Definability in Henselian Fields

Will Anscombe

Lincoln College

University of Oxford

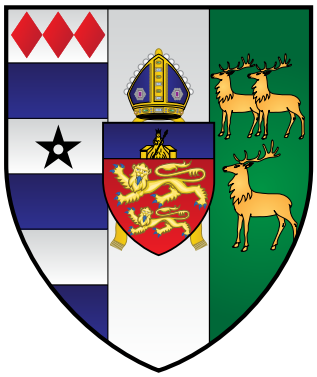

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2012

Abstract

Definability in Henselian Fields
Will Anscombe
Lincoln College, University of Oxford
Submitted for the degree of Doctor of Philosophy
Trinity Term 2012

We investigate definability in henselian fields. Specifically, we are interested in those sets and substructures that are existentially definable or definable with 'few' parameters. Our general approach is to use various versions of henselianity to understand the 'local structure' of these definable sets.

The fields in which we are most interested are those of positive characteristic, for example the local fields $\mathbb{F}_q((t))$, but many of our methods and results also apply to p-adic and real closed fields. In positive characteristic we have to deal with inseparable field extensions and we develop the method of Λ-closure to 'translate' inseparable field extensions into separable ones.

In the first part of the thesis we focus on existentially definable sets, which are projections of algebraic sets. Our main tool is the Implicit Function Theorem (for polynomials) which is equivalent to t-henselianity, by work of Prestel and Ziegler in [24]. This enables us to prove that existentially definable sets are 'large' in various senses. Using the Implicit Function Theorem, we also obtain a non-uniform local elimination of the existential quantifier. The non-uniformity and local character of this result at present forms an obstacle to full quantifier-elimination. From these technical statements we can deduce characterisations of, for example, existentially definable subfields and existentially definable transcendentals.

We prove that a dense, regular extension of t-henselian fields is existentially closed which, in particular, implies the old result of Ershov that $\mathbb{F}_p(t)^h \preceq_\exists \mathbb{F}_p((t))$. Using the existential closedness of large fields in henselian fields, we are able to apply many of these results to large fields. This answers questions for imperfect large fields that were answered in the perfect case by Fehm in [9].

In the second part of the thesis, we work with power series fields $F((t))$ and subsets which are F-definable (and not contained in F). We use a 'hensel-like' lemma to characterise F-orbits of (singleton) elements of $F((t))$. It turns out that all such orbits are $\exists$-t-definable. Consequently, we may apply our earlier results about existentially definable subsets to F-definable subsets. We can use this to characterise F-definable subfields of $F((t))$.

As a further corollary, we obtain an $\exists$-$\emptyset$-definition of $\mathbb{F}_p[[t]]$ in $\mathbb{F}_p((t))$.

For Elizabeth, min Ælflæd.

Ic eom mare þonne þes middangeard,
læsse þonne hondwyrm, leohtre þonne mona,
swiftre þonne sunne. Sæs me sind ealle
flodas on fæðmum ond þes foldan bearm,
grene wongas. Grundum ic hrine,
helle underhnige, heofonas oferstige,
wuldres eþel, wide ræce
ofer engla eard, eorþan gefylle,
ealne middangeard ond merestreamas
side mid me sylfum. Saga hwæt ic hatte.

-The Exeter Book

Acknowledgements

First and foremost I would like to thank Jochen Koenigsmann, my supervisor, for all his teaching, guidance, and friendship over the last four years. Every student owes his supervisor a debt of gratitude, but I can only try to impress upon the reader what a significant debt there is in my case. His support has been invaluable and I have truly enjoyed working with him. Thank you.

I would also like to thank everyone in the Oxford Logic group for many excellent lectures, discussions, and casual conversations over the years. Long live Logic Tea!

During my DPhil I have been lucky to talk to and get advice from many excellent mathematicians, both established and budding. In particular Boris Zilber, Jonathan Pila, Jamshid Derakhshan, Martin Bays, Franziska Jahnke, Arno Fehm, Franz-Viktor Kuhlmann, Dugald Macpherson, Tom Foster, Austin Yim, and Bernhard Elsner.

I want also to express my gratitude to those teachers, tutors, and lecturers that have taught and inspired me. In the other direction, I would like thank my students and teaching colleagues at Merton (particularly Alex Scott and Ron Reid-Edwards) and Magdalen College School (Linda and the team) for keeping my mathematical interests broad! In previous years I taught with and learned a lot (about both teaching and mathematics) from Derek Goldrei.

As someone might have said, man shall not live by maths alone. So I thank my friends for their coffee, conversation, and compansionship. Also for distracting me when I need distracting and fascinating me with their wide variety of interests.

My family have always supported and encouraged me in my studies. I would especially like to thank my parents who have had to put up with non-weekly weekly telephone calls and for nurturing my mathematical learning when I was very young.

Finally, my wife Elizabeth has been a constant support and I thank her for all her love and friendship from the bottom of my heart.

Contents

Chapter 1

Introduction

The model theory of local fields is very well developed in characteristic zero. Tarski proved, in [29], that the theory of algebraically closed fields admits quantifier-elimination in the language of rings and that the theory of real closed fields admits quantifier-elimination in the language of ordered rings. In [2], Ax and Kochen proved that the theory of the p-adics is model-complete and axiomatised by elementary sentences expressing that the valuation is a henselian p-valuation and that the value group is a $\mathbb{Z}$-group. Ershov also independently found this axiomatisation. Consequently, the theories of $\mathbb{C}$, $\mathbb{R}$, and $\mathbb{Q}_p$ (for each prime p) are decidable. This leads us by constrast to the famous open problem of whether the theory of $\mathbb{F}_p((t))$ is decidable.

The model theory of the p-adics did not end with Ax-Kochen and Ershov. In [5], Cohen gave an effective quantifier-elimination procedure (in an expanded language) by finding a cell-decomposition of definable sets; then in [18], Macintyre proved that the p-adics have quantifier-elimination in the Macintyre language, which is the ring language expanded by predicates for n-th powers for each $n \in \mathbb{N}$; and Denef reproved (in [6]) Macintyre's result by re-working and simplifying Cohen's method into a clear cell-decomposition.

In [24], Prestel and Roquette developed the 'algebraic' theory of fields elementarily equivalent to finite extensions of the p-adics; such fields are said to be *p-adically closed*. In analogy to the relationship between formally real fields and real closed fields, they defined a notion of a p-valued field and proved that each p-valued field has a p-adic closure, i.e. a 'least' p-adically closed field containing the original, which is unique up to isomorphism under the further assumption that the value group is a $\mathbb{Z}$-group. Furthermore, they prove explicit embedding theorems from which can be deduced all the

model-theoretic results described above. The analogy between p-adically closed and real closed fields is explored further in Macintyre's survey paper [19].

The study of definable sets is a central theme in model theory as well as in this thesis. The quantifier-elimination results mentioned above of course tell us a lot about definable sets in these examples. Definable sets are quantifier-free definable in the relevant language: the ring language $\mathcal{L}_{\text{ring}}$ for algebraically closed fields, the language $\mathcal{L}_{\text{ring}} \cup \{<\}$ of ordered rings for real closed fields, and the Macintyre language $\mathcal{L}_{\text{Mac}} := \mathcal{L}_{\text{ring}} \cup \{P_n \mid n \in \mathbb{N}\}$ for p-adically closed fields. Definable subsets of the (first Cartesian power of the) field have an even simpler description in each case. Furthermore, in each of the examples mentioned so far there are no proper definable subfields; and infinite definable subsets have non-empty interior in the appropriate topologies.

None of the results mentioned so far has any known analogue in the case of power series fields over finite fields. In fact, in imperfect fields there are always proper definable subfields and infinite definable subsets need not have non-empty interior. F-V Kuhlmann has studied, in [12], various possible axiomatisations for the theory of $\mathbb{F}_p((t))$ and found that the obvious translation of Ax-Kochen's axiomatisation of the theory of $\mathbb{Q}_p$ does not in this case yield a complete theory. An approach he suggests goes via the study of additive polynomials, and work has been done, by Rohwer in his thesis [26], on the theory of $\mathbb{F}_p((t))$ as a module under the ring of these polynomials. The model theory of $\mathbb{F}_p((t))$ is known to be 'wild': in [3] Cherlin proved that $\mathbb{F}_p((t))$ has the independence property.

Despite this and much other work, a great deal is left unknown. In particular, little is known about the structure of definable sets.

Henselianity

Henselianity, or more precisely t-henselianity, is a common feature of all these examples. By definition, a valued field is *henselian* if the valuation extends uniquely to every algebraic extension. There are many characterisations of henselianity, for a list see Theorem 4.1.3 in [8]. A field is called *henselian* if it admits a non-trivial henselian valuation.

In a 1978 paper of Prestel and Ziegler ([25]), t-henselianity was introduced. This is a topological generalisation of henselianity, for a precise definition see section 5.1. A non-separably closed field is t-henselian if and only if it is $\mathcal{L}_{\text{ring}}$-elementarily equivalent to a henselian field. In [25] it is proved that a field (with a field topology) is t-henselian if and only if it satisfies the Implicit Function Theorem for

polynomials, suitably re-written for the context of topological fields.

The present work

This thesis explores definability in henselian fields. In Part II we use the Implicit Function Theorem (combined with the technique of Λ-closure developed in Part I) to obtain results describing the local behaviour of existentially definable sets in non-separably closed t-henselian fields. Locally around a sufficiently generic point an existentially definable set is positive quantifier-free definable in the ring language expanded by the component maps with respect to a p-base of the field. The 'local' nature of this result at present prevents this from being converted into a true quantifier-elimination result. However, we are able to deduce various interesting corollaries including characterisations of existentially definable singletons and subfields generated by existentially definable sets.

In Part III, we prove an explicit Hensel-like Lemma for power series fields $F((t))$ in order to characterise the orbits under F-automorphisms (i.e. field automorphisms which fix F pointwise). Throughout Part III we make the assumption that F is perfect. We prove that F-orbits are in fact existentially definable (with the additional parameter t). Thus we can apply results from Part II and deduce various corollaries, for example we characterise F-definable subfields of $F((t))$ which are not subfields of F. One interesting corollary is an existential definition of $\mathbb{F}_q[[t]]$ in $\mathbb{F}_q((t))$ without using any parameters, which is exhibited in Appendix A.

Inseparability

Throughout the thesis our focus is on fields of positive characteristic, although almost everything applies to characteristic zero. Henselianity is intimately tied up with separability; thus the most major obstacle we face is the inseparability of field extensions. In Part I, we develop the tool of Λ-closure to solve this problem. The technique of Λ-closure enables us to translate inseparable field extensions into separable ones in such a way as to preserve the 'information' that we need to study definability.

Large fields

In [22], Pop defined a field L to be *large* [1] if every smooth curve defined over L has either infinitely many or no L-rational points. Equivalently, a field L is large if L is existentially closed (in $\mathcal{L}_{\text{ring}}$) in $L((t))$. This turns out to be an extremely wide class of fields. Algebraically closed fields, real closed

[1] Large fields are also known as *ample* fields.

fields, and henselian fields are large. In [9], Fehm studied existential definability in perfect large fields. He proved that perfect large fields have no proper infinite existentially definable subfields. By using the existential closure property and our work on existential definability in henselian fields, we are able to extend his results to imperfect large fields.

1.1 Notation

$\mathbb{N}, \mathbb{Z}, \mathbb{Q}$, and $\mathbb{R}$ will denote the natural numbers, integers, rational numbers, and real numbers, respectively. We use ω to denote the natural numbers when we're thinking in a more set-theoretic way, for example when we do recursive constructions. Our convention will be that $0 \in \omega$ but $0 \notin \mathbb{N}$. The symbol $\subset$ will always denote *strict* containment, otherwise we use $\subseteq$ for the subset relation.

Fields and rings We let C, D, E, F, K, L denote fields. Usually C will be a field of constants or parameters and $F/E/D/C$ will be a tower. The fields F, K, L will be our main fields of interest. For example, in Part II, K will denote a t-henselian field and, in Part III, we will be studying power series fields $F((t))$. In Appendix B, L will be a large field. We let p denote the characteristic or (occasionally) the characteristic exponent of whichever field we're working in. The real numbers are denoted by $\mathbb{R}$, the p-adic numbers by $\mathbb{Q}_p$, and finite fields by $\mathbb{F}_q$ (for q a prime power). Occasionally we may let $\mathbb{F}_p$ denote the prime subfield even in characteristic 0 for notational convenience.

Tuples and cartesian products Let A, B be any two sets. An *A-tuple* in B is a function $A \longrightarrow B$. We let ${}^{A}B$ denote the set of A-tuples in B. We do not require tuples to be either ordered or finite; although of course the size of an A-tuple is equal to the size of A. Tuples will be in bold fonts.

- $\mathbf{a}, \mathbf{b}, \mathbf{c}$ will be tuples of field elements,
- $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ will be tuples of basic open neighbourhoods around 0, and
- $\mathbf{x}, \mathbf{y}, \mathbf{z}$ will denote tuples of variables.

Tuples of elements will usually correspond to tuples of variables; for example, $\mathbf{a} \in {}^{\mathbf{x}}F$ means that $\mathbf{a}$ is an $\mathbf{x}$-tuple from F. On the other hand we may write $\mathbf{a} \subseteq A$ if we do not care to make the index set explicit. We often use juxtaposition of tuples to denote disjoint union. We will sometimes discuss partitions of tuples: if $\mathbf{x} = \mathbf{yz}$ is a partition of $\mathbf{x}$ and $\mathbf{a} \in {}^{\mathbf{x}}F$ then we will write $\mathbf{a} = \mathbf{bc} \in {}^{\mathbf{yz}}F$ for the corresponding partition of $\mathbf{a}$; i.e. the partition $\mathbf{a} = \mathbf{bc}$ such that $\mathbf{b} \in {}^{\mathbf{y}}F$ and $\mathbf{c} \in {}^{\mathbf{z}}F$.

Fixed maps Φ will denote the Frobenius map $x \longmapsto x^p$. Let $\mathbf{x}$, $\mathbf{y}$ be tuples of variables and let $\mathbf{a}$ be an $\mathbf{x}$-tuple. Then $\pi_{\mathbf{a}}$ will denote the specialisation of a rational function $f(\mathbf{x}, \mathbf{y}) \longmapsto f(\mathbf{a}, \mathbf{y})$ (where this is well-defined) and χ will denote the cross-section of the specialisation $f(\mathbf{a}, \mathbf{y}) \longmapsto f(\mathbf{x}, \mathbf{y})$ given by simply replacing the *original* occurrences of $\mathbf{a}$ with $\mathbf{x}$ (there may be other occurrences of elements of $\mathbf{a}$ as part of the coefficients of f). Note that some awkwardness arises here: there may be rational functions f, g such that $\pi_{\mathbf{a}} f = \pi_{\mathbf{a}} g$ bu $f \neq g$. This shows that χ cannot strictly be a function. We let D_y denote formal differentiation with respect to the variable y.

Basic geometry Let $\mathbf{a}$ be an $\mathbf{x}$-tuple and let $f \in C[\mathbf{x}]$ be a polynomial. Then $Z(f)$ denotes the set of zeroes of f and $I(\mathbf{a}/C)$ denotes the ideal in $C[\mathbf{x}]$ of polynomials over C which are zero at $\mathbf{a}$. Let $\mathbf{y}$ be another tuple of variables such that $\mathbf{x} \subseteq \mathbf{y}$ and let $I_{\mathbf{y}}(\mathbf{a}/C)$ be the ideal in $C[\mathbf{y}]$ generated by $I(\mathbf{a}/C)$ (thus $I_{\mathbf{y}}(\mathbf{a}/C)$ is simply those polynomials in $\mathbf{y}$ which, when thought of as polynomials in the variables $\mathbf{y} \setminus \mathbf{x}$, have coefficients from $I(\mathbf{a}/C)$). We let $\mathrm{locus}_{\mathbf{y}}(\mathbf{a}/C)$ denote the zero set of $I_{\mathbf{y}}(\mathbf{a}/C)$; thus elements of $\mathrm{locus}_{\mathbf{y}}(\mathbf{a}/C)$ are $\mathbf{y}$-tuples the $\mathbf{x}$-portions of which are zeroes of $I(\mathbf{a}/C)$. We let $\pi_{\mathbf{x}}$ denote the projection map onto the $\mathbf{x}$ co-ordinates.

Topological fields In the context of topological fields, $\mathcal{T}$ will denote a field topology, τ will denote a base of the neighbourhood filter around 0, and $\mathcal{T}_\tau$ will be the field topology generated by τ.

Let F be a topological field, let $\boldsymbol{\alpha} \in {}^{\mathbf{x}}\tau$, and let $\mathbf{a} \in {}^{\mathbf{x}}F$. Then we define $\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) := \prod\{a(x) + \alpha(x) | x \in \mathbf{x}\}$ to be the *ball* of radius $\boldsymbol{\alpha}$ with centre $\mathbf{a}$.

Imperfection and p-independence Let $E/D/C$ be a tower of fields. The following summarises definitions that appear in section 2.2.

- $\mathrm{impdeg}(E/C)$ is the relative imperfection degree of E over C.
- $pI(D; E/C)$ is the set of tuples from D which are p-independent in E over C.
- $pI(E/C) := pI(E; E/C)$ is the set of tuples from E which are p-independent in E over C.
- $pI(D; E) := pI(D; E/\mathbb{F}_p)$ is the set of tuples from D which are absolutely p-independent in E.
- $pI(E) := pI(E/\mathbb{F}_p)$ is the set of tuples from E which are absolutely p-independent in E.
- $pI_{\max}(D; E/C) \subseteq pI(D; E/C)$ is the set of tuples from $pI(D; E/C)$ which are maximal with respect to inclusion.

- $pI_{\max}(E/C) := pI_{\max}(E; E/C)$ is the set of tuples from E which are p-independent over C which are maximal with respect to inclusion.
- Maximal p-independent tuples and p-bases coincide; thus we also let $pB(E/C)$ denote the set $pI_{\max}(E/C)$ of p-bases of E over C.

Let $\mathbf{a} \in pI(E/C)$.

- $M_{\mathbf{a}} := \left\{\prod_{a \in \mathbf{a}} a^{i(a)} | i \in {}^{\mathbf{a}}\{0, ..., p-1\} \text{ has finite support}\right\}$ is the set of p-monomials of $\mathbf{a}$. This is an E^pC-linear base of $E^pC(\mathbf{a})$.

A remark on 'lambda notation'

In our study of component maps and Λ-closure, there are several different constructions and maps with similar rôles and notations. Let $\mathbf{a} \in pI(E)$. The p-th roots of the co-ordinate maps with respect to the basis $M_{\mathbf{a}}$ are called the *component maps* with respect to $\mathbf{a}$. This is a tuple of maps, each a function $E^p(\mathbf{a}) \longrightarrow E$. These definitions are further explained in chapter 4.

Component maps:

- $\boldsymbol{\lambda}^{\mathbf{a}}$ denotes the tuple of component maps with respect to a p-independent tuple $\mathbf{a}$;
- λ_m denotes a component maps with respect to the monomial $m \in M_{\mathbf{a}}$, we write $\lambda := \lambda_m$ if m is understood;

Constructing the Λ-closure:

- $\Lambda(E/D), \Lambda_n(E/D)$, and $\boldsymbol{\Lambda}(E/D)$ denote different stages of the construction of Λ-closure;
- $\boldsymbol{\Lambda}D$ will sometimes denote the Λ-closure of D where the overfield is implicit;
- $\mathbf{a}^n$ denotes the tuple constructed from $\mathbf{a}$ in the n-th stage of Λ-closure (note that this is certainly not exponentiation!);
- $\boldsymbol{\lambda}\mathbf{a}$ denotes the tuple $\mathbf{a}^n$ where $n < \omega$ is such that $\boldsymbol{\Lambda}(E/C(\mathbf{a})) = C(\mathbf{a}^n)$ (where $E = F$ in the global method and $E = C(\mathbf{ab})$ in the local method);

Global and local Λ-translation:

- $\boldsymbol{\Lambda}, \Lambda$ denote the tuples of rational maps such that $\Lambda(\mathbf{ab}) = \mathbf{a}^n$ and $\boldsymbol{\Lambda} := \Lambda \times \mathrm{id}_{\mathbf{y}}$; and

- $\boldsymbol{\Sigma}, \Sigma$ denote the tuples of polynomial maps such that $\Sigma(\boldsymbol{\lambda}\mathbf{a}) = \mathbf{a}$ and $\boldsymbol{\Sigma} := \Sigma \times \mathrm{id}_{\mathbf{y}}$.

