb{U} = E.$$

Letting c' be such that $tp(b, c/E, e) = tp(a, c'/E, e)$ we have that $c' \models tp(a/E, e)$ and $\text{acl}^{eq}(E, c') \cap \text{acl}^{eq}(E, a) \cap \mathbb{U} = E$. Thus, by maximality, $\dim_{E\langle a \rangle_{\Delta, \sigma}} c' \leq \dim_{E\langle a \rangle_{\Delta, \sigma}} b$. Hence, as \dim is automorphism invariant, $\dim_{E\langle b \rangle_{\Delta, \sigma}} c \leq \dim_{E\langle a \rangle_{\Delta, \sigma}} b$. On the other hand,

$$\dim_{E\langle b \rangle_{\Delta, \sigma}} c \geq \dim_{E\langle a, b \rangle_{\Delta, \sigma}} c = \dim_{E\langle a \rangle_{\Delta, \sigma}} c = \dim_{E\langle a \rangle_{\Delta, \sigma}} b$$

where the first equality is by Lemma 5.3.5. Hence, $\dim_{E\langle a, b \rangle_{\Delta, \sigma}} c = \dim_{E\langle b \rangle_{\Delta, \sigma}} c$, and so, by Lemma 5.3.5, $c \perp_{E, b} a$. Since $c \perp_{E, a} b$ and $\text{acl}^{eq}(E, a) \cap \text{acl}^{eq}(E, b) \cap \mathbb{U} = E$, we get $c \perp_E a, b$. In particular, $c \perp_E a$, as desired.

Now, towards a contradiction, suppose $e \notin \text{dcl}^{eq}(E)$. Then there is $a' \models tp(a/E)$ such that $f(a) \neq f(a')$. Since $tp(a'/E) = tp(a/E)$, there is $c' \models tp(a'/E)$ such that $f(c') = f(a')$ and $c' \perp_E a'$. We may assume that $c' \perp_E c$. Now observe that we have the conditions of the independence theorem (see Fact 5.3.3 (6)); that is, $E = \text{acl}(E)$, $c \perp_E c'$, $tp(a/E) = tp(a'/E)$, $a \perp_E c$ and $a' \perp_E c'$. Hence, by the independence theorem, there is a $u \perp_E c, c'$ with $tp(u/E, c) = tp(a/E, c)$ and $tp(u/E, c') = tp(a'/E, c')$. But then $f(u) = f(c) = f(a)$ and $f(u) = f(c') = f(a')$, implying that $f(a) = f(a')$ which yields the contradiction. Hence, it must be the case that $e \in \text{dcl}^{eq}(E)$. This proves the claim.

Finally, let b be a tuple from E such that $e \in \text{dcl}^{eq}(b)$. Since $b \in \text{acl}^{eq}(e)$, b has only finitely many $\text{Aut}(\mathbb{U}^{eq}/e)$ -conjugates. Say $B = \{b_1, \dots, b_n\}$. Since ACF_0 admits elimination of imaginaries, there is a tuple d from \mathbb{U} such that any field automorphism of \mathbb{U} will fix d pointwise iff it fixes B setwise. We claim that $\text{dcl}^{eq}(e) = \text{dcl}^{eq}(d)$. Let $\phi \in \text{Aut}(\mathbb{U}^{eq}/e)$. Then, by definition of B , ϕ fixes B setwise and hence d pointwise. On the other hand, let $\phi \in \text{Aut}(\mathbb{U}^{eq}/d)$ and let $\psi(x, y)$ be an $\mathcal{L}_{m, \sigma}^{eq}$ -formula such that e is the only realisation of $\psi(x, b)$ (this formula exists since $e \in \text{dcl}^{eq}(b)$). Then $\phi(e)$ is the only realisation of $\psi(x, \phi(b))$; however, since $\phi(b) \in B$, e also realises $\psi(x, \phi(b))$. This implies that $\phi(e) = e$, as desired. \square

We end this section with the following description (up to acl) of the canonical base in $DCF_{0, m}A$:

Lemma 5.3.7. *Suppose a is a tuple such that $\sigma a, \delta a \in K(a)^{alg}$ for all $\delta \in \Delta$. Let $L > K$ be an algebraically closed differential-difference field, V the Zariski locus of a over L , and F the minimal field of definition of V . Then,*

$$\text{Cb}(a/L) \subseteq \text{acl}(F, K) \quad \text{and} \quad F \subseteq \text{Cb}(a/L).$$

In particular, $\text{acl}(\text{Cb}(a/L), K) = \text{acl}(F, K)$.

Proof. We have not defined in general the canonical base in this thesis (we refer the reader to [6] for the formal definition), but it suffices to say that by definition $a \downarrow_{Cb(a/L)} L$ and that to prove this lemma we only need to prove that:

(a) $a \downarrow_{\text{acl}(F,K)} L$, and

(b) V is defined over $Cb(a/L)$.

(a) As F is the minimal field of definition of $V = \text{loc}(a/L)$, a is algebraically disjoint from L over F . Then, as $\sigma a, \delta a \in K(a)^{\text{alg}}$ for all $\delta \in \Delta$, $K\langle a \rangle_{\Delta, \sigma} \subseteq K(a)^{\text{alg}}$ is algebraically disjoint from L over $\text{acl}(F, K)$. Hence, $a \downarrow_{\text{acl}(F,K)} L$.

(b) Since $a \downarrow_{Cb(a/L)} L$, then $\dim_L a = \dim_{Cb(a/L)} a$ (see Lemma 5.3.5). It follows that a is algebraically disjoint from L over $Cb(a/L)$. Hence, the Zariski locus of a over $Cb(a/L)$ must also be V , and so V is defined over $Cb(a/L)$. \square

5.4 Zilber dichotomy for finite dimensional types

In this section we prove the canonical base property, and consequently Zilber's dichotomy, for finite dimensional types in $DCF_{0,m}A$. Our proof follows the arguments given by Pillay and Ziegler in [52], where they prove the dichotomy for DCF_0 and $ACFA_0$ using suitable jet spaces (this is also the strategy of Bustamante in [2] to prove the dichotomy for finite dimensional types in $DCF_{0,1}A$).