Note that neither $\boldsymbol{\Lambda}$ nor Λ are defined in the global method.

1.2 Model theoretic prerequisites

There are really very few model theoretic prerequisites to this work. Readers unfamiliar with the first-order languages, elementary theories, elementary extensions, and the Compactness Theorem may want to refer to Marker's textbook [20]. We frequently use the fact that, in fields, existential formulas are equivalent to existential-positive formulas. In fact, if the field is not algebraically closed, then the quantifier-free portion can be taken to be a single atomic formula. We also use the fact that C-definable sets are closed under C-automorphisms (i.e. those that fix C pointwise). We let $\mathcal{L}_{\mathrm{ring}}$ denote the ring language $\{+, \cdot, 0, 1\}$.

Part I

Λ-closure

Chapter 2

Preliminaries on separability in fields

In chapter 6 we will apply the Implicit Function Theorem (the IFT) to projections of loci in the context of t-henselian fields. We will say that the projection is separable if the corresponding field extension is separable; and, using the IFT, we will be able to understand separable projections. Thus the difficulty is inseparable projections, those that correspond to inseparable field extensions.

In this chapter and the next we develop the tool of Λ-closure to translate inseparable into separable field extensions.

2.1 The separability of field extensions

As a preliminary to studying inseparability of arbitrary field extensions, we first recall some basic facts about separable extensions. Everything in this section is standard and can be found in [17], although Lang defines separable field extensions by using separability degree which we will not use and will thus avoid.

2.1.1 Algebraic field extensions

First, we suppose that E/C is algebraic. More details can be found in Chapter V, Section 4 of [17]. Let p be the characteristic exponent, so that $p = 1$ in characteristic zero.

Definition 2.1.1. We say that an irreducible polynomial is *separable* if it has no repeated roots in an algebraic closure; equivalently, if it has no repeated roots in every algebraic closure. We say that an element $a \in E$ is *separably algebraic* (or just *separable*) over C if the minimal polynomial of a over C is separable. We say that a tuple $\mathbf{a} \subseteq E$ is *separably algebraic* over C if each $a \in \mathbf{a}$ is separable over C. If all elements of E are separable over C, then we say that E/C is *separably algebraic* over C. We may also say that E is a *separable algebraic* extension of C.

We briefly summarise some facts about separably algebraic extensions.

Fact 2.1.2. If a is algebraic over C then there exists $n \in \mathbb{N}$ such that a^{p^n} is separably algebraic over C. Thus all algebraic extensions in characteristic zero are separably algebraic. If a is separably algebraic over C and $m \in C[x]$ is the minimal polynomial of a over C then $D_x m(a) \neq 0$. An extension generated by a separably algebraic tuple is separably algebraic. Separably algebraic field extensions form a distinguished class (in the sense of [17]): the class is closed under towers and composita of arbitrarily many extensions (of the same field). Any intermediate extension of a separably algebraic extension is separably algebraic.

Next, we recall the definition of purely inseparable extensions. More details can be found in Chapter V, Section 6 of [17].

Definition 2.1.3. We say that $a \in E$ is *purely inseparable* over C if there exists $n \in \mathbb{N}$ such that $a^{p^n} \in C$. A tuple $\mathbf{a} \in E$ is said to be *purely inseparable* over C if each $a \in \mathbf{a}$ is purely inseparable over C. We say that E/C is *purely inseparable* if every element of E is purely inseparable over C.

Fact 2.1.4. An extension generated by a purely inseparable tuple is purely inseparable. Purely inseparable extensions form a distinguished class (in the sense of [17]). For each $n < \omega$, E/E^{p^n} is purely inseparable.

Proposition 2.1.5. *(Theorem 12, [30]) Let E/C be separably algebraic and purely inseparable. Then $E = C$.*

Proposition 2.1.6. *If $E = CE^{p^n}$ for some $n \in \mathbb{N}$ then $E = CE^{p^n}$ for every $n \in \mathbb{N}$.*

Proof. Let $n \in \mathbb{N}$ be such that $E = CE^{p^n}$. Certainly $E = CE^p$. Applying the Frobenius isomorphism, we have $E^p = C^p E^{p^{n+1}}$. If we take the compositum with C, we have that $E = CE^p = CE^{p^{n+1}}$. Thus $E = CE^{p^m}$ for $m \geq n$. Now suppose that $m < n$; then $E^{p^m} \supseteq E^{p^n}$ and $E = CE^{p^m}$. □

Proposition 2.1.7. *Suppose that E/C is separably algebraic. Then $E = CE^{p^n}$, for each $n \in \mathbb{N}$.*

Proof. The extension E/CE^{p^n} is purely inseparable. As E/CE^{p^n} is an intermediate extension of E/C, which is assumed to be separably algebraic, in fact E/CE^{p^n} is separably algebraic. By Proposition 2.1.5, $E = CE^{p^n}$. □

Sadly, the converse does not hold as the following example shows.

Example 2.1.8. Let E be any field of characteristic p such that $E = E^p$, i.e. E is *perfect*. Let $C \subseteq E$ be any subfield. Then $E = CE^p = E^p$ and yet E/C need not be separably algebraic.

The following definition was introduced by Kraft in [11].

Definition 2.1.9. Let E/C be algebraic. The *inseparability exponent* of E/C is the least $n \in \mathbb{N}$ such that $E^{p^n}C$ is separably algebraic over C or ∞ if no such $n \in \mathbb{N}$ exists. We denote it by $\text{inex}(E/C)$.

The converse to Proposition 2.1.7 does hold under the additional assumption that E/C is of finite inseparability exponent.

Proposition 2.1.10. *Let E/C be an algebraic extension such that $\text{inex}(E/C) < \infty$ and $E = CE^{p^n}$ for some $n \in \mathbb{N}$. Then E/C is separably algebraic.*

Proof. Let $m := \text{inex}(E/C)$ and let D be an intermediate field of E/C equal to the compositum of the separably algebraic extensions of C contained in E. By Fact 2.1.2, D/C is separably algebraic. Usually D is called the relative separably algebraic closure of C in E. Let $a \in E$. By definition of $\text{inex}(E/C)$, a^{p^m} is separably algebraic over C. Thus $a^{p^m} \in D$ and $E^{p^m} \subseteq D$. By Proposition 2.1.6, $E = CE^{p^m} \subseteq D$. Thus $E = D$ and E/C is separably algebraic. □

2.1.2 Separably generated field extensions

We study the extension of the definition of separability to transcendental extensions, so we no longer suppose that E/C is algebraic.

Linear disjointness from p-th roots

We will briefly need the notion of linear disjointness in order to properly introduce separability for arbitrary field extensions. See Chapter VIII, Section 3 of [17] for more details.

Definition 2.1.11. Let D/C be another extension. We do not suppose that $D \subseteq E$ but we do suppose that D, E lie in some common extension. We say that E is *linearly disjoint* from D over C if subsets of E which are C-linearly independent are also D-linearly independent.

Lemma 2.1.12. *If E/C is purely transcendental then E and $C^{p^{-\infty}}$ are linearly disjoint over C.*

Proof. Purely transcendental extensions (as examples of regular extensions) are linearly disjoint from algebraic extensions. □

Lemma 2.1.13. *If E/C is separably algebraic then E and $C^{p^{-\infty}}$ are linearly disjoint over C.*

Proof. A slight adaptation of the proof of Proposition 2.1.7 gives that separably algebraic and purely inseparable extensions are linearly disjoint. □

These two lemmas suggest a commonality between purely transcendental and separably algebraic extensions. Let $\mathbf{a} \subseteq E$ be any tuple.

Definition 2.1.14. We say that $\mathbf{a}$ is a *separating transcendence base* of E/C if

- $\mathbf{a}$ is a transcendence base of E/C and
- E is separably algebraic over $C(\mathbf{a})$.

The extension E/C is *separably generated* if there exists a separating transcendence base of E/C.

Proposition 2.1.15. *Let E/C be separably generated. Then E and $C^{p^{-\infty}}$ are linearly disjoint over C.*

Proof. By definition, there exists a separating transcendence base $\mathbf{a} \subseteq E$. By Lemma 2.1.12, $C(\mathbf{a})$ and $C^{p^{-\infty}}$ are linearly disjoint over C. By Lemma 2.1.13, E is linearly disjoint from $C^{p^{-\infty}}(\mathbf{a})$ over $C(\mathbf{a})$ since $E/C(\mathbf{a})$ is separably algebraic and $C^{p^{-\infty}}(\mathbf{a})/C(\mathbf{a})$ is purely inseparable. By a simple 'tower lemma' for linear disjointness (Chapter VIII, Proposition 3.1 of [17]), E is linearly disjoint from $C^{p^{-\infty}}$ over C, as required. □

Refining tuples to find separating transcendence bases

Let D/C be an extension which is linearly disjoint from $C^{1/p}$ and let $\mathbf{a} \subseteq D$ be a *finite* tuple. Our aim is to refine $\mathbf{a}$ to a separating transcendence base of $C(\mathbf{a})/C$.

We first prove a kind of 'exchange lemma' in which we swap two elements to turn algebraic dependence into separable algebraic dependence. This first appears as Lemma 1 of [14].

Lemma 2.1.16. *Let $\mathbf{d}d$ be an $\mathbf{x}x$-tuple such that $\mathbf{d}$ is algebraically independent over C, d is algebraic over $C(\mathbf{d})$, and $C(\mathbf{d}d)$ is linearly disjoint from $C^{1/p}$ over C. Then there exists $d_s \in \mathbf{d}d$ which is separably algebraic over $C(\mathbf{d}d \setminus \{d_s\})$.*

Proof. Let $m \in C[\mathbf{x}x]$ be a polynomial of least degree (in x) which is non-zero and is such that $m(\mathbf{d}d) = 0$; for example choose the numerator (in the sense of rational functions in $\mathbf{d}$) of the minimal polynomial of d over $C(\mathbf{d})$ and replace $\mathbf{d}$ by the variables $\mathbf{x}$. By definition of m and the algebraic independence of $\mathbf{d}$ over C, m cannot be constant in x.

If $m \in C[\mathbf{x}^p x^p]$ then let $n := (m(\mathbf{x}x))^{1/p} \in C^{1/p}[\mathbf{x}x]$. Then $n(\mathbf{d}d) = (m(\mathbf{d}d))^{1/p} = 0$. Note that $\deg_x(n) = \frac{1}{p}\deg_x(m) < \deg_x(m)$. This is not yet a contradiction since n is defined over $C^{1/p}$. The equation $n(\mathbf{d}d) = 0$ is a $C^{1/p}$-linear relation between the monomials of $n(\mathbf{d}d)$. By the linear disjointness of $C(\mathbf{d}d)$ and $C^{1/p}$ over C there exists a C-linear relation between the monomials of $n(\mathbf{d}d)$: but this is simply another polynomial n' of the same degree as n but defined over C. Now $\deg_x(n') = \deg_x(n) < \deg_x(m)$ and $n'(\mathbf{d}d) = 0$ contradict the minimality of m.

Thus $m \notin C[\mathbf{x}^p x^p]$. Choose any variable $x_s \in \mathbf{x}x$ such that a non p-th power of x_s appears in m, let d_s be the element of $\mathbf{d}$ corresponding to the variable x_s. Then d_s is separably algebraic over $C(\mathbf{d}d \setminus \{d_s\})$. □

Lemma 2.1.17. *Let $\mathbf{d}$ be such that $\mathbf{d}$ is not algebraically independent over C and $C(\mathbf{d})$ is linearly disjoint from $C^{1/p}$ over C. Then there exists $d_s \in \mathbf{d}$ which is separably algebraic over $C(\mathbf{d} \setminus \{d_s\})$.*

Proof. Let $\mathbf{d}' \subset \mathbf{d}$ be a transcendence base for $C(\mathbf{d})/C$. Choose any $d \in \mathbf{d} \setminus \mathbf{d}'$. Then we apply Lemma 2.1.16 to the tuple $\mathbf{d}'d$ in order to find $d_s \in \mathbf{d}'d$ which is separably algebraic over $C(\mathbf{d}'d \setminus \{d_s\})$. Thus d_s is certainly separably algebraic over $C(\mathbf{d} \setminus \{d_s\})$. □

In the next proposition and others following, we say that a tuple $\mathbf{a}$ may be *refined* to have a given property if some subtuple of $\mathbf{a}$ has the property.

Proposition 2.1.18. *Suppose that D is linearly disjoint from $C^{1/p}$ over C and let $\mathbf{a} \subseteq D$ be a finite tuple. Then $\mathbf{a}$ may be refined to a separating transcendence base of $C(\mathbf{a})/C$.*

Proof. We proceed by induction on the length of $\mathbf{a}$. The result clearly holds if $\mathbf{a}$ is the empty tuple. Now consider the finite tuple $\mathbf{a}a \subseteq D$. If $\mathbf{a}a$ is algebraically independent over C then $\mathbf{a}a$ is already a separating transcendence base. If not, then $\mathbf{a}a$ is algebraically dependent over C and we may apply Lemma 2.1.17. □

However, the result doesn't hold without the requirement that $\mathbf{a}$ be finite.

Example 2.1.19. Let D/C be linearly disjoint from $C^{1/p}$ and of infinite transcendence degree. There exists a tuple $\mathbf{a} \subseteq D$ which is infinite and **cannot** be refined to a separating transcendence base of $C(\mathbf{a})/C$.

We adapt the proof of Theorem 1 from [21].

Proof. Let $\mathbf{a} := (a_n)_{n\in\mathbb{N}} \subseteq D$ be a tuple of elements algebraically independent over C. We will find another tuple of generators for the extension $C(\mathbf{a})/C$ which cannot be refined to a separating transcendence base. Note that, as a subextension of D/C, $C(\mathbf{a})/C$ is linearly disjoint from $C^{1/p}$. Note also that $\mathbf{a}$ is already a separating transcendence base of $C(\mathbf{a})/C$. Set $\mathbf{a}_1 := (a_n a_{n+1}^p)_{n\in\mathbb{N}}$ and $\mathbf{a}_2 := (a_n^p)_{n\geq 2}$. Then set $\mathbf{a}' := \mathbf{a}_1\mathbf{a}_2$. Since $a_n = a_n a_{n+1}^p a_{n+1}^{-p}$, it is clear that $C(\mathbf{a}) = C(\mathbf{a}')$.

Let $\mathbf{b} \subseteq \mathbf{a}'$ be a separating transcendence base of $C(\mathbf{a}')/C$. By Proposition 2.1.20, $\mathbf{b}$ is both a transcendence base of $C(\mathbf{a})/C$ and a relative p-base of $C(\mathbf{a})/C$. By relative p-independence, for each $b \in \mathbf{b}$ we have $b \notin C(\mathbf{a}^p)$. Since $\mathbf{a}_2 \subseteq C(\mathbf{a}^p)$, $\mathbf{b} \subseteq \mathbf{a}_1$.

Now we claim that a_1 is transcendental over $C(\mathbf{a}_1)$. Suppose not: then a_1 is algebraic over $C_n := C(a_1 a_2^p, ..., a_n a_{n+1}^p)$ for some $n \in \mathbb{N}$. But then a_i is also algebraic over C_n for each $i \leq n+1$. Thus the transcendence degree of C_n/C is at least $n+1$, which cannot possibly be right since C_n is generated over C by n elements. The contradiction proves that a_1 is transcendental over $C(\mathbf{a}_1)$. (Actually, we have also seen that the transcendence degree of $C(\mathbf{a})/C(\mathbf{a}_1)$ is 1.) Thus $\mathbf{b}$ cannot be a transcendence base of $C(\mathbf{a})/C$, so in particular $\mathbf{b}$ cannot be a separating transcendence base. □

Proposition 2.1.20. *(Theorem 13, [13]) Let D/C be separable and let $\mathbf{a} \subseteq D$. Then $\mathbf{a}$ is a separating transcendence base of D/C if and only if $\mathbf{a}$ is both a transcendence base and a (relative) p-base of D/C.*

2.1.3 Separable field extensions

Let E/C be an arbitrary field extension. The following 'characterisation' of separability (which we present as a definition) is sometimes known as Mac Lane's Criterion. This subject was developed by Mac Lane and many others in the late 1930s while answering the question 'which field extensions are separably generated?' See Mac Lane's papers [13] and [14] for more details. Also see the more expository article [15].

Definition 2.1.21. We say that E/C is *separable* if each finite tuple $\mathbf{a} \subseteq E$ can be refined to a separating transcendence base of $C(\mathbf{a})/C$.

The following proposition is a slight rephrasing of that in Lang ([17]), although the proof (which again is from [17]) we have split up over the previous few subsections.

Proposition 2.1.22. *(Chapter VIII, Proposition 4.1, [17]) The following are equivalent.*

1. *Every finitely generated subextension D/C of E/C is separably generated.*
2. *E/C is separable.*
3. *E and $C^{p^{-1}}$ are linearly disjoint over C.*
4. *E and $C^{p^{-n}}$ are linearly disjoint over C, for some $n \in \mathbb{N}$.*
5. *E and $C^{p^{-n}}$ are linearly disjoint over C, for each $n \in \mathbb{N}$.*
6. *E and $C^{p^{-\infty}}$ are linearly disjoint over C.*

Proof. It is clear that (6) $\iff$ (5) $\implies$ (4) $\implies$ (3) and that (2) $\implies$ (1).

(3) $\implies$ (2) Let $\mathbf{a} \subseteq E$ be a finite tuple. By Proposition 2.1.18, $\mathbf{a}$ may be refined to a separating transcendence base of $C(\mathbf{a})/C$. Thus E/C is separable.

(1) $\implies$ (6) Since linear disjointness is a condition that refers to only finitely many elements at once, it suffices to show that finitely generated subextensions of E/C are linearly disjoint from $C^{p^{-\infty}}$ over C. But if $D \subseteq E$ is finitely generated over C then it is separably generated over C, by (1). We apply Proposition 2.1.15 to find that D/C is linearly disjoint from $C^{p^{-\infty}}$. □

Corollary 2.1.23. *Separably generated extensions are separable.*

Proof. Immediate from Proposition 2.1.15 and Proposition 2.1.22. □

On the other hand, separable extensions need not be separably generated as shown by the following example.

Example 2.1.24. Let $E := \mathbb{F}_p(t^{p^{-\infty}})$. Then any finitely generated subfield $C \subseteq E$ is contained in some $\mathbb{F}_p(t^{p^n})$, for some $n \in \mathbb{Z}$. By Lüroth's Theorem $C = \mathbb{F}_p(f(t^{p^n}))$ for some rational function f. In particular, C is a simple transcendental extension. Obviously a single transcendental element is a separating transcendence base for $E/\mathbb{F}_p$. and thus separably generated. Therefore $E/\mathbb{F}_p$ is separable.

On the other hand, we have that $\mathrm{trdeg}(E/\mathbb{F}_p) = 1$ so any transcendence base is a singleton; but any singleton is contained in some subfield $\mathbb{F}_p(t^{p^n})$, for some $n \in \mathbb{Z}$. But $\mathbb{F}_p(t^{p^{n-1}})$ is purely inseparable over $\mathbb{F}_p(s) \supseteq \mathbb{F}_p(t^{p^n})$. Thus $E/\mathbb{F}_p$ is not separably generated.

2.1.4 Perfect and imperfect fields

Definition 2.1.25. We say that a field C is *perfect* if all extensions of C are separable. A field which is not perfect is *imperfect*.

Proposition 2.1.26. *The following are equivalent.*

1. *C is perfect,*
2. *all algebraic extensions of C are separable, and*
3. *C is of characteristic zero or $C = C^p$.*

Proof. Extensions in characteristic zero are separable. Thus any field of characteristic zero is perfect. Suppose now that C has characteristic p. If $C \supset C^p$ then $C^{1/p}/C$ is a non-trivial purely inseparable, thus algebraic, extension. Conversely, if $C = C^p$ then also $C = C^{p^{-\infty}}$ and every extension is trivially linearly disjoint from $C^{p^{-\infty}}$. □

Imperfection degree

Let E/C be a field extension of characteristic exponent p and let $n \in \mathbb{N}$. Since E/E^{p^n} is purely inseparable the degree $[E : E^{p^n}]$ is a power of p (if it is infinite then we think of it as an infinite power of p, as a supernatural number).

Definition 2.1.27. The *(relative) imperfection degree* (of order n) is defined to be:

$$\mathrm{impdeg}_n(E/C) := \log_p[E : E^{p^n}C].$$

We will almost always be concerned with the case $n = 1$, in which case we write $\mathrm{impdeg}(E/C) := \mathrm{impdeg}_1(E/C)$. We let $\mathrm{impdeg}(E) := \mathrm{impdeg}(E/\mathbb{F})$ be the *(absolute) imperfection degree* of E, where $\mathbb{F}$ is the prime subfield.

Fact 2.1.28. C is perfect if and only if $\mathrm{impdeg}(C) = 0$.

2.2 p-independence

In this section we further investigate separable field extensions using Teichmüller's notion of p-independence, introduced in [30] and very thoroughly developed by Mac Lane in [13]. Let E/C be a field extension of characteristic exponent p.

2.2.1 p-independence, p-span, and p-bases

Definition 2.2.1. We say that a tuple $\mathbf{a} \subseteq E$ is

1. *p-independent in E over C* if, for each $a \in \mathbf{a}$, $a \notin E^pC(\mathbf{a} \setminus \{a\})$; that it
2. *p-spans E over C* if $E = E^pC(\mathbf{a})$; and that it is a
3. *p-base for E over C* if it is p-independent in E over C and p-spans E over C.

Next we give notation for sets of p-independent tuples. Later on, we shall be deliberately choosing p-independent subtuples of given tuples. We will need to keep track of the location of p-independent tuples as well as in which field extensions they are p-independent. Thus the following notation is a bit cumbersome.

Definition 2.2.2. Let $pI(D; E/C)$ be the set of tuples from D which are p-independent in E over C. Let $pI_{\max}(D; E/C)$ be the set of maximal elements (i.e. tuples) from $pI(D; E/C)$, so that $\mathbf{a} \in pI_{\max}(D; E/C)$ if and only if $\mathbf{a} \in pI(D; E/C)$ and no tuple from D which contains $\mathbf{a}$ is p-independent in E over C. If $D = E$ then we simply write $pI(E/C) := pI(E; E/C)$ and $pI_{\max}(E/C) := pI_{\max}(E; E/C)$. We also let $pB(E/C)$ be the set of tuples from E which are p-bases of E over C.

All of the above definitions are explicitely *relative* to the subfield C, but can be made *absolute* by setting C to be the prime field $\mathbb{F}$. We let $pI(E) := pI(E/\mathbb{F})$ be the set of *absolutely p-independent* tuples and we let $pB(E) := pB(E/\mathbb{F})$ be the set of *absolute p-bases* of E. Similarly we denote $pI(D; E) := pI(D; E/\mathbb{F})$ and $pI_{\max}(D; E) := pI_{\max}(D; E/\mathbb{F})$.

Proposition 2.2.3. *p-independence satisfies the exchange property. That is, for all $\mathbf{a} \subseteq E$ and $b, c \in E$, we have that $c \in E^pC(\mathbf{a}b) \setminus E^pC(\mathbf{a})$ implies that $b \in E^pC(\mathbf{a}c)$.*

Proof. We might as well set $D := C(\mathbf{a})$. We know that $c \in E^pD(b) \setminus E^pD$; thus the extension $E^pD(b)/E^pD$ is proper and of degree p. Similarly $E^pD(c)/E^pD$ is proper and also of degree p; but it is also a subextension of $E^pD(b)$, by the assumption. Thus $E^pD(c) = E^pD(b)$ and, in particular, we have that $b \in E^pD(c)$. □

Proposition 2.2.4. *Let* $\mathbf{a} \in pI(E/C)$ *and let* $b \in pI(E/C(\mathbf{a}))$. *Then* $\mathbf{a}b \in pI(E/C)$.

Proof. Let $a \in \mathbf{a}$ and set $\mathbf{a}' := \mathbf{a} \setminus \{a\}$. We aim to show that $a \notin E^pC(\mathbf{a}', b)$. Suppose, for a contradiction, that $a \in E^pC(\mathbf{a}', b)$. By assumption we have that $a \notin E^pC(\mathbf{a}')$; we apply Proposition 2.2.3, to find that $b \in E^pC(\mathbf{a}', a) = E^pC(\mathbf{a})$. This is a contradiction to our other assumption. Thus $a \notin E^pC(\mathbf{a}', b)$, as required. □

Proposition 2.2.5. *Any* $\mathbf{a} \in pI(D; E/C)$ *may be extended to* $\mathbf{b} \in pI_{\max}(D; E/C)$ *such that* $D \subseteq E^pC(\mathbf{b})$.

Proof. This is an argument by Zorn's Lemma and a simple application of Proposition 2.2.4. □

Proposition 2.2.6. *Let* $\mathbf{a} \subseteq E$. *The following are equivalent.*

1. $\mathbf{a}$ *is maximally p-independent in* E *over* C.
2. $\mathbf{a}$ *is minimally p-spanning* E *over* C.
3. $\mathbf{a}$ *is a p-base for* E *over* C.

Proof. Now, let $\mathbf{a} \in pB(E/C)$ be p-base. By definition $\mathbf{a}$ is p-independent and p-spanning. Since $\mathbf{a}$ p-spans, $E \setminus E^pC(\mathbf{a}) = \emptyset$ and no extension of $\mathbf{a}$ is p-independent. Let $a \in \mathbf{a}$. Then $a \notin E^pC(\mathbf{a} \setminus \{a\})$ by the p-independence of $\mathbf{a}$, so $\mathbf{a} \setminus \{a\}$ doesn't p-span. Thus $\mathbf{a}$ is maximal p-independent and minimal p-spanning. This proves that (3) $\implies$ (1), (2).

Conversely, a maximal p-independent tuple $\mathbf{a}$ must p-span since any element of $E \setminus E^pC(\mathbf{a})$ would be give a p-independent extension of $\mathbf{a}$. Thus $\mathbf{a} \in pB(E/C)$ and so (1) $\implies$ (3).

Finally, suppose that $\mathbf{a}$ is minimal p-spanning and let $a \in \mathbf{a}$. We know that $\mathbf{a} \setminus \{a\}$ doesn't span, i.e. $E^pC(\mathbf{a} \setminus \{a\}) \neq E$. Furthermore, we know that $\mathbf{a}$ does span! Thus $E = E^pC(\mathbf{a}) = E^pC(\mathbf{a} \setminus \{a\})(a) \supset E^pC(\mathbf{a} \setminus \{a\})$ is a proper field extension. Thus $a \notin E^pC(\mathbf{a} \setminus \{a\})$. This shows that $\mathbf{a}$ is p-independent and so $\mathbf{a} \in pB(E/C)$. Thus (2) $\implies$ (3). □

Proposition 2.2.7. *Any* $\mathbf{a} \in pI(E/C)$ *may be extended to* $\mathbf{b} \in pB(E/C)$. *Thus* $pB(E/C) \neq \emptyset$.

Proof. We apply Proposition 2.2.5 to the case $D = E$ and use Proposition 2.2.6. Applying this to the empty tuple, we have that $pB(E/C)$ is non-empty. □

2.2.2 Reinterpreting separability I: p-independence

These notions developed from p-independence will help us understand the separability of field extensions. In the next proposition, which is due to Mac Lane, we give another characterisation of separability of field extensions in terms of preservation of (absolute) p-independence.

Proposition 2.2.8. *(Theorems 7 and 10, [13]) For a field extension E/C the following are equivalent.*

1. *E/C is separable;*
2. *$pI(C) \subseteq pI(E)$, i.e. every p-independent tuple in C is p-independent in E;*
3. *$pB(C) \subseteq pI(E)$, i.e. every p-base in C is p-independent in E; and*
4. *$pB(C) \cap pI(E) \neq \emptyset$, i.e. there exists a p-base of C which is p-independent in E;*

Proof. (1) $\implies$ (2) See Theorem 10 of [13].

(2) $\implies$ (3) is clear since $pB(C) \subseteq pI(C)$. Suppose (3). By Proposition 2.2.7, there exists $\mathbf{a} \in pB(C)$. By (3), $\mathbf{a} \in pI(E)$. Thus (3) $\implies$ (4).

Suppose (4) and let $\mathbf{b} \in pB(C) \cap pI(E)$. We aim to show (2), so let $\mathbf{a} \in pI(C)$. Since $\mathbf{b}$ is a p-base in C, there exists a finite subtuple $\mathbf{b}' \subseteq \mathbf{b}$ such that $\mathbf{a} \subseteq C^p(\mathbf{b}')$. Let $\mathbf{b}'' \in pB(C^p(\mathbf{b}')/C^p(\mathbf{a}))$. Then

$$C^p(\mathbf{b}') = C^p(\mathbf{a}\mathbf{b}''\mathbf{b}'^p) = C^p(\mathbf{a}\mathbf{b}'') \tag{2.1}$$

and $\mathbf{a}\mathbf{b}''$ is p-independent over C. Furthermore $|\mathbf{b}'| = |\mathbf{a}\mathbf{b}''|$. If we take the compositum of E^p with the field in Equation 2.1 then we obtain

$$K^p(\mathbf{b}') = K^p(\mathbf{a}\mathbf{b}'').$$

Since $\mathbf{b}' \in pI(E)$, it must be that the degree of this field over K^p is $p^{|\mathbf{b}'|}$. Since $|\mathbf{a}\mathbf{b}''| = |\mathbf{b}'|$, we must have that $\mathbf{a}\mathbf{b}''$ is p-independent over C. In particular $\mathbf{a} \in pI(E)$, as required. Thus (4) $\implies$ (2). □

Proposition 2.2.9. *Let $E/D/C$ be a tower of fields and suppose that E/D and D/C are separable. Then E/C is separable.*

Proof. Actually we could have proved this from Proposition 2.1.22. Instead, we use Proposition 2.2.8. Let $\mathbf{a} \in pI(C)$. By separability of D/C, $\mathbf{a} \in pI(D)$. By separability of E/D, $\mathbf{a} \in pI(E)$. Thus E/C is separable. □

2.2.3 The component maps

We now work towards introducing component maps which are p-th roots of coordinate maps with respect to a certain linear basis. Component maps are used, for example, in the study of separably closed fields. They will enable us to describe the inseparable extensions that a field admits.

Monomials

Let $a \in pI(E/\mathbb{F})$. By definition of p-independence, the degree of a over E^p is p. Thus $\{a^i | i < p\}$ is an E^p-linear basis for $E^p(a)$. This can be generalised as follows.

Definition 2.2.10. Let $\mathbf{a} \subseteq E$. We set $M_{\mathbf{a}} := \left\{\prod_{a \in \mathbf{a}} a^{i(a)} | i \in {}^{\mathbf{a}}p = {}^{\mathbf{a}}\{0, ..., p-1\} \text{ has finite support}\right\}$ to be the set of *p-monomials of* $\mathbf{a}$.

Proposition 2.2.11. *Let $\mathbf{a} \subseteq E$. Then*

1. *$\mathbf{a} \in pI(E/C)$ if and only if $M_{\mathbf{a}}$ is an E^pC-linear base for $E^pC(\mathbf{a})$; and*
2. *$\mathbf{a} \in pB(E/C)$ if and only if $M_{\mathbf{a}}$ is an E^pC-linear base for E.*

Proof. Immediate from the usual Tower Lemma for bases of field extensions. □

The co-ordinate maps

We introduce the co-ordinate maps. Let $\mathbf{a} \in pI(E/C)$.

Definition 2.2.12. For each $m \in M_{\mathbf{a}}$, we let $\mu_m : E^pC(\mathbf{a}) \longrightarrow E^pC$ be the *co-ordinate map* with respect to the element m of the E^pC-linearly independent set $M_{\mathbf{a}}$. We let $\boldsymbol{\mu}^{\mathbf{a}} := \{\mu_m | m \in M_{\mathbf{a}}\}$ be the tuple of these maps, which are called the *co-ordinate maps of $M_{\mathbf{a}}$*. Thus $\boldsymbol{\mu}^{\mathbf{a}}$ is the unique tuple of maps such that, for each $b \in E^pC(\mathbf{a})$,

$$b = \sum_{m \in M_{\mathbf{a}}} \mu_m(b) m.$$

Note that the extension E/C and the number n are not explicitly mentioned in the notation, but will usually be obvious from the context.

These maps are well defined by Proposition 2.2.11.

Proposition 2.2.13. *Each of the maps $\mu \in \boldsymbol{\mu}^{\mathbf{a}}$ is E^pC-linear.*

The component maps

We now define an adaptation of the co-ordinate maps called *component maps*, which will help us to refine the link between p-independence and separability of field extensions, already discussed in Proposition 2.2.8. This motivation is further discussed in chapter 3, but for now we give the following simple example.

Example 2.2.14. If E is a perfect field then $\emptyset \in pB(E)$. There is only one element of $\mu \in \boldsymbol{\mu}^{\emptyset}$ and $\mu(b) = b$ for each $b \in E$. So our co-ordinate map is just the identity map and is not as useful as its p-th root $b \longmapsto b^{1/p}$ when trying to understand inseparable extensions.

This example suggests that component maps should be p-th roots of co-ordinate maps. However, there is a technicality which arises when we try to define components with respect to *relatively* p-independent tuples. We briefly illustrate this in the following example.

Example 2.2.15. Let $\mathbf{a} \in pI(E/C)$ be a tuple in E which is *relatively* p-independent over C. Then $\boldsymbol{\mu}^{\mathbf{a}}$ is the tuple of co-ordinate maps. Each map $\mu \in \boldsymbol{\mu}^{\mathbf{a}}$ (applied to the set $E^pC(\mathbf{a})$) has image E^pC which is not in general contained in E^p. Of course $E^pC \subseteq E^p$ if and only if $C \subseteq E^p$. Thus we cannot define the component maps to be p-th roots of the co-ordinate maps and still have image contained in E.

This problem does not arise when defining component maps with respect to an *absolutely* p-independent tuple $\mathbf{a}$ because the image of (each map in) $\boldsymbol{\mu}^{\mathbf{a}}$ (applied to the set $E^p(\mathbf{a})$) is E^p. The solution to this apparent problem is to avoid it! Given a relative p-independent tuple $\mathbf{a} \in pI(E/C)$, we adjoin a tuple $\mathbf{c} \in pI_{\max}(C;E)$ from C which is maximally p-independent in E; thus $\mathbf{ac} \in pI(E)$ is *absolutely* p-independent in E and the image of the tuple of co-ordinate maps $\boldsymbol{\mu}^{\mathbf{ac}}$ applied to $E^pC(\mathbf{a}) = E^p(\mathbf{ac})$ is E^p. Now we may successfully take p-th roots of the co-ordinate maps! The cost has been a non-canonical choice of the tuple $\mathbf{c}$. In practice, this cost will not matter and in fact we will often omit the tuple $\mathbf{c}$ from the notation.

This relies on the following two propositions, which are easy consequences of earlier statements.

Proposition 2.2.16. *Let* $\mathbf{c} \in pI_{\max}(C;E)$. *Then* $E^pC = E^p(\mathbf{c})$.

Proof. This is a special case (the absolute case) of Proposition 2.2.5. □

Proposition 2.2.17. *Let* $\mathbf{a} \in pI(E/C)$ *and let* $\mathbf{c} \in pI(C;E)$. *Then* $\mathbf{ac} \in pI(E)$.

Proof. This is immediate by using Proposition 2.2.4 and induction on the length of the tuple $\mathbf{a}$. $\square$

We are now in a position to define component maps. Recall that we denote by Φ the Frobenius map $x \longmapsto x^p$. First we define component maps with respect to tuples which are *absolutely* p-independent.

Definition 2.2.18. Let $\mathbf{a} \in pI(E)$ and let $m \in M_{\mathbf{a}}$. We define the *component map* with respect to the monomial m to be the map

$$\lambda_m := \Phi^{-1} \circ \mu_m : E \longrightarrow E.$$

This is well defined since the codomain of $\boldsymbol{\mu}^{\mathbf{a}}$ is E^p. We also let $\boldsymbol{\lambda}^{\mathbf{a}} := \{\lambda_m | m \in M_{\mathbf{a}}\}$ be the *tuple of component maps* with respect to $\mathbf{a}$.

Thus $\boldsymbol{\lambda}^{\mathbf{a}}$ is the unique tuple of maps $\lambda_m : E^p(\mathbf{a}) \longrightarrow E$ such that, for all $b \in E^p(\mathbf{a})$, we have

$$b = \sum_{m \in M_{\mathbf{a}}} \lambda_m(b)^p m.$$

In order to extend the definition of component maps to relatively p-independent tuples, we must rely on a non-canonical choice of a tuple $\mathbf{c} \in pI_{\max}(E/C)$, as outlined above.

Definition 2.2.19. Let $\mathbf{a} \in pI(E/C)$. We will frequently write $\boldsymbol{\lambda}^{\mathbf{a}} := \boldsymbol{\lambda}^{\mathbf{ac}}$ when there is no harm in doing so; i.e. when $\mathbf{c} \in pI_{\max}(C; E)$ is understood. We will only do this in situations where the precise choice of $\mathbf{c}$ does not matter.

2.2.4 The algebra of the component maps

We fix a field E. We will study algebraic properties of the tuple $\boldsymbol{\lambda}^{\mathbf{a}}$ of component maps, for tuples $\mathbf{a} \in pI(E)$.

Proposition 2.2.20. *Let $\mathbf{a} \in pI(E)$. Each of the maps $\lambda \in \boldsymbol{\lambda}^{\mathbf{a}}$ is an additive homomorphism.*

Proof. Let $\mu \in \boldsymbol{\mu}^{\mathbf{a}}$ be the co-ordinate maps corresponding to λ. By Proposition 2.2.13, μ is E^p-linear, thus certainly μ is an additive homomorphism. In exponential characteristic p the Frobenius map Φ is a ring-homomorphism. Thus $\lambda := \Phi^{-1} \circ \mu$ is an additive homomorphism. $\square$

Proposition 2.2.21. *Let $\mathbf{a} \in pI(E)$, let $\lambda \in \boldsymbol{\lambda}^{\mathbf{a}}$, and let $b, c \in E^p(\mathbf{a})$. Then $\lambda(bc) \in [\boldsymbol{\lambda}^{\mathbf{a}}(b), \boldsymbol{\lambda}^{\mathbf{a}}(c), \mathbf{a}]$.*

Proof. We calculate! Of course we may write $b = \sum_{m \in M_{\mathbf{a}}} \lambda_m(b)^p m$ and $c = \sum_{m \in M_{\mathbf{a}}} \lambda_m(c)^p m$. Then

$$\begin{aligned} bc &= \left(\sum_m \lambda_m(b)^p m \right) \left(\sum_m \lambda_m(c)^p m \right) \\ &= \sum_m \left(\sum_{ln = mm'^p} \lambda_l(b)\lambda_n(c)m' \right)^p m, \end{aligned}$$

where the second sum ranges through monomials $l, n \in M_{\mathbf{a}}$ such that we may find any monomial m' so that $ln = m'^p m$. Thus $\lambda_m(bc) = \sum_{ln=mm'^p} \lambda_l(b)\lambda_n(c)m'$. Since $M_{\mathbf{a}} \subseteq [\mathbf{a}]$, we must have that $\lambda_m(bc) \in [\boldsymbol{\lambda}^{\mathbf{a}}(b), \boldsymbol{\lambda}^{\mathbf{a}}(c), \mathbf{a}]$, as required. □

Finding co-ordinates and components

Lemma 2.2.22. *Let D be a field and suppose that $\mathbf{a}^p \in D$. Then $D(\mathbf{a}) = D \cdot M_{\mathbf{a}} := \bigoplus_{m \in M_{\mathbf{a}}} Dm$, meaning a direct sum of D-vector spaces.*

Proof. Elementary. Since $\mathbf{a}^p \in D$, $D(\mathbf{a})/D$ is purely inseparable of exponent p (or 1) and the degree of $\mathbf{a}$ over D is no more than $p|\mathbf{a}|$. Therefore $M(\mathbf{a})$ D-spans $D(\mathbf{a})$. □

Lemma 2.2.23. *Let $\mathbf{a} \in pI(E/C)$, let $X \subseteq E^p C$, and let $b \in X \cdot M_{\mathbf{a}}$. Then $\boldsymbol{\mu}^{\mathbf{a}}(b) \subseteq X$.*

Proof. Since $\mathbf{a} \in pI(E/C)$ and by Proposition 2.2.11, $M_{\mathbf{a}}$ is an E^pC-linear base for $E^pC(\mathbf{a})$ and is certainly E^pC-linearly independent. Since $X \subseteq E^pC$, $M_{\mathbf{a}}$ is therefore X-linearly independent and $\boldsymbol{\mu}^{\mathbf{a}}(b) \subseteq X$, as required. □

Proposition 2.2.24. *Let $\mathbf{a} \in pI(E)$, let D be a subfield such that $\mathbf{a}^p \in D \subseteq E^p$, and let $b \in D(\mathbf{a})$. Then $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq D^{1/p}$.*

Proof. By Lemma 2.2.22 and Lemma 2.2.23, we have that $\boldsymbol{\mu}^{\mathbf{a}}(b) \subseteq D \subseteq E^p$. Thus Φ^{-1} is well-defined on $\boldsymbol{\mu}^{\mathbf{a}}(b)$ and we have that $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq D^{1/p}$. □

Proposition 2.2.25. *Let $\mathbf{a} \in pI(E)$ and let $\mathbf{b} \subseteq E^p(\mathbf{a})$ be in the p-span of $\mathbf{a}$ in E. Then*

1. $\boldsymbol{\lambda}^{\mathbf{a}}([\mathbf{b}]) \subseteq [\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a}]$, *and*
2. $\boldsymbol{\lambda}^{\mathbf{a}}((\mathbf{b})) \subseteq (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a})$.