Fix a sufficiently large saturated $(\mathbb{U}, \Delta, \sigma) \models DCF_{0,m}A$, and an algebraically closed differential-difference subfield K of \mathbb{U} . We first recall the theory of jet spaces from algebraic geometry. We refer the reader to §5 of [39] for a more detailed treatment of this classical material. The reason we are interested in jet spaces, which are higher order analogues of tangent spaces, is that they effect a linearisation of algebraic varieties.

Let V be an irreducible affine algebraic variety defined over K . Let $\mathbb{U}[V] = \mathbb{U}[x]/\mathcal{I}(V/\mathbb{U})$ denote the coordinate ring of V over \mathbb{U} . For each $a \in V$, let

$$\mathcal{M}_{V,a} = \{f \in \mathbb{U}[V] : f(a) = 0\}.$$

Let $r > 0$, the r -th jet space of V at a , denoted by $j_r V_a$, is defined as the dual space of the finite dimensional \mathbb{U} -vector space $\mathcal{M}_{V,a}/\mathcal{M}_{V,a}^{r+1}$.

The following gives explicit equations for the r -th jet space and allows us to consider it as an affine algebraic variety.

Fact 5.4.1. Let $V \subseteq \mathbb{U}^n$ be an irreducible affine algebraic variety defined over K . Fix $r > 0$. Let \mathcal{D} be the set of operator of the form

$$\frac{\partial^s}{\partial x_{i_1}^{s_1} \cdots \partial x_{i_k}^{s_k}}$$

where $0 < s \leq r$, $1 \leq i_1 < \cdots < i_k \leq n$, $s_1 + \cdots + s_k = s$, and $0 < s_i$. Let $a \in V$ and $d = |\mathcal{D}|$. Then $j_r V_a$ can be identified with the \mathbb{U} -vector subspace

$$\{(u_D)_{D \in \mathcal{D}} \in \mathbb{U}^d : \sum_{D \in \mathcal{D}} Df(a)u_D = 0, \text{ for all } f \in \mathcal{I}(V/K)\}.$$

Let $X \subseteq V$ be an irreducible algebraic subvariety and $a \in X$. The containment of X in V yields a canonical linear embedding of $j_r X_a$ into $j_r V_a$ for all r . We therefore identify $j_r X_a$ with its image in $j_r V_a$.

The following fact makes precise what we said earlier about linearisation of algebraic varieties.

Fact 5.4.2. Let X and Y be irreducible algebraic subvarieties of V and $a \in X \cap Y$. Then, $X = Y$ if and only if $j_r X_a = j_r Y_a$, as subspaces of $j_r V_a$, for all r .

We want to develop a (Δ, σ) analogue of the notions of differential and difference modules from [52].

Definition 5.4.3. By a (Δ, σ) -module over $(\mathbb{U}, \Delta, \sigma)$ we mean a triple (E, Ω, Σ) such that E is a finite dimensional \mathbb{U} -vector space, $\Omega = \{\partial_1, \dots, \partial_m\}$ is a family of additive endomorphisms of E and Σ is an additive automorphism of E such that

$$\partial_i(\alpha e) = \delta_i(\alpha)e + \alpha \partial_i(e)$$

and

$$\Sigma(\alpha e) = \sigma(\alpha)\Sigma(e)$$

for all $\alpha \in \mathbb{U}$ and $e \in E$, and the operators in $\Omega \cup \{\Sigma\}$ commute. If we omit σ and Σ we obtain Pillay and Ziegler's definition of a Δ -module over (\mathbb{U}, Δ) . Similarly, if we omit Δ and Ω we obtain the definition of a σ -module over (\mathbb{U}, σ) .

The following is for us the key property of (Δ, σ) -modules.

Lemma 5.4.4. *Let (E, Ω, Σ) be a (Δ, σ) -module over \mathbb{U} . Let*

$$E^\# = \{e \in E : \Sigma(e) = e \text{ and } \partial(e) = 0 \text{ for all } \partial \in \Omega\}.$$

Then $E^\#$ is a $\mathcal{C}_{\mathbb{U}}$ -vector space (recall that $\mathcal{C}_{\mathbb{U}} = \mathbb{U}^\sigma \cap \mathbb{U}^\Delta$) and there is a $\mathcal{C}_{\mathbb{U}}$ -basis of $E^\#$ which is also a \mathbb{U} -basis of E .

Proof. Since $\partial_1, \dots, \partial_m$ and Σ are $\mathcal{C}_{\mathbb{U}}$ -linear, $E^\#$ is a $\mathcal{C}_{\mathbb{U}}$ -vector space. Let $\{e_1, \dots, e_d\}$ be a \mathbb{U} -basis of E . With respect to this basis, let A_i be the matrix of ∂_i , $i = 1, \dots, m$, and B the matrix of Σ . By this we mean that A_i is the matrix whose j -th column consists of the coefficients of the linear combination of $\partial_i(e_j)$ in terms of the basis, and similarly for B . Under the linear transformation that takes the basis $\{e_1, \dots, e_d\}$ to the standard basis of \mathbb{U}^d , the (Δ, σ) -module (E, Ω, Σ) is transformed into the (Δ, σ) -module

$$(\mathbb{U}^d, \{\delta_1 + A_1, \dots, \delta_m + A_m\}, B\sigma).$$

It suffices to prove the result for this (Δ, σ) -module. As Σ is an additive automorphism of E , the matrix B is invertible. Also, the commutativity of $\Omega \cup \{\Sigma\}$ yields:

$$\delta_j A_i - \delta_i A_j = [A_i, A_j], \quad i = 1, \dots, m \quad (5.4.1)$$

and

$$B\sigma(A_i) = \delta_i(B) + A_i B, \quad i = 1, \dots, m. \quad (5.4.2)$$

Since $\delta_i(B^{-1}) = -B^{-1}\delta_i(B)B^{-1}$, the previous equation yields

$$B^{-1}A_i = \delta_i(B^{-1}) + \sigma(A_i)B^{-1}, \quad i = 1, \dots, m. \quad (5.4.3)$$