Proof. 1. This is a simple induction using Proposition 2.2.20 and Proposition 2.2.21.

2. We note that the field of fractions of $[\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a}]^p(\mathbf{a})$ is $(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a})^p(\mathbf{a})$. By part 1

$$[\mathbf{b}] \subseteq [\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a}]^p(\mathbf{a}).$$

Passing to the field of fractions on each side, we find that

$$(\mathbf{b}) \subseteq (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a})^p(\mathbf{a}).$$

Applying Proposition 2.2.24 with $D := (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a})^p$, we find that $\boldsymbol{\lambda}^{\mathbf{a}}((\mathbf{b})) \subseteq D^{1/p} = (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})\mathbf{a})$.

□

Relationships between component maps

Proposition 2.2.26. *Let $\mathbf{a}, \mathbf{b} \in pI(E/C)$ and let $c \in E^pC(\mathbf{a}) \cap E^pC(\mathbf{b})$ be in the p-span of both $\mathbf{a}, \mathbf{b}$ in E over C.*

1. *If $\mathbf{b} \subseteq E^pC(\mathbf{a})$, then $\boldsymbol{\lambda}^{\mathbf{a}}(c) \subseteq (\boldsymbol{\lambda}^{\mathbf{b}}(c), \boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b}), \mathbf{a})$;*
2. *if $E^pC(\mathbf{a}) = E^pC(\mathbf{b})$, then $(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})) = (\boldsymbol{\lambda}^{\mathbf{b}}(\mathbf{a}))$; and*
3. *if $E^pC(\mathbf{a}) = E^pC(\mathbf{b})$, then $\boldsymbol{\lambda}^{\mathbf{a}}(c) \subseteq (\boldsymbol{\lambda}^{\mathbf{b}}(\mathbf{a}c), \mathbf{a})$.*

Proof. 1. We write c as a linear sum along the base $M_{\mathbf{b}}$.

$$c = \sum_{m \in M_{\mathbf{b}}} \lambda_m(c)^p \cdot m.$$

Since $\mathbf{b} \subseteq E^pC(\mathbf{a})$, then $\boldsymbol{\lambda}^{\mathbf{a}}$ is defined on $M_{\mathbf{b}}$. Thus we re-write the monomials $m \in M_{\mathbf{b}}$ as E^pC-linear sums with respect to the basis $M_{\mathbf{a}}$.

$$\begin{aligned}
&= \sum_{m \in M_{\mathbf{b}}} \lambda_m(c)^p \cdot \left(\sum_{n \in M_{\mathbf{a}}} \lambda_n(m)^p \cdot n \right) \\
&= \sum_{m \in M_{\mathbf{b}}} \sum_{n \in M_{\mathbf{a}}} \lambda_m(c)^p \cdot \lambda_n(m)^p \cdot n \\
&= \sum_{n \in M_{\mathbf{a}}} \sum_{m \in M_{\mathbf{b}}} \lambda_m(c)^p \cdot \lambda_n(m)^p \cdot n \\
&= \sum_{n \in M_{\mathbf{a}}} \left(\sum_{m \in M_{\mathbf{b}}} \lambda_m(c)^p \cdot \lambda_n(m)^p \right) \cdot n \\
&= \sum_{n \in M_{\mathbf{a}}} \left(\sum_{m \in M_{\mathbf{b}}} \lambda_m(c) \cdot \lambda_n(m) \right)^p \cdot n.
\end{aligned}$$

Thus

$$\begin{aligned}
\lambda_n(c) &= \sum_{m \in M_{\mathbf{b}}} \lambda_m(c) \cdot \lambda_n(m) \\
&\in (\boldsymbol{\lambda}^{\mathbf{b}}(c), \lambda_n(M_{\mathbf{b}}))
\end{aligned}$$

and

$$\boldsymbol{\lambda}^{\mathbf{a}}(c) \in (\boldsymbol{\lambda}^{\mathbf{b}}(c), \boldsymbol{\lambda}^{\mathbf{a}}(M_{\mathbf{b}})).$$

By Proposition 2.2.21 and since $M_{\mathbf{b}} \subseteq [\mathbf{b}]$, we have that $\boldsymbol{\lambda}^{\mathbf{a}}(M_{\mathbf{b}}) \subseteq [\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b}), \mathbf{a}]$. Thus $\boldsymbol{\lambda}^{\mathbf{a}}(c) \in (\boldsymbol{\lambda}^{\mathbf{b}}(c), \boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b}), \mathbf{a})$, as required.

2. Let L be the matrix representing the identity map on $E^pC(\mathbf{a}) = E^pC(\mathbf{b})$ written with respect to the bases $M_{\mathbf{a}}$ and $M_{\mathbf{b}}$. Clearly L is invertible. Furthermore, the coefficients of L are the $\mathbf{b}$-components of $\mathbf{a}$, i.e. the tuple $\boldsymbol{\lambda}^{\mathbf{b}}(\mathbf{a})$; and the coefficients of the inverse L^{-1} are the $\mathbf{a}$-components of $\mathbf{b}$, i.e. $\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b})$. Finally, the coefficients of L^{-1} are contained in the field generated by the coefficients of L.

3. We apply parts 1. and 2.

□

2.3 Factorising a field extension

Let E/C be a field extension. In this section we study *factorisations* of E/C by which we mean intermediate fields D such that E/D has some specified property and D/C has some other. This is intended as motivation for our later study of relative inseparable closures. Everything in this section is well known and most of it can be found in Section 6, Chapter V of [17].

2.3.1 Purely inseparable over separable

One way of understanding an inseparable extension is to factorise it into a separable extension over a purely inseparable extension. First we tackle algebraic extensions.

Proposition 2.3.1. *Let E/C be an algebraic extension of fields. There exists a unique subextension D such that*

- *E/D is purely inseparable and*
- *D/C is separably algebraic.*

Proof. Let D be the compositum of the separably algebraic subextensions D/C of E/C. By Fact 2.1.2, D/C is separably algebraic and E/D is purely inseparable. The uniqueness is obvious from the definition of D. □

We can extend this to non-algebraic extensions, although we lose the uniqueness.

Proposition 2.3.2. *Let E/C be an extension of fields. There exists a subextension D such that*

- *E/D is purely inseparable and*
- *D/C is separable.*

Proof. Let $\mathbf{e} \subseteq E$ be a transcendence base for E/C. Then $E/C(\mathbf{e})$ is algebraic and, by Proposition 2.3.1, there exists $D \supseteq C(\mathbf{e})$ such that E/D is purely inseparable and $D/C(\mathbf{e})$ is separably algebraic. Observe that $C(\mathbf{e})/C$ is a separably generated extension; thus, by Corollary 2.1.23, $C(\mathbf{e})/C$ is separable. By Proposition 2.2.9, a tower of separable extensions is separable. Thus the extension D/C is separable, as required. □

2.3.2 A partial 'inverse'

Much of what follows in section 3.2 and chapter 4 is motivated by the problem of trying to 'invert' Proposition 2.3.2 by finding a subextension D such that

- E/D is separable and
- D/C is purely inseparable.

In fact, it is not always possible to find such a D even for algebraic extensions E/C. However in the case of normal algebraic extensions, there is the following result. The following proposition appears as Proposition 6.11 in Chapter V of [17] where more details of the proof can also be found.

Proposition 2.3.3. *Let E/C be an algebraic normal extension of fields. Then there exists a subextension D such that*

- *E/D is separably algebraic and*
- *D/C is purely inseparable.*

Proof. We define $D := \mathrm{Fix}(\mathrm{Aut}(E/C))$. Since $\mathrm{Aut}(D/C)$ is trivial, D/C is purely inseparable. Furthermore, it turns out that E/D is separable. □

If we drop the second requirement in the short list above, and instead look for the 'least' D such that E/D is separable; then in fact there does always exist such a field D. We will explore this in chapter 3.

Chapter 3

The problem of inseparability

In this chapter we investigate and solve the problem posed at the end of the last chapter: given any field extension E/C we aim to find an intermediate field D such that E/D is separable and D is 'minimal' with this property. We will call such a field D a *relative inseparable closure* of C in E, since the only inseparable subextension of E/D is trivial.

Fix a field extension E/C and let $\mathrm{Int}(E/C)$ the set of intermediate fields of E/C. Before motivating the study of relative inseparable closures, we make the definition a little more precise, including explaining the word 'minimal'.

Definition 3.0.4. We say that $D \in \mathrm{Int}(E/C)$ is a *relative inseparable closure* of C in E if

1. E/D is separable and
2. for all $D' \in \mathrm{Int}(E/C)$, if E/D' is separable then $D \subseteq D'$.

It is not clear *a priori* that such a field exists. If it does exist then it is obviously unique.

3.1 Seeking the relative inseparable closure

First we give a little motivation for why we seek the relative inseparable closure. In the first subsection we explain how it relates to the idea of a 'relative perfect hull', and in the second subsection we explain its importance for our later applications.

3.1.1 Two relative versions of the perfect hull

We say that a field is *inseparably closed* if it has no proper inseparable extensions and that it is *purely-inseparably closed* if it has no proper purely-inseparable extensions. Since all purely-inseparable extensions are inseparable but the converse is far from true; we might imagine that the former property is strictly stronger than the latter. However, this is not at all the case. Recall that a field is said to be *perfect* if all of its extensions are separable, i.e. if it is inseparably closed. But a field E is perfect if and only if $\mathrm{char}(E) = 0$ or $E^p = E$, by Proposition 2.1.26. Thus a purely inseparably closed field is inseparably closed; and the two properties are equivalent. This observation is the content of the following proposition, which (although almost self-evident) must be considered quite an important co-incidence.

Proposition 3.1.1. *The following properties of a field E are equivalent.*

1. *E is perfect.*
2. *E is inseparably closed.*
3. *E is purely inseparably closed.*

Closing a field E under all p-th roots, we obtain the field $E^{\mathrm{perf}} := \bigcup_{n\in\mathbb{N}} E^{p^{-n}}$. This is called the *perfect hull* of E. The perfect hull is the smallest perfect field containing E.

Like other notions of closure, we may try to understand how to 'relativise' the perfect hull. Fix an extension E/C. We investigate how to find something akin to the perfect hull of C which lies within a given field E; but precisely what we mean by this depends on which of the two properties of perfect fields we wish to relativise. Thus there are two questions to answer.

1. What is the relative inseparable closure of C in E (as in Definition 3.0.4)?
2. What is the relative purely-inseparable closure of C in E?

And perhaps a third.

3. Do the notions of relative inseparable closure and relative purely-inseparable closure coincide?

The second questions is the easiest to answer.

Proposition 3.1.2. *$E \cap C^{p^{-\infty}}$ is the relative purely-inseparable closure of C in E.*

Proof. It is obvious that any proper extension of $C^{\text{perf}} \cap E$ lying inside E cannot be purely-inseparable and that $C^{\text{perf}} \cap E$ is the least extension of C with this property. □

However if D is the relative purely-inseparable closure of E/C, it is well known that E/D need not be separable, even when E/C is algebraic. Identifying the relative inseparable closure is not quite such an easy task.

3.1.2 The Implicit Function Theorem and separable extensions

Our wider goal is the study of existentially definable sets in t-henselian fields. In particular, we aim to understand their local behaviour using the Implicit Function Theorem for polynomials. (See section 5.1 for definitions.)

Let F/C be a field extension and suppose that we may apply the Implicit Function Theorem in F. Let $f \in C[xy]$ be a polynomial and let $ab \in Z(f)$ be a zero. The set defined in F by the formula

$$\exists y \; f \doteq 0$$

is the projection of $Z(f)$, the set of zeroes of f, and can be studied by applying the Implicity Function Theorem. If $D_y f(ab) \neq 0$ then $Z(f)$ is *locally* the graph of a continuous function with domain an open set around a.

The crucial assumption is that $D_y f(ab) \neq 0$. If b were separably algebraic over $C(a)$ and f were the minimal polynomial of b over $C(a)$ (with the denominators (in a) suitably multiplied out) then indeed $D_y f(ab) \neq 0$.

In section 5.2 we will prove the link between the minimal polynomial of b over $C(a)$ and other polynomials of which ab is a zero. In this way, the local behaviour of any algebraic set containing ab is controlled by the local behaviour of the relevant minimal polynomial. These minimal polynomials can, in turn, be understood locally using the Implicit Function Theorem if $C(ab)/C(a)$ is separable. A further complication will be that we have to consider an arbitrary finite number of existential quantifiers.

Thus the problem is redued to being able to 'translate' an inseparable extension into a separable one, but in such a way that we can recover a local understanding of the original algebraic set. So far, we've discussed two methods of finding a separable extension from an arbitrary extension E/C:

1. factorise E/C by finding D such that E/D is purely inseparable and D/C is separable, as in Proposition 2.3.2; or

2. let D be the relative inseparable closure of C in E; then E/D is separable.

The problem with the first method is that we cannot recover a local understanding of the original set, as the following example shows.

Example 3.1.3. Let F be a field in which we can apply the Implicit Function Theorem. Let $f := x - y^p$ be a polynomial (defined over $C = \mathbb{F}_p$) and let $ab \in Z(f)$ so that $a = b^p$. Since $D_y f(ab) = 0$, we cannot apply the Implicit Function Theorem directly. The problem is the inseparability of the extension $C(ab)/C(a)$. In this case the extension is purely inseparable. If we factorise the extension (as in 1, above), we find D such that $C(ab)/D$ is purely inseparable and $D/C(a)$ is separable. Of course in this very simple example, $D = C(a)$ and we haven't achieved anything.

In chapter 4, we prove the the relative inseparable closure is the right tool to use: we can recover a local understanding of the original algebraic set using the method of Λ-alteration.

3.1.3 The paper of Deveney and Mordeson

The existence of the relative inseparable closure is contained in the 1977 paper of Deveney and Mordeson, [7].

Proposition 3.1.4. *(cf. Theorem 1.1, [7]) The relative inseparable closure of C in E exists.*

We will sketch the proof given in [7], but first we give a simple definition.

Definition 3.1.5. Let $D \in \mathrm{Int}(E/C)$. We say that D is *coseparable* in E if E/D is separable.

In [7] the relative inseparable closure (which is called C^* in the notation of that paper) is constructed as the intersection of all intermediate fields $D \in \mathrm{Int}(E/C)$ which are coseparable in E. Since E/E is trivially separable, E is coseparable in itself and the intersection is taken over a non-empty family; thus it is well-defined. Deveney and Mordeson prove that the intersection of coseparable intermediate fields is coseparable by arguing that the intersection of a family of extensions which are linearly disjoint from a fixed extension is linearly disjoint from the same fixed extension (see the proof of Theorem 1.1 in [7] for details). Therefore C^* really is a relative inseparable closure. The uniqueness also follows since C^* is an intersection: the intersetion of any two relative inseparable closures must

also be a relative inseparable closure, thus they must be equal.

That the intersection of a family of coseparable fields is coseparable can be proved by an alternative argument using p-independence, which we give here to illustrate the techniques used later in this chapter.

Lemma 3.1.6. *Let $(D_i|i \in I)$ be the family of coseparable intermediate fields of E/C. We claim that $D := \bigcap\{D_i|i \in I\}$ is coseparable in E.*

Proof. By Proposition 3.2.1 (which has not yet been stated or proved), it suffices to prove that D is closed under the component maps with respect to tuples from D which are p-independent in E. Let $\mathbf{a} \in pI(E) \cap D$ and let $b \in D \cap E^p(\mathbf{a})$. Then, for each $i \in I$, $\mathbf{a} \in pI(E) \cap D_i$ and $b \in D_i \cap E^p(\mathbf{a})$. By the separability of E/D_i, $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq D_i$. This holds for each $i \in I$, so $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq D$. Thus D is closed under the relevant maps, so E/D is separable. □

Now we can give an alternative proof of Proposition 3.1.4.

Proof. This follows from Lemma 3.1.6 in the same way as in [7]. By Lemma 3.1.6, D is coseparable. Let $D' \in \mathrm{Int}(E/C)$ be coseparable in E. Then $D \subseteq D'$. Thus D satisfies both of the required properties. □

However, for our purposes (which will be further elaborated on in chapter 4) we require a much more constructive method of finding the relative inseparable closure. The method of Λ-closure is our solution to this problem.

3.2 Λ-closure

The goal of this section is to constructively describe the relative inseparable closure using the component maps, which were defined in the previous chapter. Since these component maps are usually denoted by λ, we call the subfield so constructed the *Λ-closure*. We develop the construction of the Λ-closure and prove that the Λ-closure of D in E is the relative inseparable closure, as defined in Definition 3.0.4. The process of Λ-closure is a very constructive one, which is important for our applications to the study of existential definability. The construction is a recursion up to the ordinal ω.

Later in the section we will suppose that D is generated (intially, not necessarily finitely) over a subfield C which is already Λ-closed in E. The construction of the Λ-closure will simplify dramatically: we can express it in terms of the generators of D over C. This will be very useful later in applications.

Lastly, we suppose in addition that E is finitely generated over D. The construction of the Λ-closure simplifies again to be a finite recursion.

3.2.1 Reinterpreting separability II: Λ-closedness

Let us now come to our first goal: to reinterpret the notion of separability of a field extension in terms of the extension being Λ-closed. In Proposition 2.2.8 we saw a link between p-independence and the separability of field extensions. This first proposition extends this relationship to include the component maps associated to p-independent tuples. This motivates the subsequent definition of Λ-closedness.

Recall that $pI(D;E)$ denotes the set of tuples from D which are absolutely p-independent in E.

Proposition 3.2.1. *For a field extension E/D, the following are equivalent.*

1. *E/D is separable,*
2. *for each $\mathbf{a} \in pI(D;E)$, we have that $E^p(\mathbf{a}) \cap D = D^p(\mathbf{a})$, and*
3. *for each $\mathbf{a} \in pI(D;E)$, we have that $\boldsymbol{\lambda}^{\mathbf{a}}(E^p(\mathbf{a}) \cap D) \subseteq D$.*

Proof. By Proposition 2.2.8, the separability of E/D is equivalent to the inclusion $pI(D) \subseteq pI(E)$.

(1) $\longrightarrow$ (2) Observe that for any tuple $\mathbf{a} \subseteq D$, we trivially have that $E^p(\mathbf{a}) \cap D \supseteq D^p(\mathbf{a})$. We aim to prove the converse inclusion. Let $\mathbf{a} \in pI(D;E)$ and let $b \in (E^p(\mathbf{a}) \cap D) \setminus D^p(\mathbf{a})$. Then $\mathbf{a}b \in pI(D) \setminus pI(E)$, but this contradicts $pI(D) \subseteq pI(E)$.

(2) $\longrightarrow$ (1) We prove that $pI(D) \subseteq pI(E)$ by induction on the length of members of $pI(D)$. The empty tuple is a member of $pI(D) \cap pI(E)$. We let $\mathbf{a}a \in pI(D)$ and suppose as an inductive hypothesis that $\mathbf{a} \in pI(E)$. Since $\mathbf{a}a \in pI(D)$, we have $a \notin D^p(\mathbf{a}) = E^p(\mathbf{a}) \cap D$; and since $a \notin D$, we have $a \notin E^p(\mathbf{a})$. Thus $\mathbf{a}a \in pI(E)$, as required.

(2) $\longrightarrow$ (3) Let $\mathbf{a} \in pI(D;E)$. Observe that by swapping D in the statement of Proposition 2.2.24 with D^p, we have $\boldsymbol{\lambda}^{\mathbf{a}}(D^p(\mathbf{a})) \subseteq D$. Then $\boldsymbol{\lambda}^{\mathbf{a}}(E^p(\mathbf{a}) \cap D) = \boldsymbol{\lambda}^{\mathbf{a}}(D^p(\mathbf{a})) \subseteq D$, by assumption.

(3) $\longrightarrow$ (2) Let $\mathbf{a} \in pI(D;E)$. As before the inclusion $E^p(\mathbf{a}) \cap D \supseteq D^p(\mathbf{a})$ is trivially true. Aiming to prove the converse inclusion, we let $b \in E^p(\mathbf{a}) \cap D$. Using our assumption that $\boldsymbol{\lambda}^{\mathbf{a}}(E^p(\mathbf{a}) \cap D) \subseteq D$, we have that $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq D$. Thus $b \in (\boldsymbol{\lambda}^{\mathbf{a}}(b))^p(\mathbf{a}) \subseteq D^p(\mathbf{a})$, as required. □

3.2.2 Λ-closedness and Λ-closure

By the equivalence of the first and third statements in Proposition 3.2.1, E/D is separable if and only if D is closed under the component maps $\boldsymbol{\lambda}^{\mathbf{a}}$, for all $\mathbf{a} \in pI(D;E)$. This motivates the following definition.

Definition 3.2.2. We say that D is *Λ-closed* in E if, for all $\mathbf{a} \in pI(D;E)$, we have that $\boldsymbol{\lambda}^{\mathbf{a}}(E^p(\mathbf{a}) \cap D) \subseteq D$.

Note that, for $\mathbf{a} \in pI(D;E)$, the domain of the tuple of maps $\boldsymbol{\lambda}^{\mathbf{a}}$ is $E^p(\mathbf{a})$.

Corollary 3.2.3. *An extension E/D is separable if and only if D is Λ-closed in E.*

Proof. This is a simple rephrasing of the equivalence of the first and third statements from Proposition 3.2.1. □

Any property which may be expressed as being *closed* under some maps automatically suggests a *closure* operation.

Definition 3.2.4. We say that a subfield $\boldsymbol{\Lambda} \in \mathrm{Int}(E/C)$ is a *Λ-closure* of C in E if

1. $\boldsymbol{\Lambda}$ is Λ-closed and
2. for all $D \in \mathrm{Int}(E/C)$, if E/D is Λ-closed then $\boldsymbol{\Lambda} \subseteq D$.

As with relative inseparable closure, it is obviously the case that the Λ-closure is unique, if it exists. This definition allows us to reinterpret our goal of finding the relative inseparable closure.

Proposition 3.2.5. *Let E/C be any extension and let $\boldsymbol{\Lambda} \in \mathrm{Int}(E/C)$. Then $\boldsymbol{\Lambda}$ is the Λ-closure of E/C if and only if $\boldsymbol{\Lambda}$ is the relative inseparable closure of E/C.*

Proof. The equivalence of Definition 3.2.4 and Definition 3.0.4 is obvious in the light of Corollary 3.2.3. □

3.2.3 General construction of Λ-closure

We seek a construction of the relative inseparable closure $\boldsymbol{\Lambda}$ of D in E. By Proposition 3.2.5, our task is equal to constructing the Λ-closure of D in E. Specifically, we are seeking the *least* intermediate field $\boldsymbol{\Lambda} \in \mathrm{Int}(E/D)$ such that $\boldsymbol{\Lambda}$ is closed under the component maps $\boldsymbol{\lambda}^{\mathbf{a}}$, for all possible tuples $\mathbf{a} \in pI(D;E)$.

With any such notion of closure, the usual way to try to construct it is to adjoin whatever it is that may be lacking. In this situation, the images of the maps $\boldsymbol{\lambda}^{\mathbf{a}}$ may be lacking; so we adjoin them. However, this construction is not complete after one step: there may now be more tuples $\mathbf{a}$ which are p-independent in E, and thus more maps $\boldsymbol{\lambda}^{\mathbf{a}}$ which respect to which our subfield needs to be closed. Therefore, we proceed recursively. By the finite character of the generation of fields (i.e. the field generated by a set X is the union of the fields generated by finite subsets of X), we can stop the recursion after the ordinal ω.

Definition 3.2.6. Let

$$\Lambda(E/D) := D(\lambda(d) | d \in E^p(\mathbf{a}) \cap D, \lambda \in \boldsymbol{\lambda}^{\mathbf{a}}, \mathbf{a} \in pI(D;E))$$

be the field generated over D by the image of D under the component maps with respect to all tuples from D which are p-independent in E.

Definition 3.2.7. We give the recursive construction of an increasing chain of subfields of E which contain D.

1. First we denote $\Lambda_0(E/D) := D$.
2. Let
$$\Lambda_{n+1}(E/D) := \Lambda(E/\Lambda_n D).$$

Proposition 3.2.8. *The sequence $(\Lambda_n(E/D))_{n<\omega}$ is an increasing chain of subfields of E.*

Definition 3.2.9. We let

$$\boldsymbol{\Lambda}(E/D) := \bigcup_{n<\omega} \Lambda_n(E/D).$$

Of course, we now aim to show that $\boldsymbol{\Lambda}(E/D)$ really is the Λ-closure of D in E. This involves proving both that $\boldsymbol{\Lambda}(E/D)$ is Λ-closed and that it is contained in all other Λ-closed $D' \in \mathrm{Int}(E/D)$.

Lemma 3.2.10. *$\boldsymbol{\Lambda}(E/D)$ is Λ-closed in E.*

Proof. Let $\mathbf{a} \in pI(\boldsymbol{\Lambda}(E/D); E)$ and let $b \in E^p(\mathbf{a}) \cap \boldsymbol{\Lambda}(E/D)$. Let $\mathbf{a}' \subseteq \mathbf{a}$ be a finite subtuple of $\mathbf{a}$ such that $b \in E^p(\mathbf{a}')$. By Proposition 3.2.1, we must show that $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq \boldsymbol{\Lambda}(E/D)$. There exists $n < \omega$ such that $b \in \Lambda_n(E/D)$ and there exists $n' < \omega$ such that $\mathbf{a}' \subseteq \Lambda_{n'}(E/D)$. Let $N := \max\{n, n'\}$. Then $\boldsymbol{\lambda}^{\mathbf{a}'}(b) \subseteq \Lambda(E/\Lambda_N(E/D)) = \Lambda_{N+1}(E/D)$, by Definition 3.2.7. Since $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq \boldsymbol{\lambda}^{\mathbf{a}'}(b) \cup \{0\}$, we have that $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq \Lambda_{N+1}(E/D) \subseteq \boldsymbol{\Lambda}(E/D)$, as required. Thus $\boldsymbol{\Lambda}(E/D)$ is Λ-closed. □

Lemma 3.2.11. *Let $D' \in \mathrm{Int}(E/D)$ be Λ-closed in E. Then $\Lambda(E/D) \subseteq D'$.*

Proof. Let $\mathbf{a} \in pI(D; E) \subseteq pI(D'; E)$ and let $b \in E^p(\mathbf{a}) \cap D \subseteq E^p(\mathbf{a}) \cap D'$. Since D' is Λ-closed in E, we have $\boldsymbol{\lambda}^{\mathbf{a}}(b) \subseteq D'$. Thus $\Lambda(E/D) = D(\boldsymbol{\lambda}^{\mathbf{a}}(b) | \mathbf{a} \in pI(E) \cap D, b \in E^p(\mathbf{a}) \cap D) \subseteq D'$. □

Lemma 3.2.12. *Let E/D be a field extension and let $D' \in \mathrm{Int}(E/D)$ be Λ-closed in E. Then $\boldsymbol{\Lambda}(E/D) \subseteq D'$.*

Proof. By Lemma 3.2.11 and an easy induction, for each $n < \omega$ we have that $\Lambda_n(E/D) \subseteq D'$. Thus $\boldsymbol{\Lambda}(E/D) = \bigcup_{n<\omega} \Lambda_n(E/D) \subseteq D'$. □

Proposition 3.2.13. *$\boldsymbol{\Lambda}(E/D)$ is the Λ-closure of D in E.*

Proof. By Lemma 3.2.10, $\boldsymbol{\Lambda}(E/D)$ is Λ-closed in E; and, by Lemma 3.2.11, $\boldsymbol{\Lambda}(E/D)$ is contained in any $D' \in \mathrm{Int}(E/D)$ which is Λ-closed in E. Thus $\boldsymbol{\Lambda}(E/D)$ is the Λ-closure of D in E. □

3.2.4 Generated extensions: simplifying the construction

Notice that in the definition of $\Lambda(E/D)$ we adjoin to D, for each tuple $\mathbf{a} \in pI(D; E)$, the image of $E^p(\mathbf{a}) \cap D$ under the tuple of maps $\boldsymbol{\lambda}^{\mathbf{a}}$. Thus we have quantified both over tuples $\mathbf{a} \in pI(D; E)$ and then over elements of $E^p(\mathbf{a}) \cap D$. This is not very useful in practice and there is a lot of redundancy. We can replace both quantifications by instead considering a single maximal p-independent tuple $\mathbf{a}$ and a set of generators $\mathbf{b}$ of D. This will greatly improve the construction.

In the next proposition, we show that the extension $\Lambda(E/D)$ can be constructed by closing under components with respect to any given tuple of D which is maximal with respect to being p-independent in E.

Proposition 3.2.14. *Let $\mathbf{a} \in pI_{\max}(D; E)$. Then $\Lambda(E/D) = D(\boldsymbol{\lambda}^{\mathbf{a}}(D))$.*

Proof. Observe first that $D \subseteq E^p(\mathbf{a})$ so that the domain of $\boldsymbol{\lambda}^{\mathbf{a}}$ is $E^p(\mathbf{a}) \cap D = D$. It is clear that $\Lambda(E/D) \supseteq D(\boldsymbol{\lambda}^{\mathbf{a}}(D))$. Let $\mathbf{b} \in pI(D;E)$ and let $c \in E^p(\mathbf{b}) \cap D$. We want to show that $\boldsymbol{\lambda}^{\mathbf{b}}(c) \in D(\boldsymbol{\lambda}^{\mathbf{a}}(D))$. We might as well assume that $\mathbf{b} \in pI_{\max}(D;E)$. In fact we have that $E^p(\mathbf{a}) = E^p(\mathbf{b}) = E^pD$. By part 3 of Proposition 2.2.26 we have that

$$\boldsymbol{\lambda}^{\mathbf{b}}(c) \subseteq (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{b}c), \mathbf{b}) \subseteq D(\boldsymbol{\lambda}^{\mathbf{a}}(D)).$$

Thus $\lambda(E/D) \subseteq D(\boldsymbol{\lambda}^{\mathbf{a}}(D))$, as required. □

Now we suppose that D is generated by a tuple $\mathbf{d}$; since we have so far not insisted that tuples be finite, there is nothing stopping $\mathbf{d}$ being the tuple of all elements of D! We show that the extension $\Lambda(E/D)$ is generated by $\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d})$.

Remark 3.2.15. Let $\mathbf{a} \in pI(E)$ and let $\mathbf{d} \subseteq E^p(\mathbf{a})$ be any tuple. By Proposition 2.2.25, we have that $\boldsymbol{\lambda}^{\mathbf{a}}((\mathbf{d})) \subseteq (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d}), \mathbf{a})$. We now apply this to our task of simplifying the construction of the Λ-closure.

Proposition 3.2.16. *Let* $\mathbf{a} \in pI_{\max}(D;E)$*, and let* $\mathbf{d}$ *be such that* $D = (\mathbf{d})$*. Then* $\Lambda(E/D) = D(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d}))$*.*

Proof. Since $\mathbf{a}$ is a maximal tuple from D which is p-independent in E, $D \subseteq E^p(\mathbf{a})$. In particular, $\mathbf{d} \subseteq E^p(\mathbf{a})$. Therefore we may apply Proposition 2.2.25 (as in the remark above) to find that:

$$\begin{aligned} \boldsymbol{\lambda}^{\mathbf{a}}(D) &= \boldsymbol{\lambda}^{\mathbf{a}}((\mathbf{d})) \\ &\subseteq (\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d})\mathbf{a}) \\ &\subseteq D(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d})). \end{aligned}$$

Since $\mathbf{d} \subseteq D$, it is obvious that $\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d}) \subseteq \boldsymbol{\lambda}^{\mathbf{a}}(D)$. Therefore $D(\boldsymbol{\lambda}^{\mathbf{a}}(D)) \subseteq D(D(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d}))) = D(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d})) \subseteq D(\boldsymbol{\lambda}^{\mathbf{a}}(D))$. Using Proposition 3.2.14, we have that $\Lambda(E/D) = D(\boldsymbol{\lambda}^{\mathbf{a}}(D)) = D(\boldsymbol{\lambda}^{\mathbf{a}}(\mathbf{d}))$, as required. □

3.2.5 Splitting points

Definition 3.2.17. Let E/C be a separable field extension and let $\mathbf{a} \subseteq E$. We say that a partition $\mathbf{a} = \mathbf{a}_p\mathbf{a}_s$ is a *splitting* of $\mathbf{a}$ over C if $\mathbf{a}_p \in pI_{\max}(\mathbf{a}; E/C)$.

Proposition 3.2.18. *Let* $\mathbf{a} \subseteq E$*. We may find a splitting of* $\mathbf{a}$ *over* C*. Furthermore* $\mathbf{a}_p$ *will be algebraically independent over* C*.*

Proof. Simply choose $\mathbf{a}_p$ to be any subtuple of $\mathbf{a}$ which is maximal with respect to the property of being p-independent in E over C. Note that the empty tuple is always p-independent in and over any field; thus such a maximal subtuple exists. □

Proposition 3.2.19. *Suppose that E/C is Λ-closed and let $\mathbf{a} \subseteq E$. Let $\mathbf{a} = \mathbf{a}_p\mathbf{a}_s$ be a splitting of $\mathbf{a}$ over C. Then*

$$\Lambda(E/C(\mathbf{a})) = C(\mathbf{a}_p, \boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a}_s)).$$

Proof. First we observe that the notation $\boldsymbol{\lambda}^{\mathbf{a}_p}$ implicitely includes mention of a tuple $\mathbf{c} \in pB(C)$. It won't matter which one we choose, so let us fix a tuple $\mathbf{c} \in pB(C)$ and remember that $\boldsymbol{\lambda}^{\mathbf{a}_p} = \boldsymbol{\lambda}^{\mathbf{a}_p\mathbf{c}}$. Thus $\boldsymbol{\lambda}^{\mathbf{a}_p}(C) = \boldsymbol{\lambda}^{\mathbf{a}_p\mathbf{c}}(C) = \boldsymbol{\lambda}^{\mathbf{c}}(C) = C$. Also $\boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a}_p) = \{0, 1\}$, thus $\boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a}) = \{0, 1\} \cup \boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a}_s)$. Set $D := C(\mathbf{a})$. By Proposition 3.2.16

$$\begin{aligned}
\Lambda(E/D) &= D(\boldsymbol{\lambda}^{\mathbf{a}_p}(D)) \\
&= C(\mathbf{a})(\boldsymbol{\lambda}^{\mathbf{a}_p}(C(\mathbf{a}))) \\
&= C(\mathbf{a}, \boldsymbol{\lambda}^{\mathbf{a}_p}(C), \boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a})) \\
&= C(\mathbf{a}, \boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a}_s)) \\
&= C(\mathbf{a}_p, \boldsymbol{\lambda}^{\mathbf{a}_p}(\mathbf{a}_s)).
\end{aligned}$$

□

3.2.6 Recursively splitting points

Given $\mathbf{a}$, we aim to understand the Λ-closure of $C(\mathbf{a})$ in E. We will recursively apply the idea of splitting from the previous subsection and define a sequence of tuples $(\mathbf{a}^n)_{n<\omega}$ such that $C(\mathbf{a}^n | n < \omega)$ is equal to the Λ-closure.

Definition 3.2.20. Write $\mathbf{a}^0 := \mathbf{a}$. Suppose we already have $\mathbf{a}^n = \mathbf{a}^n_p\mathbf{a}^n_s$; we let $\mathbf{a}^{n+1} = \mathbf{a}^n_p\boldsymbol{\lambda}^{\mathbf{a}^n_p}(\mathbf{a}^n_s)$.

Lemma 3.2.21. *$\mathbf{a}^n$ is a polynomial image of the tuple $\mathbf{a}^{n+1}$ over C.*

Proof. This amounts to showing that $\mathbf{a}^n_s$ is a polynomial image of $\mathbf{a}^{n+1}$ over C. This is clear from the definition of the component maps. In fact $\mathbf{a}^n_s$ is even a polynomial image of $\mathbf{a}^n_p(\boldsymbol{\lambda}^{a^n_p}a^n_s)^p$ over C! □

Proposition 3.2.22. *$\Lambda(E/C(\mathbf{a}^n)) = C(\mathbf{a}^{n+1})$.*

Proof. This is just re-writing Proposition 3.2.19 in our new notation. □

Proposition 3.2.23. $\boldsymbol{\Lambda}(E/C(\mathbf{a})) = \bigcup_{n<\omega} C(\mathbf{a}^n) = C(\mathbf{a}^n | n < \omega)$.

Proof. We combine Definition 3.2.9 and Proposition 3.2.22. □

The usefulness of the sequence $(\mathbf{a}^n)_{n<\omega}$ suggests the following definition.

Definition 3.2.24. Let $\mathbf{a}^\infty := \bigcup_{n<\omega} \mathbf{a}^n$ and let $\mathbf{a}_p^\infty := \bigcup_{n<\omega} \mathbf{a}_p^n$.

Note that $\mathbf{a}_p^\infty$ really is a subtuple of $\mathbf{a}^\infty$.

Proposition 3.2.25. *The tuple* $\mathbf{a}_p^\infty$ *is a p-base for* $\boldsymbol{\Lambda}(E/C(\mathbf{a})) = C(\mathbf{a}^\infty)$.

This justifies the notation which suggests that $\mathbf{a}_p^\infty$ is part of a splitting of $\mathbf{a}^\infty$ over C. In fact this can be made more precise: let $\mathbf{a}_s^\infty := \mathbf{a}^\infty \setminus \mathbf{a}_p^\infty$. Then $\mathbf{a}_s^\infty$ is separable over $C(\mathbf{a}_p^\infty)$. Note however that $\mathbf{a}_s^\infty$ may be infinite, and it may be transcendental over $C(\mathbf{a}_p^\infty)$. Thus $\mathbf{a}_s^\infty$ may not be separably generated over $C(\mathbf{a}_p^\infty)$.

Proof. Since $\mathbf{a}_p^\infty$ is p-independent in E, thus it is p-independent in $C(\mathbf{a}^\infty)$. Let $n < \omega$. By Definition 3.2.6 and Proposition 3.2.22 and since $\mathbf{a}_p^n \in pI(C(\mathbf{a}^n); E)$, we have that $\boldsymbol{\lambda}^{\mathbf{a}_p^n}(C(\mathbf{a}^n)) \subseteq \Lambda(E/C(\mathbf{a}^n)) = C(\mathbf{a}^{n+1})$. Since $\boldsymbol{\lambda}^{\mathbf{a}_p^n}$ is a tuple of *relative* component maps, we have that

$$\begin{aligned} C(\mathbf{a}^n) &\subseteq C(C(\mathbf{a}^{n+1}))^p(\mathbf{a}_p^n) \\ &\subseteq C(\mathbf{a}^{n+1})^p(\mathbf{a}_p^n) \\ &\subseteq C(\mathbf{a}^\infty)^p(\mathbf{a}_p^n) \\ &\subseteq C(\mathbf{a}^\infty)^p(\mathbf{a}_p^\infty). \end{aligned}$$

Thus $C(\mathbf{a}^\infty) = C(\mathbf{a}^\infty)^p(\mathbf{a}_p^\infty)$, so $\mathbf{a}_p^\infty$ p-spans $C(\mathbf{a}^\infty)$ over C. □

3.2.7 Finitely generated extensions

Let $\mathbf{a} \subseteq E$ be given as before. Finally for this section, we are interested in the Λ-closure of $C(\mathbf{a})$ in E under the assumption that E is finitely generated over $C(\mathbf{a})$.