Now, note that $(\mathbb{U}^d)^\# = \{u \in \mathbb{U}^d : B\sigma u = u \text{ and } \delta_i u + A_i u = 0, i = 1, \dots, m\}$, and thus it suffices to find a nonsingular $n \times n$ matrix M over \mathbb{U} such that $B\sigma(M) = M$ and $\delta_i M + A_i M = 0$. Let X and Y be $n \times n$ matrices of variables. It follows from (5.4.1), and Rosenfeld's criterion (see Fact 1.2.6), that the set $\Lambda = \{\delta_i X + A_i X : i = 1, \dots, m\}$ is a characteristic set of a prime Δ -ideal of $\mathbb{U}\{X\}_\Delta$. Since in this case $H_\Lambda = 1$, the corresponding prime Δ -ideal is simply

$$\mathcal{P} = [\delta_i X + A_i X : i = 1, \dots, m] \subset \mathbb{U}\{X\}_\Delta.$$

Then $V = \mathcal{V}(\mathcal{P})$ is an irreducible Δ -variety. Let W be the irreducible Δ -variety defined by

$$X = BY \quad \text{and} \quad Y \in V^\sigma.$$

Then, by (5.4.2), $W \subseteq V \times V^\sigma$. Clearly, W projects onto V^σ and, by (5.4.3), it also projects onto V . Hence, by Fact 5.2.1, there is a nonsingular matrix M over \mathbb{U} such that $M \in V$ and $(M, \sigma M) \in W$. This M satisfies the desired properties. \square

Notice that if $\{e_1, \dots, e_d\} \subset E^\#$ is a \mathbb{U} -basis of E , which exists by the previous lemma, then, under the linear transformation that takes this basis to the standard basis of \mathbb{U}^d , the (Δ, σ) -module (E, Ω, Σ) is transformed into the (Δ, σ) -module $(\mathbb{U}^d, \Delta, \sigma)$.

Remark 5.4.5. Let (E, Ω) be a Δ -module over (\mathbb{U}, Δ) and E^* be the dual space of E . If we define the *dual operators* $\Omega^* = \{\partial_1^*, \dots, \partial_m^*\}$ on E^* by

$$\partial_i^*(\lambda)(e) = \delta_i(\lambda(e)) - \lambda(\partial_i(e))$$

for all $\lambda \in E^*$ and $e \in E$, then (E^*, Ω^*) becomes a Δ -module over \mathbb{U} . Indeed, by Remark 3.3 of [52], (E^*, ∂_i^*) is a $\{\delta_i\}$ -module. Hence, all we need to verify is that the ∂_i^* 's commute:

$$\begin{aligned} \partial_j^*(\partial_i^*(\lambda))(e) &= \delta_j(\delta_i(\lambda(e))) - \delta_j(\lambda(\partial_i(e))) - \delta_i(\lambda(\partial_j(e))) + \lambda(\partial_i(\partial_j(e))) \\ &= \delta_i(\delta_j(\lambda(e))) - \delta_i(\lambda(\partial_j(e))) - \delta_j(\lambda(\partial_i(e))) + \lambda(\partial_j(\partial_i(e))) \\ &= \partial_i^*(\partial_j^*(\lambda))(e). \end{aligned}$$

We now describe a natural (Δ, σ) -module structure on the jet spaces of an algebraic D-variety equipped with a finite-to-finite correspondence with its σ -transform.

By a D-variety we will mean a relative D-variety w.r.t. Δ/\emptyset in the sense of Section 2.4. That is, the classical non relative notion studied by Buium [1]. Let (V, s) be an irreducible affine D-variety defined over K . Then we can extend the derivations Δ to the coordinate ring $\mathbb{U}[V]$ by defining $\delta_i(x) = s_i(x)$ where $s = (\text{Id}, s_1, \dots, s_m)$ and $x = (x_1, \dots, x_n)$ are the coordinate functions of $\mathbb{U}[V]$. The integrability condition (2.4.1) shows that these extensions commute with each other. Indeed, we have

$$\delta_j \delta_i(x) = \delta_j s_i(x) = d_{\delta_j/\emptyset} s_i(x, s_j(x)) = d_{\delta_i/\emptyset} s_j(x, s_i(x)) = \delta_i s_j(x) = \delta_i \delta_j(x).$$

Hence, $(\mathbb{U}[V], \Delta)$ becomes a Δ -ring (having a Δ -ring structure on the coordinate ring is the approach of Buium to D-varieties [1]).

Let $a \in (V, s)^\#$. Then, $\mathcal{M}_{V,a}^r$ is a δ_i -ideal of the δ_i -ring $\mathbb{U}[V]$ for all $r > 0$. This is shown explicitly in Lemma 3.7 of [52]. Hence, $(\mathcal{M}_{V,a}/\mathcal{M}_{V,a}^r, \Delta)$ becomes a Δ -module over (\mathbb{U}, Δ) . By Remark 5.4.5, $(j_r V_a, \Delta^*)$ is a Δ -module over (\mathbb{U}, Δ) .