A finite tuple $\mathbf{b} \subseteq E$ is separable over some finite portion of the construction of the Λ-closure of $C(\mathbf{a})$ in E. This will be the case relevant to most of our applications.

Proposition 3.2.26. *Let* $\mathbf{b}$ *be a finite tuple from* E. *Then there exists* $n < \omega$ *such that* $\mathbf{b}$ *is separable over* $C(\mathbf{a}^n)$.

Proof. Denote $\boldsymbol{\Lambda} := \boldsymbol{\Lambda}(E/C(\mathbf{a}))$. Thus $E/\boldsymbol{\Lambda}$ is separable. By Proposition 3.2.23, we already know that $\boldsymbol{\Lambda} = \bigcup_{n<\omega} C(\mathbf{a}^n)$. Since $\mathbf{b}$ is finite, we may find a separating partition $\mathbf{b} = \mathbf{b}_i\mathbf{b}_s$ over $\boldsymbol{\Lambda}$. Thus $\mathbf{b}_i$ is algebraically independent over $C(\mathbf{a}^n)$, for each $n < \omega$.

For each $b \in \mathbf{b}_s$, we let $m_b := \min(b/\boldsymbol{\Lambda}(\mathbf{b}_i))$ be the minimal polynomial of b over $\boldsymbol{\Lambda}(\mathbf{b}_i)$. Each m_b could be thought of as a rational function in $x\mathbf{b}_i$ (i.e. the tuple of variables corresponding to the concatenation of the single variable x with the tuple $\mathbf{b}_i$ of 'variables'). Let $\mathbf{c}_b$ be the tuple of coefficients of m_b (when thought of in this way) which lie in $\boldsymbol{\Lambda}$. We observe that $(\mathbf{c}_b | b \in \mathbf{b}_s)$ is a finitely generated subfield of $\boldsymbol{\Lambda}$. Thus there exists $n < \omega$ such that $(\mathbf{c}_b | b \in \mathbf{b}_s) \subseteq C(\mathbf{a}^n)$. Thus, for each $b \in \mathbf{b}_s$, we have that $m_b \in C(\mathbf{a}^n\mathbf{b}_i)[x]$. Each polynomial m_b is still irreducible and separable in $C(\mathbf{a}^n\mathbf{b}_i)[x]$; so b is separably algebraic over $C(\mathbf{a}^n\mathbf{b}_s)$. Therefore $\mathbf{b} = \mathbf{b}_i\mathbf{b}_s$ is a separating partition over $C(\mathbf{a}^n)$ and so $\mathbf{b}$ is separable over $C(\mathbf{a}^n)$. □

Proposition 3.2.27. *Suppose that E is finitely generated over $C(\mathbf{a})$. Then the recursive construction of $\boldsymbol{\Lambda}(E/C(\mathbf{a}))$ stops after finitely many steps.*

As before let $\boldsymbol{\Lambda} := \boldsymbol{\Lambda}(E/C(\mathbf{a})) = \bigcup_{n<\omega} C(\mathbf{a}^n)$. We must show that this is a finite union, i.e. that the chain of fields $C(\mathbf{a}) = C(\mathbf{a}^0) \subseteq ... \subseteq C(\mathbf{a}^n) \subseteq ...$ stabilises at some point. Note that we are **not** aiming to find $n < \omega$ such that $\mathbf{a}^n = \mathbf{a}^{n+1}$ (in fact this can only happen if $\mathbf{a}^n$ is the union of a tuple p-independent in E over C and elements of $\mathbb{F}_p$).

Proof. Let $\mathbf{b}$ be a finite tuple such that $E = C(\mathbf{ab})$. Using Proposition 3.2.26, find $n < \omega$ such that $\mathbf{b}$ is separable over $C(\mathbf{a}^n)$; then already E is separable over $C(\mathbf{a}^n)$, i.e. $C(\mathbf{a}^n)$ is Λ-closed in E. Since $\boldsymbol{\Lambda}$ is the Λ-closure of $C(\mathbf{a})$ in E, it must be the case that $C(\mathbf{a}^n) = \boldsymbol{\Lambda}$, as required. □

Chapter 4

Λ-closure in application

Our main goal is to understand the local behaviour of existentially definable sets in t-henselian fields. In the first part of chapter 6, we will see that this local behaviour can be readily understood when a certain associated field extension is separable. In this chapter we show how the technique of Λ-closure (developed in chapter 3) provides the machinery to understand the local behaviour of existentially definable sets even when the associated extension is inseparable.

The technique of Λ-closure is used to turn a field extension into a separable one, roughly speaking. This can also be interpreted geometrically because existentially definable sets are just projections of algebraic sets. We refer to the pair of an algebraic set with a projection map simply as a *projection*. Given one projection, we find another projection (which we call a *translation* of the original) of which the associated extension is separable.

Note that we do not need to be in the context of t-henselian fields for this; only for the later application (in chapter 6) of the Implicit Function Theorem. Let F/C be any field extension. We work within F and view C as a field of parameters.

4.1 A local understanding of projections

In this section we present an explanation of how the methods of this chapter will be applied in chapter 6 to give a local understanding of existentially definable sets despite inseparability of the associated extension.

Let $X \subseteq {}^{\mathbf{xy}}F$ be an algebraic set defined over C and let $\mathbf{ab} \in X$ be any point of it. Let $\mathrm{pr}_{\mathbf{x}}$:

${}^{\mathbf{xy}}F \longrightarrow {}^{\mathbf{x}}F$ be the *projection map* onto the $\mathbf{x}$-coordinates.

We would like to understand $\mathrm{pr}_{\mathbf{x}}(X)$ *locally* around $\mathbf{a}$.

$$\begin{array}{c} X \\ \Big\downarrow \mathrm{pr}_{\mathbf{x}} \\ \mathrm{pr}_{\mathbf{x}}(X) \end{array}$$

To understand $\mathrm{pr}_{\mathbf{x}}(X)$ locally means that we seek an open set $U \subseteq {}^{\mathbf{x}}F$ around $\mathbf{a}$ such that we may 'nicely' describe the set $\mathrm{pr}_{\mathbf{x}}(X) \cap U$. However, in this work we aim at a slightly weaker goal: to understand the projection itself locally around $\mathbf{ab}$, i.e. we seek an open set $V \subseteq {}^{\mathbf{xy}}F$ around $\mathbf{ab}$ such that we may 'nicely' describe $\mathrm{pr}_{\mathbf{x}}(X \cap V)$. Of course $\mathrm{pr}_{\mathbf{x}}(X \cap V) \subseteq \mathrm{pr}_{\mathbf{x}}(X)$. To describe 'nicely' means to have some other description of the set which is useful and does not simply amount to the original existential definition.

So far we have not made clear in what topology on ${}^{\mathbf{xy}}F$ should V be an open set: in later chapters we will make the additional assumption that F is t-henselian and it is in the t-henselian topology that we look for V to be an open set. The t-henselian topology is certainly a field topology, thus polynomials and (well-defined) rational functions are continuous. In fact, for much of this chapter, we will be concerened only with the Zariski topology on ${}^{\mathbf{xy}}F$.

It will be much easier for us to come to a precise local understanding of this projection if the algebraic set X is irreducible in a certain sense. To make this sense precise, we recall the notion of the locus of a point.

Definition 4.1.1. Let $\mathbf{a} \in {}^{\mathbf{x}}F$. We let $I(\mathbf{a}/C) := \{f \in C[\mathbf{x}] | f(\mathbf{a}) = 0\}$ be the ideal of polynomials over C which vanish at $\mathbf{a}$. We define the *locus* of $\mathbf{a}$ over C to be

$$\begin{aligned} \mathrm{locus}(\mathbf{a}/C) &:= Z(I(\mathbf{a}/C)) \\ &= \{\mathbf{b} \in {}^{\mathbf{x}}F \mid \forall f \in I(\mathbf{a}/C)\ f(\mathbf{b}) = 0\}, \end{aligned}$$

i.e. the zero set of the ideal of polynomials over C which vanish at $\mathbf{a}$.

It is in this sense in which we require X to be irreducible. With this simplification, our aim is now

to find a local understanding of the following projection around $\mathbf{ab}$.

$$\begin{array}{c} \mathrm{locus}(\mathbf{ab}/C) \\ \Big\downarrow \mathrm{pr}_{\mathbf{x}} \\ \mathrm{pr}_{\mathbf{x}}(\mathrm{locus}(\mathbf{ab}/C)) \end{array}$$

The *associated extension* of this projection is the field extension $C(\mathbf{ab})/C(\mathbf{a})$. As mentioned above, if this extension is separable then we will obtain a local understanding fairly straightforwardly in the first part of chapter 6. If it is not separable we will have to find a *translation*, which is roughly defined to be another projection satisfying two requirements:

1. the associated extension of the translation is separable and
2. from a local understanding of the translated projection we must be able to recover a useful amount of information about the local behaviour of the original projection.

This will be made precise in subsection 4.1.3. We use the technique of Λ-closure to find the translation.

4.1.1 Irreducible algebraic sets

We briefly recall some simple terminology from the geometry of affine algebraic sets in order to formalise the above motivation and also to check that loci really are irreducible. We say that $X \subseteq {}^{\mathbf{x}}F$ is an *algebraic set* defined over C if there exists a subset $I \subseteq C[\mathbf{x}]$ such that $X = Z(I)$ is the zero-set of I. Obviously, we may choose I to be an ideal of $C[\mathbf{x}]$. By Noetherianity of $C[\mathbf{x}]$, any ideal I will be finitely generated.

Recall that the Zariski topology (over C) on an algebraic set $X \subseteq {}^{\mathbf{x}}F$ is given by defining algebraic subsets of X, defined over C, to be closed. A topological space X is *irreducible* if it is not the union of two proper closed subsets. An algebraic set X is *irreducible* over C if it is irreducible as a topological space with the Zariski topology over C. Note that this is **not** the same as X being the F-points of a geometrically irreducible algebraic set.

Lemma 4.1.2. *If X is an algebraic set such that $\mathbf{a} \in X \subseteq \mathrm{locus}(\mathbf{a}/C)$, then $X = \mathrm{locus}(\mathbf{a}/C)$.*

Proof. Obvious from the definition of locus. □

So the locus is the smallest algebraic set, defined over C, which contains $\mathbf{a}$.

Lemma 4.1.3. *A locus is irreducible and an irreducible set is the locus of some point* $\mathbf{a}$ *which may lie in an elementary extension.*

Proof. Suppose that $X = \text{locus}(\mathbf{a}/C)$ is a locus and that $X = X_1 \sqcup X_2$ is a partition of X into disjoint closed sets. Wlog $\mathbf{a} \in X_1$. By Lemma 4.1.2, $X_1 = \text{locus}(\mathbf{a}/C)$ and so X is irreducible.

Suppose that X is irreducible and let I be an ideal in $C[\mathbf{x}]$ corresponding to X, thus $X = Z(I)$. In fact choose I to be maximal with this property. Let $\mathcal{J}$ be a finite family of ideals of $C[\mathbf{x}]$ which properly contain I. For $J \in \mathcal{J}$, let $X_J := Z(J)$, be the family of subsets of X corresponding to the ideals $J \in \mathcal{J}$. Then $X_J \subset X$ is a proper subset because J is a proper superideal of I, for each $J \in \mathcal{J}$. Since X is irreducible, X cannot be equal to the union $\bigcup_{J \in \mathcal{J}} X_J$. Applying the compactness theorem, we find an elementary extension $F \preceq F^*$ such that X^* is not equal to the union of the zero-sets of all the proper superideals of I. Thus there exists $\mathbf{a} \in X$ which is not contained in any proper closed (i.e. algebraic over C) subset of X. The ideal $I(\mathbf{a}/C)$ must contain I (since $\mathbf{a} \in X$) but cannot properly contain I. Thus $I = I(\mathbf{a}/C)$ and so $X = Z(I) = Z(I(\mathbf{a}/C)) = \text{locus}(\mathbf{a}/C)$. □

Thus loci and irreducible sets are basically the same thing, after passing to an elementary extension.

4.1.2 Rational maps between loci

The translations, mentioned above, will be tuples of polynomial and rational functions. For the sake of clarity, the phrases *polynomial map* and *rational map* will refer to tuples of polynomial and rational maps, respectively. We briefly study polynomial and rational maps between Zariski-open subsets of loci.

Let $g : {}^{\mathbf{x}}F \longrightarrow {}^{\mathbf{y}}F$ be a polynomial map defined over C and let $h : {}^{\mathbf{y}}F \longrightarrow {}^{\mathbf{x}}F$ be a rational function defined over C. Note that h is not a true function with domain ${}^{\mathbf{y}}F$ despite the notation of the previous sentence: let $V \subseteq {}^{\mathbf{y}}F$ be a Zariski-open set on which h is well-defined. Then $g \circ h$ is well-defined on V and $h \circ g$ is well-defined on $U := g^{-1}(V)$ (which may be empty!). We will write the rational map h as a (coordinate-wise) quotient of polynomial maps: thus $h = h'/h''$ where h', h'' are polynomial maps and h'' is non-zero (in any coordinate) on V. In fact we may as well let this condition define V.

Let n be the greatest degree of any single variable in any coordinate of g. Then the power to which each polynomial in the tuple h'' occurs in any coordinate of $g(h'/h'')$ is greater than or equal to $-n$. Let H be the product of the polynomials in the tuple h'', thus H is not a tuple but a single polynomial.

Then $g(h'/h'')H^n$ is a polynomial map (we mean that each coordinate of $g(h'/h'')$ is multiplied by H^n) and we may write $g \circ h$ as a ratio of polynomial maps:

$$g \circ h = g(h'/h'') = g(h'/h'')H^n \ / \ H^n.$$

Therefore on V the function $h \circ g$ is the quotient of two polynomial maps where the denominator is not zero in V.

The map $h \circ g$ can be written $h'(g)/h''(g)$, which is already a quotient of two polynomial maps with the denominator not zero (in any coordinate) on $U = g^{-1}(V)$.

Proposition 4.1.4. *Let g, h, U, V be as above. Let $\mathbf{ab} \in {}^{\mathbf{xy}}F$ and suppose that $\mathbf{b} \in V$, $g(\mathbf{a}) = \mathbf{b}$, and $h(\mathbf{b}) = \mathbf{a}$. Consequently, $\mathbf{a} \in U$. Then*

$$g : U \cap \text{locus}(\mathbf{a}/C) \longrightarrow V \cap \text{locus}(\mathbf{b}/C)$$

and

$$h : V \cap \text{locus}(\mathbf{b}/C) \longrightarrow U \cap \text{locus}(\mathbf{a}/C)$$

are well-defined and mutually inverse birational maps. In particular, these maps are mutually inverse homeomorphisms with respect to any field topology (for then polynomials and (well-defined) rational functions are continuous maps).

For a polynomial $f \in C[\mathbf{x}]$, the equality $f \circ h = f(h'/h'')H^n \ / \ H^n$ (where H is defined in the same way as above, but with respect to the greatest degree of any variable in f) holds and is well-defined on the open set V.

Proof. First we claim that g may be restricted to a function $U \cap \text{locus}(\mathbf{a}/C) \longrightarrow V \cap \text{locus}(\mathbf{b}/C)$. Clearly $g(U) = g(g^{-1}(V)) \subseteq V$. Let $f \in I(\mathbf{b}/C)$. Thus $(f \circ g)(\mathbf{a}) = f(g(\mathbf{a})) = f(\mathbf{b}) = 0$ and so $f \circ g \in I(\mathbf{a}/C)$. Let $\mathbf{a}' \in \text{locus}(\mathbf{a}/C)$. Then $(f \circ g)(\mathbf{a}') = 0$ and so $f(g(\mathbf{a}')) = 0$. Therefore $g(\text{locus}(\mathbf{a}/C)) \subseteq \text{locus}(\mathbf{b}/C)$, proving the claim.

Next we show that h may be restricted to a (well-defined) rational function $V \cap \text{locus}(\mathbf{b}/C) \longrightarrow U \cap \text{locus}(\mathbf{a}/C)$. Observe that $(g \circ h)(\mathbf{b}) = \mathbf{b}$ and $(g(h'/h'')H^n)(\mathbf{b}) - \mathbf{b} \cdot H^n(\mathbf{b}) = 0$ (where again we mean that each coordinate of $\mathbf{b}$ is multiplied by the single element $H^n(\mathbf{b})$). Of course, these are equations of $\mathbf{y}$-tuples. Therefore $g(h'/h'')H^n - \mathbf{x} \cdot H^n$ is a tuple of elements of $I(\mathbf{b}/C)$. Let $\mathbf{b}' \in V \cap \text{locus}(\mathbf{b}/C)$. Then $(g(h', h'')H^n)(\mathbf{b}') - \mathbf{b}' \cdot H^n(\mathbf{b}') = 0$. Note that $H^n(\mathbf{b}') \neq 0$ since $\mathbf{b}' \in V$. Thus $(g \circ h)(\mathbf{b}') = \mathbf{b}'$.

This proves that $h(\mathbf{b}') \in g^{-1}(V) = U$. Let $f \in I(\mathbf{a}/C)$. Then $(f(h'/h'')H^n/H^n)(\mathbf{b}) = (f \circ h)(\mathbf{b}) = 0$. Thus $f(h'/h'')H^n \in I(\mathbf{b}/C)$; and so $(f(h'/h'')H^n)(\mathbf{b}') = 0$. Again since $H^n(\mathbf{b}') \neq 0$, $(f \circ h)(\mathbf{b}') = 0$ is a well-defined equality. Thus $h(\mathbf{b}') \in \text{locus}(\mathbf{a}/C)$.

In fact we've already shown that $g \circ h$ is the identity on $V \cap \text{locus}(\mathbf{b}/C)$. Let $\mathbf{a}' \in U \cap \text{locus}(\mathbf{a}/C)$. Then $g(\mathbf{a}') \in V$, so that $H^n(g(\mathbf{a}')) \neq 0$. As remarked above, $h \circ g = h'(g)/h''(g)$ is a coordinate-wise ratio of polynomials which is well-defined on U. Thus $(h' \circ g)(\mathbf{a}) = (h'' \circ g)(\mathbf{a}) \cdot \mathbf{a}$, where this time the product is a coordinate-wise multiplication between $\mathbf{x}$-tuples. Since $(h' \circ g) - (h'' \circ g) \cdot \mathbf{x} \in I(\mathbf{a}/C)$ and $(h'' \circ g)(\mathbf{a}') \neq 0$, we have that $(h \circ g)(\mathbf{a}') = \mathbf{a}'$.

Thus the maps g and h are mutually inverse rational maps (in fact g is a polynomial map) between Zariski-open subsets of algebraic sets. In any field topology (which will refine the Zariski topology), well-defined rational maps are continuous and thus g and h will be a pair of homeomorphisms. □

4.1.3 Rational and birational translation of algebraic sets

In this section, we formalise the notion of translation that was mentioned above. Fix algebraic sets $V \subseteq {}^{\mathbf{vw}}F$ and $X \subseteq {}^{\mathbf{xy}}F$. Let $\mathbf{abcd} \in {}^{\mathbf{vwxy}}F$ and suppose that $\mathbf{ab} \in V$ and $\mathbf{cd} \in X$. In fact, translations really apply to *pointed algebraic sets*, i.e. pairs comprising algebraic sets with points on them. We will be thinking about the pointed algebraic sets $(V, \mathbf{ab})$ and $(X, \mathbf{cd})$.