Suppose now $a \in (V, s)^\#$ is a generic point of V over K and $\sigma a \in K(a)^{alg}$. Let W be the Zariski locus of $(a, \sigma a)$ over K . Then $W \subseteq V \times V^\sigma$ projects dominantly and generically finite-to-one onto both V and V^σ . Moreover, for each $r > 0$, $j_r W_{(a, \sigma a)} \subseteq j_r V_a \times j_r V_{\sigma a}^\sigma$ is the graph of an isomorphism $f : j_r V_a \rightarrow j_r V_{\sigma a}^\sigma$ and the map $\sigma^* = f^{-1} \circ \sigma$ equips $j_r V_a$ with the structure of a σ -module over (\mathbb{U}, σ) (see Lemma 4.3 of [52] for details). Furthermore, Lemma 4.4 of [2] shows that $(j_r V_a, \delta_i^*, \sigma^*)$ is a (δ_i, σ) -module over $(\mathbb{U}, \delta_i, \sigma)$ for all $i = 1, \dots, m$. Thus, since we have already seen that the dual operators Δ^* commute, $(j_r V_a, \Delta^*, \sigma^*)$ is a (Δ, σ) -module over $(\mathbb{U}, \Delta, \sigma)$.

Remark 5.4.6. Let V be an irreducible affine D-variety defined over K and suppose that $a \in (V, s)^\#$ is a generic point of V over K such that $\sigma a \in K(a)^{alg}$. Suppose $L > K$ is an algebraically closed differential-difference field and W is the Zariski locus of a over L . Then $(W, s|_W)$ is a D -subvariety of (V, s) and, under the identification of $j_r W_a$ as a subspace of $j_r V_a$, we have that $j_r W_a$ is a (Δ, σ) -submodule of $(j_r V_a, \Delta^*, \sigma^*)$. Indeed, by Lemma 4.7 of [2], $j_r W_a$ is a (δ_i, σ) -submodule of $(j_r V_a, \delta_i^*, \sigma^*)$ for all $i = 1, \dots, r$.

Definition 5.4.7. A type $p = tp(a/K)$ is said to be *finite dimensional* if the transcendence degree of $K\langle a \rangle_{\Delta, \sigma}$ over K is finite.

Lemma 5.4.8. *Suppose $p = tp(a/K)$ is finite dimensional. Then there is an irreducible affine D-variety (V, s) over K and a generic point $c \in (V, s)^\#$ of V over K such that $K\langle a \rangle_{\Delta, \sigma} = K\langle c \rangle_{\Delta, \sigma}$ and $\sigma c \in K(c)^{alg}$.*

Proof. Since p is finite dimensional then there is $s < \omega$ such that $K\langle a \rangle_{\Delta, \sigma} \subseteq K(\Theta_s a)^{alg}$ where (as in Definition 5.3.5)

$$\Theta_s a = \{\delta_m^{e_m} \cdots \delta_1^{e_1} \sigma^k a : e_i, k \in \omega \text{ and } e_1 + \cdots + e_m + k \leq s\}.$$

In particular, $\sigma^{s+1} a \in K(\Theta_s a)^{alg}$. Hence, if we let $b = (a, \sigma a, \dots, \sigma^s a)$, then $K\langle a \rangle_{\Delta, \sigma} = K\langle b \rangle_{\Delta, \sigma}$ and $\sigma b \in K\langle b \rangle_{\Delta}^{alg}$. Also, since $K\langle b \rangle_{\Delta} \subseteq K(\Theta_s b)$, we have that $\Delta\text{-type}(b/K) = 0$. So, by Proposition 2.4.6, there is an irreducible affine D-variety (V, s) over K and a generic point $c \in (V, s)^\#$ of V over K such that $K\langle b \rangle_{\Delta} = K(c)$. Hence, $K\langle a \rangle_{\Delta, \sigma} = K\langle c \rangle_{\Delta, \sigma}$ and $\sigma c \in K(c)^{alg}$. \square

Let us make explicit what the notion of almost internality means in $DCF_{0,m}A$. We say that $p = tp(a/K)$ is *almost \mathcal{C}_U -internal* if there is a differential-difference field extension L of K with $a \downarrow_K L$ such that $a \in \text{acl}(L, c) = L\langle c \rangle_{\Delta, \sigma}^{alg}$ for some tuple c from \mathcal{C}_U .

We are now in position to prove the canonical base property for finite dimensional types.

Theorem 5.4.9. *Suppose $tp(a/K)$ is finite dimensional and $L > K$ is an algebraically closed differential-difference field. Let c be a tuple such that $\text{acl}(c) = \text{acl}(Cb(a/L))$. Then $tp(c/K\langle a \rangle_{\Delta, \sigma})$ is almost \mathcal{C}_U -internal.*

Proof. We may replace a by anything interdefinable with it over K . Hence, by Lemma 5.4.8, we may assume that $\sigma a \in K(a)^{alg}$ and that there is an irreducible affine D-variety (V, s) defined over K such that $a \in (V, s)^\#$ is a generic point of V over K . Let W be the locus of a over L . Then $(W, s|_W)$ is a D -subvariety of (V, s) and $a \in (W, s|_W)^\#$. By Lemma 5.3.7,

if d is a tuple generating the minimal field of definition of W then $\text{acl}(d, K) = \text{acl}(c, K)$. Thus, it suffices to prove that $tp(d/K\langle a \rangle_{\Delta, \sigma})$ is almost $\mathcal{C}_{\mathbb{U}}$ -internal.