Definition 4.1.5. Let $f : V \longrightarrow X$ be a rational map (as above, this means a tuple of rational functions) and let $U_f \subseteq V$ be a Zariski-open set on which f is well-defined. We say that f is *projectable* if the function $\text{pr}_{\mathbf{x}} \circ f$ is a rational function of only the variables $\mathbf{v}$. Thus $\text{pr}_{\mathbf{x}} \circ f$ is constant on fibres above elements $\mathbf{a} \in {}^{\mathbf{v}}F$. If f is projectable then we define $\text{pr}(f) : \text{pr}_{\mathbf{v}}(V) \longrightarrow \text{pr}_{\mathbf{x}}(X)$ by $\mathbf{v} \longmapsto \text{pr}_{\mathbf{x}}(f(\mathbf{vw}))$, for $\mathbf{w}$ such that $\mathbf{vw} \in U_f$. Note that $\text{pr}(f)$ is a rational map.

As usual with rational maps, we have abused notation when it comes to the domain. In fact the domain on which $\text{pr}(f)$ is well-defined is the set $\text{pr}_{\mathbf{v}}(U_f)$. Note that the image of $\text{pr}(f)$ really is contained in the set $\text{pr}_{\mathbf{x}}(X)$ because it is just the projection of f and the image of f is contained in X.

Fix $\mathbf{ab} \in V$ and $\mathbf{cd} \in X$.

Definition 4.1.6. A *translation* from $(V, \mathbf{ab})$ to $(X, \mathbf{cd})$ is a rational map $f : V \longrightarrow X$ (well-defined on some open set $U_f \subseteq V$ such that $\mathbf{ab} \in U_f$) such that

1. $f(\mathbf{ab}) = \mathbf{cd}$ and

2. f is projectable.

We write $f : (V, \mathbf{ab}) \Longrightarrow (X, \mathbf{cd})$.

The translation f is illustrated by the following diagram.

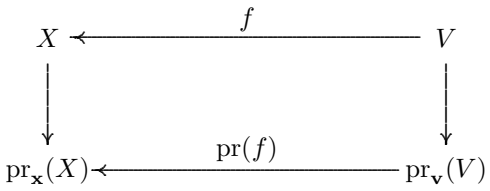

Definition 4.1.7. A *bi-translation* from $(V, \mathbf{ab})$ to $(X, \mathbf{cd})$ is a pair (f, g) of rational maps $f : V \longrightarrow X$ and $g : X \longrightarrow V$ such that

1. f and g are mutually inverse well-defined birational maps between Zariski-open sets $U_f \subseteq V$ and $U_g \subseteq X$ which are such that $\mathbf{ab} \in U_f$ and $\mathbf{cd} \in U_g$.

2. $f : (V, \mathbf{ab}) \Longrightarrow (X, \mathbf{cd})$ is a translation.

We write $(f, g) : (V, \mathbf{ab}) \Longleftrightarrow (X, \mathbf{cd})$. We say that (f, g) is *defined over* C if f and g are defined over C.

The bi-translation (f, g) is illustrated by the following diagram.

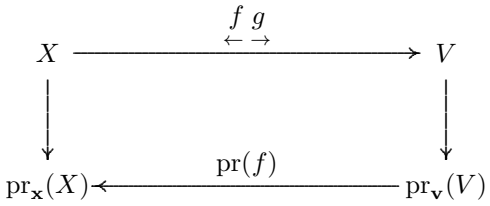

In later applications F will be a topological field, that is it will be endowed with a topology with respect to which addition, sutraction, multiplication, and division (by non-zero elements) are continuous. Such a topology will refine the Zariski topology and with respect to it well-defined rational maps will be continuous. See chapter 5 for more details; for the moment we make two simple observations.

Suppose that $\mathcal{T}$ is such a field topology on F (so that rational maps are continuous). If $(f, g) : (V, \mathbf{ab}) \Longleftrightarrow (X, \mathbf{cd})$, then f and g are mutually inverse homeomorphisms between U_f and U_g when these sets are endowed with the subspace topology inherited from T. Let $U^f \subseteq U_f$ and $U^g \subseteq U_g$ be two open sets in this topology between which f and g are homeomorphisms.

Proposition 4.1.8. $\mathrm{pr}(f)\mathrm{pr}_{\mathbf{v}}(U^f) \subseteq \mathrm{pr}_{\mathbf{x}}(U^g)$.

Proof. Let $\mathbf{a}'\mathbf{b}' \in U^f$. By definition of $\mathrm{pr}(f)$, $\mathrm{pr}(f)(\mathbf{a}') = (\mathrm{pr}_{\mathbf{x}} \circ f)(\mathbf{a}'\mathbf{b}')$ which is well-defined since f is projectable. If we write $\mathbf{c}'\mathbf{d}' := f(\mathbf{a}'\mathbf{b}')$ then $\mathbf{c}'\mathbf{d}' \in U^g$ since $f : U^f \longrightarrow U^g$. Then $\mathrm{pr}(f)(\mathbf{a}') = \mathrm{pr}_{\mathbf{x}}(f(\mathbf{a}'\mathbf{b}')) = \mathrm{pr}_{\mathbf{x}}(\mathbf{c}'\mathbf{d}') \in \mathrm{pr}_{\mathbf{x}}(U^g)$. □

Proposition 4.1.9. $\mathrm{pr}(f)\mathrm{pr}_{\mathbf{v}}(U^f) = \mathrm{pr}_{\mathbf{x}}(U^g)$.

Proof. Let $\mathbf{c}'\mathbf{d}' \in U^g$ and let $\mathbf{a}'\mathbf{b}' := g(\mathbf{c}'\mathbf{d}') \in U^f$. Then $\mathbf{c}'\mathbf{d}' = f(\mathbf{a}'\mathbf{b}')$ (as f and g are mutually inverse between U^f and U^g) and so

$$\begin{aligned} \mathbf{c}' &= \mathrm{pr}_{\mathbf{x}}(\mathbf{c}'\mathbf{d}') \\ &= \mathrm{pr}_{\mathbf{x}}(f(\mathbf{a}'\mathbf{b}')) \\ &= (\mathrm{pr}_{\mathbf{x}} \circ f)(\mathbf{a}'\mathbf{b}') \\ &= \mathrm{pr}(f)(\mathbf{a}'\mathbf{b}') \\ &\in \mathrm{pr}(f)(U^f). \end{aligned}$$

Thus $\mathrm{pr}(f)\mathrm{pr}_{\mathbf{v}}(U^f) \supseteq \mathrm{pr}_{\mathbf{x}}(U^g)$. □

The idea is that $\mathrm{pr}(f)$ translates the set $\mathrm{pr}_{\mathbf{v}}(U^f)$ into the set $\mathrm{pr}_{\mathbf{x}}(U^g)$, so that the map $\mathrm{pr}(f)$ takes a local understanding of $\mathrm{pr}_{\mathbf{v}}(V)$ and gives us a local understanding of $\mathrm{pr}_{\mathbf{x}}(X)$.

4.2 Λ-alterations

We now put the framework of translations and bi-translations into action. Using the technique of Λ-closure, we construct one translation (called *global Λ-alteration*) and one bitranslation (called *local Λ-alteration*). We just say *Λ-alteration* when referring to both of these translations. Our aim is to end up with the picture:

$$\begin{array}{ccc} \mathrm{locus}(\mathbf{ab}/C) & \xrightarrow{\substack{\Sigma\ \Lambda \\ \leftarrow\ \rightarrow}} & \mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C) \\ \downarrow & & \downarrow \mathrm{sep} \\ \mathrm{locus}(\mathbf{a}/C) & \xleftarrow{\Sigma} & \mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/C) \end{array}$$

We will also indicate why the local method gives us the best local understanding of the projection, but is less amenable to use in possible future model completeness or quantifier-elimination arguments.

On the other hand, the global method gives us only an incomplete local understanding but is simpler and perhaps more useful.

4.2.1 Global and local Λ-alteration

Let $F/E/D/C$ be a tower of fields and suppose that C is Λ-closed in F, i.e. F/C is separable. The subfield C should be thought of as a field of parameters. Using Λ-closure we can transform D so that E is separable over it. In fact there are two ways of doing this: either we take the Λ-closure of D in E or we take it in F.

Definition 4.2.1. The phrase *global Λ-alteration* refers to defining $\boldsymbol{\Lambda}D := \boldsymbol{\Lambda}(F/D)$ and $\boldsymbol{\Lambda}E := E \cdot \boldsymbol{\Lambda}(F/D)$ (i.e. the compositum of E with $\boldsymbol{\Lambda}(F/D)$). The phrase *local Λ-alteration* refers to defining $\boldsymbol{\Lambda}D := \boldsymbol{\Lambda}(E/D)$ and $\boldsymbol{\Lambda}E := E$. These phrases will also be associated to particular maps and pairs of maps later on.

Both of these alterations are illustrated by the diagram below.

Global Λ-alteration

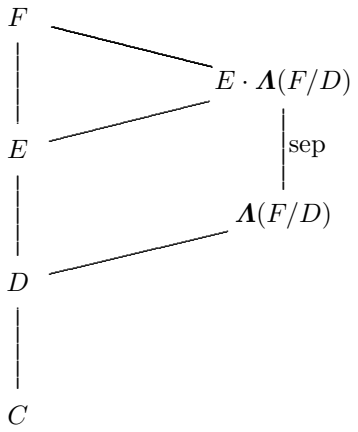

Local Λ-alteration

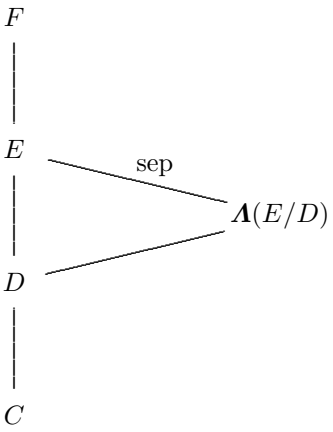

We now focus on defining the Λ-alteration of pairs of tuples rather than simply of field extensions. Let $\mathbf{ab} \in {}^{\mathbf{xy}}F$ and let $D = C(\mathbf{a})$ and $E = C(\mathbf{ab})$. In subsection 3.2.7 we saw how to simplify the construction of $\boldsymbol{\Lambda}(E/C(\mathbf{a}))$ and $\boldsymbol{\Lambda}(F/C(\mathbf{a}))$ to just involve the tuple $\mathbf{a}$.

Definition 4.2.2. In both the global and the local Λ-alterations of $F/E/D/C$, the construction of $\boldsymbol{\Lambda}C(\mathbf{a})$ is complete after only finitely many steps. Thus there exists $n \in \mathbb{N} \cup \{0\}$ such that $\boldsymbol{\Lambda}C(\mathbf{a}) = C(\mathbf{a}^n)$. We write $\boldsymbol{\lambda}\mathbf{a} := \mathbf{a}^n$, where n is as just described. We call $\boldsymbol{\lambda}\mathbf{a}$ the *λ-splitting* of $\mathbf{a}$. We will say that $\boldsymbol{\lambda}\mathbf{a}$ is the *global* or *local* splitting of $\mathbf{a}$ depending on whether $\boldsymbol{\lambda}\mathbf{a}$ was defined with respect to the

global or local constructions of $\boldsymbol{\Lambda}D$ and $\boldsymbol{\Lambda}E$ in Definition 4.2.1. We will usually let $\mathbf{v}$ be a tuple of variables corresponding to the tuple $\boldsymbol{\lambda}\mathbf{a}$.

By Lemma 3.2.21, $\mathbf{a}$ is a polynomial image of $\boldsymbol{\lambda}\mathbf{a}$. Let $\Sigma \in C[\mathbf{v}]$ be a polynomial map such that $\Sigma(\boldsymbol{\lambda}\mathbf{a}) = \mathbf{a}$. Let $\boldsymbol{\Sigma}$ be the polynomial map

$$\boldsymbol{\Sigma} := \Sigma \times \mathrm{id}_{\mathbf{y}} : \mathbf{xy} \longmapsto (\Sigma(\mathbf{x}), \mathbf{y})$$

so that $\boldsymbol{\Sigma}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}) = \mathbf{ab}$. We refer to both Σ and $\boldsymbol{\Sigma}$ as *Σ-maps*. Note that Σ-maps will be different in the local and global settings.

In the situation of the local Λ-alteration we have that $E = C(\mathbf{ab})$. Thus $\boldsymbol{\lambda}\mathbf{a} \in C(\mathbf{ab})$. Let $\Lambda \in C(\mathbf{xy})$ be the rational map which is well-defined at $\mathbf{ab}$ and is such that $\boldsymbol{\Lambda}(\mathbf{ab}) = \boldsymbol{\lambda}\mathbf{a}$. Let $\boldsymbol{\Lambda}$ be the rational map

$$\boldsymbol{\Lambda} := \Lambda \times \mathrm{id}_{\mathbf{y}} : \mathbf{vy} \longmapsto (\Lambda(\mathbf{vy}), \mathbf{y})$$

so that $\boldsymbol{\Lambda}(\mathbf{ab}) = (\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})$. We refer to both Λ and $\boldsymbol{\Lambda}$ as *Λ-maps*. Note that Λ-maps are not defined in the global setting.

The polynomial map $\boldsymbol{\Sigma}$ (defined with respect to the global splitting of $\mathbf{a}$) is called the *global Λ-alteration* of $\mathbf{b}/\mathbf{a}$ (with respect to F/C). The pair $(\boldsymbol{\Sigma}, \boldsymbol{\Lambda})$ (defined with respect to the local splitting of $\mathbf{a}$) is called the *local Λ-alteration* of $\mathbf{b}/\mathbf{a}$ (with respect to C).

Note that the field F has nothing to do with the local Λ-alteration of $\mathbf{b}/\mathbf{a}$.

Remark 4.2.3. In fact we may define $\boldsymbol{\Lambda}$ (and Λ) in even the global setting but it would not be a rational map and does not restrict to a birational map between open subsets of loci.

4.2.2 Λ-alterations give translations

Let $\mathbf{ab} \in {}^{\mathbf{xy}}F$.

Proposition 4.2.4. *Let $\boldsymbol{\Sigma}$ be the global Λ-alteration of $\mathbf{b}/\mathbf{a}$ with respect to F/C. Then*

$$\boldsymbol{\Sigma} : (\mathrm{locus}(\lambda\mathbf{a}, \mathbf{b}/C), (\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})) \Longrightarrow (\mathrm{locus}(\mathbf{ab}/C), \mathbf{ab})$$

is a translation defined over C.

Proof. Let $\mathbf{l} \in \mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/C)$ and let $f \in I(\mathbf{a}/C)$. Since $\boldsymbol{\Sigma}$ is a polynomial map, $f \circ \boldsymbol{\Sigma}$ is a polynomial.

In fact $(f \circ \boldsymbol{\Sigma})(\boldsymbol{\lambda}\mathbf{a}) = f(\boldsymbol{\Sigma}(\boldsymbol{\lambda}\mathbf{a})) = f(\mathbf{a}) = 0$, thus $(f \circ \boldsymbol{\Sigma})(\mathbf{l}) = 0$. Therefore $\boldsymbol{\Sigma}$ restricts to a map $\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/C) \longrightarrow \mathrm{locus}(\mathbf{a}/C)$. Also observe that $\mathrm{pr}_{\mathbf{x}} \circ \boldsymbol{\Sigma} = \Sigma$ is a polynomial map in the variables $\mathbf{v}$. Thus $\boldsymbol{\Sigma}$ is a translation. □

Proposition 4.2.5. *Let $(\boldsymbol{\Sigma}, \boldsymbol{\Lambda})$ be the local Λ-alteration of $\mathbf{b}/\mathbf{a}$ with respect to C. Then*

$$(\boldsymbol{\Sigma}, \boldsymbol{\Lambda}) : (\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C), (\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})) \Longleftrightarrow (\mathrm{locus}(\mathbf{ab}/C), \mathbf{ab})$$

is a bi-translation defined over C.

Proof. Since $\boldsymbol{\Sigma}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}) = \mathbf{ab}$ and $\boldsymbol{\Lambda}(\mathbf{ab}) = (\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})$ and by Proposition 4.1.4, there exist Zariski open subsets $U \subseteq \mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C)$ and $V \subseteq \mathrm{locus}(\mathbf{ab}/C)$ such that $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}$ are mutually inverse birational maps between U and V. Note that $\mathrm{pr}_{\mathbf{x}} \circ \boldsymbol{\Sigma} = \Sigma$ is a polynomial map in the variables $\mathbf{v}$ over C; thus $\boldsymbol{\Sigma}$ is projectable. Therefore $(\boldsymbol{\Sigma}, \boldsymbol{\Lambda})$ is a bi-translation. □

We return to the situation where F is endowed with a topology refining the Zariski-topology with respect to which rational functions are continuous. Let $\mathcal{T}$ and $\mathcal{T}_{\boldsymbol{\lambda}}$ denote the topologies induced on the subspaces $\mathrm{locus}(\mathbf{ab}/C)$ and $\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C)$, respectively.

Proposition 4.2.6. *Let $\boldsymbol{\Sigma}$ be the global Λ-alteration of $\mathbf{b}/\mathbf{a}$ with respect to F/C. Let $U \in \mathcal{T}_{\boldsymbol{\lambda}}$ be a $\mathcal{T}_{\boldsymbol{\lambda}}$-open set. Then*

$$\Sigma(\mathrm{pr}_{\mathbf{v}}(U)) \subseteq \mathrm{pr}_{\mathbf{x}}(\mathrm{locus}(\mathbf{ab}/C)).$$

Note that if $\boldsymbol{\lambda}\mathbf{a} \in \mathrm{pr}_{\mathbf{v}}(U)$ then $\mathbf{a} \in \Sigma(\mathrm{pr}_{\mathbf{v}}(U))$.

Proof. Since $\boldsymbol{\Sigma}$ is a polynomial map, we may as well let $U_{\boldsymbol{\Sigma}}$ (i.e. the Zariski-open set containing $(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})$ on which $\boldsymbol{\Sigma}$ is defined) be the whole of $\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C)$. Since $\boldsymbol{\Sigma} : (\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C), (\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})) \Longrightarrow (\mathrm{locus}(\mathbf{ab}/C), \mathbf{ab})$, the image of any $\mathcal{T}_{\boldsymbol{\lambda}}$-open set $U \subseteq U_{\boldsymbol{\Sigma}}$ under the map $\boldsymbol{\Sigma}$ is contained in $\mathrm{locus}(\mathbf{ab}/C)$. Thus the image of $\mathrm{pr}_{\mathbf{v}}(U)$ under $\Sigma = \mathrm{pr}(\boldsymbol{\Sigma})$ is contained in $\mathrm{pr}_{\mathbf{x}}(\mathrm{locus}(\mathbf{ab}/C))$. □

Proposition 4.2.7. *Let $(\boldsymbol{\Sigma}, \boldsymbol{\Lambda})$ be the local Λ-alteration of $\mathbf{b}/\mathbf{a}$ with respect to C. Let $U^{\boldsymbol{\Sigma}} \in \mathcal{T}_{\boldsymbol{\lambda}}$ be a $\mathcal{T}_{\boldsymbol{\lambda}}$-open set such that $U^{\boldsymbol{\Sigma}} \subseteq U_{\boldsymbol{\Sigma}}$. There exists a $\mathcal{T}$-open set $U^{\boldsymbol{\Lambda}} \in \mathcal{T}$ such that $U^{\boldsymbol{\Lambda}} \subseteq U_{\boldsymbol{\Lambda}}$ and*

$$\Sigma(\mathrm{pr}_{\mathbf{v}}(U^{\boldsymbol{\Sigma}})) = \mathrm{pr}_{\mathbf{x}}(U^{\boldsymbol{\Lambda}}) \subseteq \mathrm{pr}_{\mathbf{x}}(\mathrm{locus}(\mathbf{ab}/C)).$$

Note that $(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}) \in U^{\boldsymbol{\Sigma}}$ if and only if $\mathbf{ab} \in U^{\boldsymbol{\Lambda}}$.