Consider the (Δ, σ) -module $(j_r V_a, \Delta^*, \sigma^*)$ and recall that $j_r W_a$ is a (Δ, σ) -submodule. For each r , let b_r a $\mathcal{C}_{\mathbb{U}}$ -basis of $j_r V_a^\#$ which is also a \mathbb{U} -basis of $j_r V_a$. Let $B = \cup_{r=1}^\infty b_r$, we may choose the b_r 's such that $d \downarrow_{K\langle a \rangle_{\Delta, \sigma}} B$. The basis b_r yields a (Δ, σ) -module isomorphism between $(j_r V_a, \Delta^*, \sigma^*)$ and $(\mathbb{U}^{d_r}, \Delta, \sigma)$ which therefore takes $j_r W_a$ into a (Δ, σ) -submodule S_r of $(\mathbb{U}^{d_r}, \Delta, \sigma)$. We can find a $\mathcal{C}_{\mathbb{U}}$ -basis e_r of $S_r^\# \subseteq \mathcal{C}_{\mathbb{U}}^{d_r}$ which is also a \mathbb{U} -basis of S_r , so S_r is defined over $e_r \subset \mathcal{C}_{\mathbb{U}}^{d_r}$. Let $E = \cup_{r=1}^\infty e_r$.

It suffices to show that $d \in \text{dcl}(a, K, B, E)$. To see this, let ϕ be an automorphism of $(\mathbb{U}, \Delta, \sigma)$ fixing a, K, B, E pointwise. Since $j_r V_a$ is defined over $K\langle a \rangle_{\Delta, \sigma}$, then $\phi(j_r V_a) = j_r V_a$. Also, as each S_r is defined over E and the isomorphism between S_r and $j_r W_a$ is defined over B , $\phi(j_r W_a) = j_r W_a$ for all $r > 0$. Since V is defined over K and ϕ fixes a pointwise, Fact 5.4.2 implies that $\phi(W) = W$. But d is in the minimal field of definition of W , thus ϕ fixes d pointwise, as desired. \square

Remark 5.4.10. Even though we did not mention it explicitly in the proof of Theorem 5.4.9, the key construction in it is that of

$$\begin{aligned} (j_r W_a, \Delta^*, \sigma^*)^\# &= \{\lambda \in j_r W_a : \sigma^*(\lambda) = \lambda \text{ and } \delta_i^*(\lambda) = 0, i = 1, \dots, m\} \\ &= \{\lambda \in j_r W_a : \sigma(\lambda) = f(\lambda) \text{ and } \delta_i \circ \lambda = \lambda \circ \delta_i, i = 1 \dots, m\}, \end{aligned}$$

which is a finite dimensional $\mathcal{C}_{\mathbb{U}}$ -vector space. We could define this vector space to be the r -th (Δ, σ) -jet space at a of the (Δ, σ) -locus of a over L . But (Δ, σ) -jet spaces for arbitrary (Δ, σ) -varieties have already been defined, they are a special case of the generalized Hasse-Schmidt jet spaces of Moosa and Scanlon from [38]. We expect that in this context (namely a is a Δ -generic point of the sharp points of an algebraic D-variety over L and $\sigma a \in L(a)^{\text{alg}}$) the two constructions agree. This has been checked in the differential case, see the end of §4.3 of [38]. But we have not checked our case here, because we can prove the desired results directly.

As a consequence of the canonical base property, we can prove Zilber dichotomy for finite dimensional types. But first let us give the definition of a one-based type in the context of $DCF_{0,m}A$. A type $p = tp(a/K)$ is said to be *one-based* if for every algebraically closed differential-difference field extension L of K we have that $Cb(a/L) \subseteq \text{acl}(K, a) = K\langle a \rangle_{\Delta, \sigma}^{\text{alg}}$.

Corollary 5.4.11. *Let p be a finite dimensional complete type over K with $SU(p) = 1$. Then p is either one-based or almost $\mathcal{C}_{\mathbb{U}}$ -internal.*

Proof. Suppose $p = tp(a/K)$ is not one-based. Then there is an algebraically closed differential-difference field $L > K$ such that $Cb(a/L)$ is not contained in $\text{acl}(K, a)$. Let c be a tuple such that $\text{acl}(c) = \text{acl}(Cb(a/L))$, then $c \notin \text{acl}(K, a)$. By Theorem 5.4.9, $tp(c/K\langle a \rangle_{\Delta, \sigma})$ is almost \mathcal{C}_U -internal. Hence, $c \in \text{acl}(K, a, b, e)$ where $c \perp_{K, a} b$ and e is a tuple from \mathcal{C}_U . Therefore we get $c \not\perp_{K, a, b} e$.

Now, by a general fact about canonical bases in supersimple theories we get that $c \in \text{acl}(a_1, \dots, a_s)$ where the a_i 's are independent realisations of p (see Chap. 17 of [6]). We may assume the a_i 's are independent from a, b over K . We get that

$$a_1, \dots, a_s \not\perp_{K, a, b} e,$$

and so, since $SU(p) = 1$, $a_{i+1} \in \text{acl}(K, a, b, a_1, \dots, a_i, e)$ for some $i < s$. Since

$$a_{i+1} \perp_K a, b, a_1, \dots, a_i$$

and e is a tuple from \mathcal{C}_U , p is almost \mathcal{C}_U -internal. □

Remark 5.4.12. The assumption in Corollary 5.4.11 that p is finite dimensional should not be necessary, though a different proof is needed. To prove the general case one could work out the theory of arc spaces of Moosa, Pillay, Scanlon [37] in the (Δ, σ) setting, and apply their weak dichotomy for regular types. This is done in the (δ, σ) setting (i.e., ordinary differential-difference fields) by Bustamante in [5], and the arguments there should work here as well.

Appendix

Some U -rank computations in $DCF_{0,m}$

In this appendix, we present some results on computing and finding lower bounds for the U -rank of certain definable subgroups of the additive group. We work inside a sufficiently saturated $(\mathbb{U}, \Delta) \models DCF_{0,m}$ and fix a base Δ -field K .

Let $p \in S_n(K)$, $\tau = \Delta\text{-type}(p)$ and $d = \Delta\text{-dim}(p)$. In [35], McGrail proved that

$$U(p) < \omega^\tau(d + 1). \quad (5.4.4)$$

She also claimed the lower bound $\omega^\tau d \leq U(p)$. However, in [57], Suer shows that this lower bound is in general false. He showed that, in $DCF_{0,2}$, if p is the generic type of $\delta_2^k x = \delta_1 x$ then $U(p) = \omega$ while $\Delta\text{-type}(p) = 1$ and $\Delta\text{-dim}(p) = k$ for all $k > 0$. So far there is no known lower bound for the U -rank in terms of the Δ -type; in fact, it is not even known if $U(p) = 1$ implies $\Delta\text{-type}(p) = 0$.

On the other hand, it is generally known that the generic type of $DCF_{0,m}$ has U -rank ω^m . However, the only written argument (for $m > 1$) found by the author uses McGrail's incorrect lower bound. We record a correct (and very standard) proof.

Lemma 5.4.13. *If $p \in S_1(K)$ is the generic type of \mathbb{U} over K then $U(p) = \omega^m$.*

Proof. For $i = 0, 1, \dots, m$, let $\Delta_i = \{\delta_1, \dots, \delta_{m-i}\}$. Hence, $\mathbb{U}^{\Delta_0} = \mathbb{U}^\Delta$ and $\mathbb{U}^{\Delta_m} = \mathbb{U}$. We show, by induction on i , that $U(\mathbb{U}^{\Delta_i}) = \omega^i$. The base case holds trivially, as in this case we are dealing with the (total) constants \mathbb{U}^Δ . Now suppose the result holds for i . Let $n > 0$ and consider the definable group

$$G = \{a \in \mathbb{U}^{\Delta_{i+1}} : \delta_{m-i}^n a = 0\},$$

which is an n dimensional \mathbb{U}^{Δ_i} -vector subspace of $\mathbb{U}^{\Delta_{i+1}}$. By the induction hypothesis $U(G) = \omega^i n$. As n was arbitrary and $G \subset \mathbb{U}^{\Delta_{i+1}}$, we get that $U(\mathbb{U}^{\Delta_{i+1}}) \geq \omega^{i+1}$. Now we need to show that $U(\mathbb{U}^{\Delta_{i+1}}) \leq \omega^{i+1}$; that is, we have to show that for every forking extension $tp(a/L)$ of the generic type of $\mathbb{U}^{\Delta_{i+1}}$ over K we have that $U(a/L) < \omega^{i+1}$. But, since $a \in \mathbb{U}^{\Delta_{i+1}}$ is $\Delta \setminus \Delta_i$ -algebraic over L (as $tp(a/L)$ forks over K), Δ -type(a/L) $\leq i$. Hence, by (5.4.4), $U(a/L) < \omega^{i+1}$. This completes the induction.

Applying this to the case $i = m$ shows that the generic type of \mathbb{U} over K is ω^m . \square

Let us observe that the idea here was to find a family of definable subgroups of the additive group $(\mathbb{U}, +)$ whose image under the U -rank map was a cofinal subset of ω^m . This cofinal subset was the set of ordinals of the form $\omega^i n$. In the following example we observe that, through a more detailed argument, one can produce a family of subgroups of the additive group such that for every $\alpha < \omega^m$ there is an element in the family whose U -rank is “close” to α .

Example 5.4.14. We construct a family $(G_r)_{r \in \mathbb{N}^m}$ of definable subgroups of the additive group with the following property. For every α of the form

$$\alpha = \omega^n a_n + \omega^{n-1} a_{n-1} + \cdots + \omega a_1 + a_0$$

with $n < m$, $a_i < \omega$ and $a_n \neq 0$, there is $r \in \mathbb{N}^m$ such that

$$\alpha \leq U(G_r) < \omega^n (a_n + 1).$$

Let G_r be defined by the homogeneous system of linear differential equations

$$\delta_1^{r_1+1} x = 0, \delta_2^{r_2+1} \delta_1^{r_1} x = 0, \dots, \delta_{m-1}^{r_{m-1}+1} \delta_{m-2}^{r_{m-2}} \cdots \delta_1^{r_1} x = 0, \delta_m^{r_m} \delta_{m-1}^{r_{m-1}} \cdots \delta_1^{r_1} x = 0.$$

We first show that

$$U(G_r) \geq \omega^{m-1} r_1 + \omega^{m-2} r_2 + \cdots + \omega r_{m-1} + r_m.$$

We prove this by transfinite induction on $r = (r_1, \dots, r_m)$ in the lexicographical order. The base case holds trivially. Suppose first that $r_m \neq 0$ (i.e., the successor ordinal case). Consider the definable group homomorphism $f : (G_r, +) \rightarrow (G_r, +)$ given by $f(a) = \delta_m^{r_m-1} a$. Then the generic type of the generic fibre of f is a forking extension of the generic type of G_r . Since f is a definable group homomorphism, the U -rank of the generic fibre is the same as the U -rank of $\text{Ker}(f) = G_{r'}$ where $r' = (r_1, \dots, r_{m-1}, r_m - 1)$. By induction,

$$U(G_{r'}) \geq \omega^{m-1} r_1 + \omega^{m-2} r_2 + \cdots + \omega r_{m-1} + (r_m - 1).$$

Hence,

$$U(G_r) \geq \omega^{m-1}r_1 + \omega^{m-2}r_2 + \cdots + \omega r_{m-1} + r_m.$$

Now suppose $r_m = 0$ (i.e., the limit ordinal case) and that k is the largest such that $r_k \neq 0$. Let $n \in \omega$ and $r' = (r_1, \dots, r_k - 1, n, 0, \dots, 0)$. Then $G_{r'} \subset G_r$ and, by induction,

$$U(G_{r'}) \geq \omega^{m-1}r_1 + \omega^{m-2}r_2 + \cdots + \omega^{m-k}(r_k - 1) + \omega^{m-k-1}n.$$

Since n was arbitrary,

$$U(G_r) \geq \omega^{m-1}r_1 + \omega^{m-2}r_2 + \cdots + \omega^{m-k}r_k.$$

This complete the induction.

Next we show that if k is the smallest such that $r_k > 0$ then $\Delta\text{-type}(G) = m - k$ and $\Delta\text{-dim}(G) = r_k$. Let $tp(a/K)$ be the generic type of G_r . As $r_1 = \cdots = r_{k-1} = 0$, we have $\delta_1 a = 0, \dots, \delta_{k-1} a = 0, \delta_k^{r_k+1} a = 0$ and $\delta_k^{r_k} a$ is Δ'_k -algebraic over K where $\Delta'_k = \{\delta_{k+1}, \dots, \delta_m\}$. It suffices to show that $a, \delta_k a, \dots, \delta_k^{r_k-1} a$ are Δ'_k -algebraically independent over K . Let $f \in K\{x_0, \dots, x_{k-1}\}_{\Delta'_k}$ be nonzero, and let $g(x) = f(x, \delta_k x, \dots, \delta_k^{r_k-1} x) \in K\{x\}_{\Delta}$. Then g is a nonzero Δ -polynomial over K reduced with respect to all $f \in \mathcal{I}(G_r/K)_{\Delta}$. Thus, as a is a generic point of G_r over K ,

$$0 \neq g(a) = f(a, \delta_k a, \dots, \delta_k^{r_k-1} a),$$

as desired.

Applying this, and (5.4.4), in the case when $r_1 = 0, \dots, r_{m-n-1} = 0, r_{m-n} = a_n, r_{m-n+1} = a_{n-1}, \dots, r_{m-1} = a_1, r_m = a_0$, we obtain the desired inequalities

$$\alpha \leq U(G_r) < \omega^n(a_n + 1).$$

Note that, in particular, each G_r satisfies McGrail's lower bound.

References

- [1] A. Buium. *Differential algebraic groups of finite dimension*. Lecture notes in mathematics, Vol. 1506, Springer-Verlag, 1992.
- [2] R. Bustamante. *Algebraic jet spaces and Zilber's dichotomy in DCFA*. Revista de Matemática: Teoría y Aplicaciones 17 (1), pp.1-12, 2010.
- [3] R. Bustamante. *Differentially closed fields of characteristic zero with a generic automorphism*. Revista de Matemática: Teoría y Aplicaciones 14(1), pp.81-100, 2007.
- [4] R. Bustamante. *Rank and dimension in difference-differential fields*. Notre Dame Journal of Formal Logic, Vol. 52 (4), pp.403-414, 2011.
- [5] R. Bustamante. *Théorie des modèles des corps différentiellement clos avec un automorphisme générique*. PhD Thesis. Université Paris 7, 2005.
- [6] E. Casanovas. *Simple theories and hyperimaginaries*. Lecture Notes in Logic. Cambridge University Press, 2011.
- [7] P. Cassidy and M. Singer. *Galois theory of parametrized differential equations and linear differential algebraic groups*. Differential Equations and Quantum Groups (IRMA Lectures in Mathematics and Theoretical Physics Vol. 9). EMS Publishing house, pp.113- 157, 2006.
- [8] Z. Chatzidakis and E. Hrushovski. *Model theory of difference fields*. Trans. Amer. Math. Soc.351 (8), pp.2997-3071, 1999.
- [9] Z. Chatzidakis and A. Pillay. *Generic structures and simple theories*. Annals of Pure and Applied Logic 95, pp.71-92, 1998.
- [10] R. Cohn. *Difference algebra*. Interscience tracts in pure and applied mathematics, No. 17, 1965.

- [11] D. Evans and E. Hrushovski. *On the automorphism groups of finite covers*. Annals of Pure and Applied Logic 62, pp.83-112, 1993.
- [12] J. Freitag. *Indecomposability for differential algebraic groups*. Preprint 2011 (<http://arxiv.org/pdf/1106.0695v2.pdf>).
- [13] J. Freitag. *Model theory and differential algebraic geometry*. PhD Thesis. University of Illinois at Chicago, 2012.
- [14] O. Golubitsky, M. Kondratieva, and A. Ovchinnikov. *On the Generalised Ritt Problem as a Computational Problem*. Journal of Mathematical Sciences. Vol. 163, No. 5, pp. 515-522 (2009).
- [15] N. Guzy and C. Riviere. *On existentially closed partial differential fields with an automorphism*. Preprint 2007 ([citeseerx.ist.psu.edu/viewdoc/download?doi](http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.1.1.1.1.1.1.1)).
- [16] M. Harrison-Trainor, J. Klys, and R. Moosa. *Non standard methods for bounds in differential polynomial rings*. Journal of Algebra 360, pp.71-86, 2012.
- [17] B. Hart, B. Kim, and A. Pillay. *Coordinatisation and canonical basis in simple theories*. Journal of Symbolic Logic. Vol. 65, No. 1, pp.293-309, 2000.
- [18] M. Hils. *Generic automorphisms and green fields*. Preprint 2011 (<http://arxiv.org/pdf/1003.5149.pdf>).
- [19] E. Hrushovski. *The Mordell-Lang conjecture for function fields*. Journal of the American Mathematical Society. Vol. 9, No. 3, 1996.
- [20] E. Hrushovski. *Unidimensional theories are superstable*. Annals of Pure and Applied Logic 50, pp.117-138, 1990.
- [21] I. Kaplansky. *An introduction to differential algebra*. L'institut de mathématique de l'université de Nancago. Hermann, 1957.
- [22] B. Kim and A. Pillay. *Simple theories*. Annals of Pure and Applied Logic. Vol. 88, pp.149-164, 1997.
- [23] E. R. Kolchin. *Algebraic matrix groups and the Picard-Vessiot theory of homogeneous linear ordinary differential equations*. Annals of Mathematics, Vol. 49, No. 1, 1948.
- [24] E. R. Kolchin. *Differential algebra and algebraic groups*. Academic Press. New York, New York 1973.

- [25] E. R. Kolchin. *Differential algebraic groups*. Academic Press, Inc. 1985.
- [26] E. R. Kolchin. *On the basis theorem for infinite systems of differential polynomials*. Bull. Amer. Math. Soc. 45, pp.923-926, 1939.
- [27] J. Kovacic. *Differential Schemes*. Differential algebra and related topics. Proc. Internat. Workshop. pp.71-94, 2002.
- [28] P. Kowalski. *Geometric axioms for existentially closed Hasse fields*. Annals of Pure and Applied Logic 135, pp.286-302, 2005.
- [29] P. Kowalski and A. Pillay. *Quantifier elimination for algebraic D -groups*. Trans. Amer. Math. Soc. 358, pp.167-181, 2005.
- [30] P. Landesman. *Generalized differential Galois theory*. Trans. Amer. Math.Soc. 360, pp.4441-4495, 2008.
- [31] D. Marker. *Embedding differential algebraic groups in algebraic groups*. www.sci.cny.cuny.edu/ksda/PostedPapers/Marker050109.pdf.
- [32] D. Marker. *Manin kernels*. Connections between model theory and algebraic and analytic geometry, pp.1-21, 2000.
- [33] D. Marker. *Model Theory: An Introduction*. Springer-Verlag. New York, Inc, 2002.
- [34] D. Marker, M. Messmer, and A. Pillay. *Model theory of fields*. Lecture Notes in Logic 5, ASL, 2005.
- [35] T. McGrail. *The model theory of differential fields with finitely many commuting derivations*. The Journal of Symbolic Logic. Vol. 65, No. 2, pp.885-913, 2000.
- [36] R. Moosa. *Model theory and complex geometry*. Notices of the American Mathematical Society. Vol. 57, No. 2, pp.230-235, 2010.
- [37] R. Moosa, A. Pillay, and T. Scanlon. *Differential arcs and regular types in differential fields*. J. reine angew. Math. 620, pp.35-54, 2008.
- [38] R. Moosa and T. Scanlon. *Generalized Hasse-Schmidt varieties and their jet spaces*. Proc. London Math. Soc. 3, pp.197-234, 2011.
- [39] R. Moosa and T. Scanlon. *Jet and prolongation spaces*. Journal of the Inst. of Math. Jussieu 9, pp.391-430, 2010.

- [40] R. Moosa and T. Scanlon. *Model theory of fields with free operators in characteristic zero*. Preprint 2012 (<http://arxiv.org/pdf/1212.5838v1.pdf>).
- [41] D. Pierce. *Fields with several commuting derivations*. Preprint, 2007, <http://arxiv.org/pdf/0708.2769v1.pdf>.
- [42] D. Pierce. *Geometric characterizations of existentially closed fields with operators*. Illinois J. Math., Vol. 48, No. 4, pp.1321-1343, 2004.
- [43] D. Pierce and A. Pillay. *A note on the axioms for differentially closed fields of characteristic zero*. Journal of Algebra 204, pp.108-115, 1998.
- [44] A. Pillay. *Algebraic D-groups and differential Galois theory*. Pacific Journal of Mathematics. Vol. 216, No. 2, 2004.
- [45] A. Pillay. *Differential Galois theory I*. Illinois Journal of Mathematics. Vol. 42, N. 4, 1998.
- [46] A. Pillay. *Geometric stability theory*. Oxford Science Publications, 1996.
- [47] A. Pillay. *An introduction to stability theory*. Claredon Press. Oxford 1983.
- [48] A. Pillay. *Remarks on algebraic D-varieties and the model theory of differential fields*. Lecture Notes in Logic 26, Logic in Tehran, pp.256-269, 2006.
- [49] A. Pillay. *Remarks on Galois cohomology and definability*. The Journal of Symbolic Logic. Vol. 62, No. 2, pp. 487-492, 1997.
- [50] A. Pillay. *Some foundational questions concerning differential algebraic groups*. Pacific Journal of Mathematics. Vol. 179, No.1, 1997.
- [51] A. Pillay. *Two remarks on differential fields*. Model Theory and Applications. Quaderni di Matematica. Volume 11, 2003. Edited by Dipartimento di Matematica, Seconda Università di Napoli.
- [52] A. Pillay and M. Ziegler. *Jet spaces of varieties over differential and difference fields*. Selecta Mathematica. New series 9, pp.579-599, 2003.
- [53] B. Poizat. *Stable Groups*. Mathematical Surveys and Monographs Vol. 87. American Mathematical Society 2001.
- [54] B. Poizat. *Une théorie de Galois imaginaire*. Journal of Symbolic Logic, Vol. 48, No. 4, pp.1151-1170, 1983.

- [55] S. Shelah. *Classification theory and the number of non-isomorphic models*. Studies in Logic and the Foundations of Mathematics, Vol. 92. Elsevier Science Publishing Company, 1990.
- [56] S. Suer. *Model theory of differentially closed fields with several commuting derivations*. PhD Thesis. University of Illinois at Urbana-Champaign, 2007.
- [57] S. Suer. *On subgroups of the additive group in differentially closed fields*. The Journal of Symbolic Logic. Vol. 77, No. 2, 2012.
- [58] M. Tressl. *The Uniform Companion for Large Differential Fields of Characteristic 0*. Transactions of the American Mathematical Society. Vol. 357, No. 10, pp. 3933-3951 (2005).
- [59] L. van den Dries and K. Schmidt. *Bounds in the Theory of Polynomial Rings over Fields. A Nonstandard Approach*. Inventiones Mathematicae 76, pp. 77-91 (1984).
- [60] M. van der Put and M. Singer. *Galois theory of linear differential equations*. Springer-Verlag Berlin Heidelberg 2003.