Proof. Since $(\boldsymbol{\Sigma}, \boldsymbol{\Lambda}) : (\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C), (\boldsymbol{\lambda}\mathbf{a}, \mathbf{b})) \iff (\mathrm{locus}(\mathbf{ab}/C), \mathbf{ab})$, we have that $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}$ are mutually inverse well-defined rational maps between Zariski-open sets $U_{\boldsymbol{\Sigma}} \subseteq \mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/C)$ and $U_{\boldsymbol{\Lambda}} \subseteq \mathrm{locus}(\mathbf{ab}/C)$ which are such that $(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}) \in U_{\boldsymbol{\Sigma}}$ and $(\mathbf{ab}/C) \in U_{\boldsymbol{\Lambda}}$. Since well-defined rational maps are $(\mathcal{T}_{\boldsymbol{\lambda}}, \mathcal{T})$-continuous by assumption, $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}$ are mutually inverse homeomorphisms between $U_{\boldsymbol{\Sigma}}$ and $U_{\boldsymbol{\Lambda}}$ with respect to $\mathcal{T}_{\boldsymbol{\lambda}}$ and $\mathcal{T}$. Thus, for each $\mathcal{T}_{\boldsymbol{\lambda}}$-open set $U^{\boldsymbol{\Sigma}} \subseteq U_{\boldsymbol{\Sigma}}$, there exists a unique $\mathcal{T}$-open set $U^{\boldsymbol{\Lambda}} \subseteq U_{\boldsymbol{\Lambda}}$ such that $\boldsymbol{\Sigma}$ and $\boldsymbol{\Lambda}$ restrict to homeomorphisms between them. Applying Proposition 4.1.9, we have that $\Sigma(\mathrm{pr}_{\mathbf{v}}(U^{\boldsymbol{\Sigma}})) = \mathrm{pr}(\boldsymbol{\Sigma})(\mathrm{pr}_{\mathbf{v}}(U^{\boldsymbol{\Sigma}})) = \mathrm{pr}_{\mathbf{x}}(U_{\boldsymbol{\Lambda}})$. □

Part II

$\exists$-definability in t-henselian fields

Chapter 5

Topological fields

5.1 Prestel and Ziegler's paper

The entirety of Part II will be set in the context of t-henselian fields. In this section we recall the theory of topological fields and particularly t-henselian fields from the 1978 paper of Prestel and Ziegler, [25]. Everything in this section can be found in that paper or in the book [8].

5.1.1 Topological fields, filtered fields, and V-topological fields

Definition 5.1.1. Let F be a field and let $\mathcal{T}$ be a topology on F. We say that $(F, \mathcal{T})$ is a *topological field* if $\mathcal{T}$ is Hausdorff and not discrete and addition, subtraction, multiplication, and (non-zero) division are continuous.

Let τ be a filter base of $(\mathcal{P}(F), \subseteq)$ such that $0 \in \alpha$, for each $\alpha \in \tau$, and that $\{0\} \notin \tau$. The set

$$\mathcal{T}_\tau := \{A \in \mathcal{P}(F) | \forall a \in A \; \exists \alpha \in \tau : a + \alpha \subseteq A\}$$

is a topology on F; we call it the topology *generated by* τ. We say that (F, τ) is a *filtered field* if $(F, \mathcal{T}_\tau)$ is a topological field.

We consider a two sorted language $\mathcal{L}_{\text{filt}}$ with the ring language on the first sort and a relation '$\in$' between the sorts. A formula is said to be in *negation normal form* if the symbols for implication and 'if and only if' do not occur and negations only appear directly in front of atomic formulas. Every formula is logically equivalent to a formula in negation normal form. To avoid overly cumbersome

notation we will use roman letters for (field) variables of the first sort and greek letters for variables of the second sort. An $\mathcal{L}_{\text{filt}}$-formula is said to be *local* if its negation normal form has the property that, for any second-sort variable α, within the scope of $\forall\alpha$ the atomic formulas containing α occur positively and within the scope of $\exists\alpha$ the atomic formulas containing α occur negatively. A *local sentence* is a sentence which is a local formula. A class of structures axiomatised by a set of local sentences is called a *local class*. Two structures (F, τ) and (F', τ') are said to be *locally equivalent* if they model the same set of local sentences.

Proposition 5.1.2. *(***Theorem 1.1(a)**, [25]) Let τ, τ' be two bases of the same filter. Then (F, τ) and (F, τ') are locally equivalent.

Since formulas in just the ring language with variables from only the field sort are certainly local formulas; if (F, τ) and (F', τ') are locally equivalent then F and F' are elementarily equivalent in the language of rings. In just the same way, any $\mathcal{L}_{\text{ring}}$-theory is trivially a local $\mathcal{L}_{\text{filt}}$-theory.

Let $\mathbf{T}_{\text{fields}}$ be any axiomatisation of the theory of fields in the language of rings. The local $\mathcal{L}_{\text{filt}}$-theory

$$\mathbf{T}_0 := \left\{ \begin{array}{l} \forall\alpha\ \forall\beta\ \exists\gamma\ \forall x\ (x \in \gamma \longrightarrow (x \in \alpha \wedge x \in \beta)), \\ \forall\alpha\ \exists x\ (\neg x \doteq 0 \wedge x \in \alpha), \\ \forall x\ (x \doteq 0 \longleftrightarrow \forall\beta\ x \in \beta) \end{array} \right\}$$

is an axiomatisation of filters around 0 which generate non-discrete Hausdorff topologies. The class of filtered fields is axiomatised in this language by the theory $\mathbf{T}_{\text{fields}} \cup \mathbf{T}_0$ together with the following sentences:

1. $\forall\alpha\ \exists\beta\ \beta - \beta \subseteq \alpha$,
2. $\forall\alpha\ \forall a, b\ \exists\beta\ (a + \beta)(b + \beta) \subseteq ab + \alpha$, and
3. $\forall\alpha\ \exists\beta\ (1 + \beta)^{-1} \subseteq 1 + \alpha$.

We can re-write these as local $\mathcal{L}_{\text{filt}}$-sentences in our formal language as follows:

1. $\phi_+ := \forall\alpha\ \exists\beta\ \forall x\ \forall y\ ((x \in \beta \wedge y \in \beta) \longrightarrow x - y \in \alpha)$,
2. $\phi_\times := \forall\alpha\ \forall a\ \forall b\ \exists\beta\ \forall u, v\ ((u \in \beta \wedge v \in \beta) \longrightarrow \exists w\ (w \in \alpha \wedge (x + u)(y + v) = xy + w))$, and
3. $\phi_\div := \forall\alpha\ \exists\beta\ \forall y (y \in \beta \longrightarrow \exists x\ (x \in \alpha \wedge (1 + y)(1 + x) = 1))$.

Let $\mathbf{T}_{\text{filt}} := \mathbf{T}_{\text{field}} \cup \mathbf{T}_0 \cup \{\phi_+, \phi_\times, \phi_\div\}$. The proof that $\mathbf{T}_{\text{filt}}$ axiomatises the class of filtered fields is fairly straightforward. Roughly, ϕ_+ expresses the continuity of addition and subtraction, $\phi_\times$ the continuity of multiplication, and $\phi_\div$ the continuity of the multiplicative inverse. See Appendix B of [8] for more details. Since these are local sentences, the class of filtered fields is a local class.

Definition 5.1.3. We say that a filtered field (F, τ) is *ω-complete* if the filter is closed under countable intersections.

Thus a filtered field is ω-complete if it realises certain countable types.

Proposition 5.1.4. *(***Theorem 1.1(b)**, [25]) Every filtered field has an $\mathcal{L}_{\text{filt}}$-elementary extension which is an ω-complete filtered field.

Proof. By the usual model-theoretic construction (using a cardinal like $\aleph_1$), there exists an $\aleph_0$-saturated elementary extension. □

Any absolute value or valuation on a field F will induce a field topology on F. Let τ be any filter base of the filter of open neighbourhoods of zero in this induced topology. Then τ will satisfy the following additional property:

$$\forall \alpha \, \exists \beta \, (F \setminus \alpha) \cdot (F \setminus \alpha) \subseteq F \setminus \beta,$$

i.e. if a product of two elements lies in β, at least one of the factors must lie in α. Roughly speaking, this axiom expresses the property that the product of two large elements is large.

We can re-write this last property as a local $\mathcal{L}_{\text{filt}}$-sentence: let

$$\phi_V := \forall \alpha \, \exists \beta \, \forall a \, \forall b \, (ab \in \beta \longrightarrow (a \in \alpha \vee b \in \alpha)).$$

Definition 5.1.5. A filtered field (F, τ) is *V-topological* if (F, τ) models ϕ_V. The theory $\mathbf{T}_{\text{V}} := \mathbf{T}_{\text{filt}} \cup \{\phi_V\}$ is (an axiomatisation of) the theory of V-topological fields. The class of V-topological fields is a local class since ϕ_V is a local sentence.

A *valuation ideal* is a set $\alpha \subseteq F$ such that $\alpha + \alpha \subseteq \alpha$, $\alpha \cdot \alpha \subseteq \alpha$, $1 \notin \alpha$ and if $xy \in \alpha$ then $x \in \alpha$ or $y \in \alpha$. Any maximal ideal of a valuation ring is a valuation ideal, and conversely any valuation ideal gives a valuation ring of F (as $\alpha \cup \{x \in F | x, x^{-1} \notin \alpha\}$). Thus there is a bijection between valuation ideals and valuation rings.

Proposition 5.1.6. *If (F, τ) is an ω-complete V-topological field then τ admits a base of valuation ideals. Of course, these ideals correspond to valuations which induce the same topology as τ.*

The following important characterisation of V-topological fields is due to Kowalsky-Dürbaum and also Fleischer.

Proposition 5.1.7. *(Theorem 3.1 of [25]) A filtered field (F, τ) is V-topological if and only if there is an absolute value or a valuation on F which induces τ.*

For a proof, see Appendix B of [8]. These absolute values and valuations are always non-trivial since the trivial valuation induces the indiscrete topology, which is not Hausdorff and has already been excluded in the definition of filtered fields.

5.1.2 t-henselian fields and the Implicit Function Theorem

A filtered field (F, τ) is *henselian* if F admits a non-trivial henselian valuation which induces the topology $\mathcal{T}_\tau$.

Definition 5.1.8. We say that a filtered field (F, τ) is *t-henselian* if it is locally equivalent to a filtered field in which the topology induced by a henselian valuation.

Especially in later chapters we will say that a field F is *t-henselian* if (F, τ) is t-henselian for some τ. Similarly we will say that F is *henselian* if F admits a non-trivial henselian valuation. Therefore t-henselian fields are $\mathcal{L}_{\text{ring}}$-elementarily equivalent to henselian fields. In fact, any field which is not separably closed and $\mathcal{L}_{\text{ring}}$-elementarily equivalent to a henselian field is t-henselian because the topology will be an $\mathcal{L}_{\text{ring}}$-definable family. Note that any t-henselian filtered field is V-topological.

The following proposition characterises t-henselian filtered fields by a 'non-uniform' version of a common characterisation of henselianity. For $\alpha \in \tau$, let $\alpha[x]^n$ denote the set of polynomials in x of degree at most n with coefficients from α. We let $x^{n+1} + x^n + \alpha[x]^{n-1}$ denote the set given by addition of $x^{n+1} + x^n$ with each element of $\alpha[x]^{n-1}$.

Proposition 5.1.9. *(***Theorem 7.2(a)**, [25]) Let (F, τ) be a filtered field. Then (F, τ) is t-henselian if and only if for each $n \in \mathbb{N}$ there exists $\alpha \in \tau$ such that each $f \in x^{n+1} + x^n + \alpha[x]^{n-1}$ has a zero in F.

We saw above that in sufficiently saturated V-topological fields the topology is induced by valuations. In the case of a sufficiently saturated t-henselian field, the topology is induced by henselian valuations.

Proposition 5.1.10. *An ω-complete t-henselian field admits a base of valuation ideals which belong to henselian valuations.*

We now introduce some notation that will be commonplace in the subsequent chapters. Let $\mathbf{a} = (a_x)_{x\in\mathbf{x}} \in {}^{\mathbf{x}}F$ and let $\boldsymbol{\alpha} = (\alpha_x)_{x\in\mathbf{x}} \in {}^{\mathbf{x}}\tau$. We set

$$\mathbf{B}(\boldsymbol{\alpha};\mathbf{a}) := \prod \{\mathbf{a}(x) + \boldsymbol{\alpha}(x) \,|\, x \in \mathbf{x}\} = \{(b_x)_{x\in\mathbf{x}} \in {}^{\mathbf{x}}F \,|\, a_x - b_x \in \alpha_x, \text{ for each } x \in \mathbf{x}\}$$

to be the *ball* around $\mathbf{a}$ of *radius* $\boldsymbol{\alpha}$. In later chapters we will also use the notation $\mathbf{B}^n(\boldsymbol{\alpha};\mathbf{a}) := \mathbf{B}(\boldsymbol{\alpha};\mathbf{a}^{p^{-n}})^{p^n}$.

For filtered fields (F,τ), we define the following property in analogy to the usual Implicit Function Theorem of Real Analysis. Let D_y denote the formal derivative with respect to the variable y.

Property 5.1.11 (The Implicit Function Theorem)
Let $f \in F[\mathbf{x}y]$ and let $\mathbf{a}b \in {}^{\mathbf{x}y}F$. Suppose that $D_y f(\mathbf{a}b) \neq 0$. Then there exist $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ such that

$$Z(f) \cap \mathbf{B}(\boldsymbol{\alpha}\beta;\mathbf{a}b)$$

is the graph of a continuous algebraic function

$$\mathbf{B}(\boldsymbol{\alpha};\mathbf{a}) \longrightarrow \mathbf{B}(\beta, b).$$

Here *algebraic* means both that the function satisfies an algebraic equation (after all, the graph is the zero-set of f) and that $f(\mathbf{a}')$ is algebraic over $\mathbf{a}'$, for every $\mathbf{a}' \in \mathbf{B}(\boldsymbol{\alpha};\mathbf{a})$.

Proposition 5.1.12. *(Theorem 7.4, [25]) A filtered field (F,τ) is t-henselian if and only if it satisfies the* **IFT**.

Note that the Implicit Function Theorem is an infinite set of local $\mathcal{L}_{\text{filt}}$-sentences.

5.1.3 Non-separably closed t-henselian fields

In analogue to the well-known theorem of Schmidt, Prestel and Ziegler prove that a non-separably closed field admits at most one t-henselian topology. Furthermore, Prestel-Ziegler (in [25]) proved that (if F is t-henselian and not separably closed) there exists an $\exists$-definable (in $\mathcal{L}_{\text{ring}}$) base for the filter of neighbourhoods of zero in the t-henselian topology. A corrected definition for a base of the

neighbourhood filter of 0 was given by Prestel in [23].

Let $D := D_x$ denote the formal derivative.

Proposition 5.1.13. *(Remark 7.11, [25]; Lemma, [23]) Suppose that F is t-henselian. Let $f \in F[x]$ be a separable irreducible polynomial without a zero in F and let $a \in F$ be such that $Df(a) \neq 0$. Let $U_{f,a} := f(F)^{-1} - f(a)^{-1}$. Then $\tau := \{c \cdot U_{f,a} | c \in F^\times\}$ is a base for the unique t-henselian topology on F.*

Proof. The polynomial f is irreducible and non-linear, and so has no zeroes; i.e. $0 \notin f(F)$. Let $S := f(F)^{-1}$. By Lemma 7.5(ii) in [25], $S^{-1} = f(F)$ is bounded away from zero. By statement equivalent to the definition of V-topological, S is bounded. In fact the statement we use here is taken to be the definition of V-topological in [25]. Note that 'bounded' has a special definition in this context.

Let $b := f(a)$. Since $a \notin Z(Df)$, we may apply the Implicit Function Theorem to $\hat{f} := f(x) - y$ at the point ab: we find $\alpha\beta \in {}^{xy}\tau$ such that $B(\alpha\beta; ab) \cap Z(\hat{f})$ is the graph of a continuous function $B(\beta; b) \longrightarrow B(\alpha; a)$. Thus $B(\beta; b) \subseteq f(F)$.

By the continuity of multiplicative inverse away from zero, there exists $\gamma \in \tau$ such that $\big(B(\gamma, b^{-1})\big)^{-1} \subseteq B(\beta; b) \subseteq f(F)$. Thus $B(\gamma; b^{-1}) \subseteq f(F)^{-1}$ and $B(\gamma; 0) \subseteq f(F)^{-1} - f(a)^{-1} = U_{f,a}$.

Thus $U_{f,a}$ is a bounded neighbourhood of 0 in the t-henselian topology. Clearly, taking (non-zero) multiples of the set $U_{f,a}$ gives a base for the filter of neighbourhoods of 0. □

Suppose that $(F, \mathcal{O})$ is a henselian valued field with value group vF. Then a base for the filter of neighbourhoods of zero is given by the sets $B(n; 0) := \{a \in F \,|\, va > n\}$. Thus for each $n \in vF$, there exists $c \in F^\times$ such that $c \cdot U_{f,a} \subseteq B(n; 0)$; whence $U_{f,a} \subseteq B(n - vc; 0)$. Therefore $U_{f,a}$ is bounded in the usual valuation theoretical sense with respect to any valuation inducing the same topology.

5.2 Local equality of algebraic sets

Let (F, τ) be a filtered field. For a ball $\mathbf{B}$ and two sets A, B, we write $A \equiv_{\mathbf{B}} B$ to mean that $A \cap \mathbf{B} = B \cap \mathbf{B}$. This section contains several technical propositions which describe how loci are locally controlled by the zero-sets of minimal polynomials. More precisely, we are interested in the 'numerators' of minimal polynomials: that is, if m is the minimal polynomial of b over $C(\mathbf{a})$, then we may view m as a polynomial in b and a rational function in $\mathbf{a}$. Thus $m = m'(\mathbf{a}y)/m''(\mathbf{a})$ where $m' \in C[\mathbf{x}y]$ and $m'' \in C[\mathbf{x}]$. Roughly speaking, it is the zero-set of m' that will control the locus of $\mathbf{a}b$.

Recall that, for a polynomial $f \in C[\mathbf{x}]$, $Z(f) := \{\mathbf{a} \in {}^{\mathbf{x}}F \mid f(\mathbf{a}) = 0\}$ denotes the zero-set of f.

5.2.1 Solving equations locally

Proposition 5.2.1. *Let $\mathbf{a} \in {}^{\mathbf{x}}F$ and let $f \in C[\mathbf{x}] \setminus I(\mathbf{a}/C)$. Then there exists $\boldsymbol{\alpha} \in {}^{\mathbf{x}}\tau$ such that $Z(f) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) = \emptyset$.*

Proof. The set $Z(f)$ is closed in the Zariski topology on ${}^{\mathbf{x}}F$, thus its complement is open. Since the field topology refines the Zariski-topology, there exists $\boldsymbol{\alpha} \in {}^{\mathbf{x}}\tau$ such that $\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \subseteq {}^{\mathbf{x}}F \setminus Z(f)$. □

Recall that $\pi_{\mathbf{a}}$ denotes the map 'evaluate at $\mathbf{a}$'. Note that we must be careful where $\pi_{\mathbf{a}}$ is defined: it will not be defined on rational functions with polynomials from $I(\mathbf{a}/C)$ in the denominator.

Proposition 5.2.2. *Let $\mathbf{a} \in {}^{\mathbf{x}}F$ and let $f \in C[\mathbf{xy}]$ be such that $\pi_{\mathbf{a}} f = 0$. Then, for every $\mathbf{a}' \in$ locus$(\mathbf{a}/C)$, we have that $f(\mathbf{a}'\mathbf{y}) = 0$.*

The idea is that the vanishing of the coefficients of $\pi_{\mathbf{a}} f \in C[\mathbf{a}][\mathbf{y}]$ is encoded in the locus of $\mathbf{a}$ over C.

Proof. Since $C[\mathbf{xy}] \cong C[\mathbf{x}][\mathbf{y}]$, we think of f as a polynomial in the variables $\mathbf{y}$ over the ring $C[\mathbf{x}]$. From this viewpoint, the coefficients of f are polynomials in $\mathbf{x}$. On our assumption that the equation $\pi_{\mathbf{a}} f = 0$ holds in the polynomial ring $C[\mathbf{ay}]$, it must be the case that each of the finitely many coefficients of f vanishes when evaluated at $\mathbf{a}$. This is clearly encoded in locus$(\mathbf{a}/C)$ and the result follows. □

5.2.2 Minimal polynomials

The following proposition is part of the explanation that the zero set of any polynomial is locally equal to the zero set of the minimal polynomial. Let χ denote the map that re-writes a polynomial or rational function $f(\mathbf{ay}) \longmapsto f(\mathbf{xy})$ simply by replacing occurrences of $\mathbf{a}$ with $\mathbf{x}$. This is well-defined on all rational functions. We briefly re-write $\pi := \pi_{\mathbf{a}}$.

Proposition 5.2.3. *Let $\mathbf{a}b \in {}^{\mathbf{x}y}F$ and suppose that b is algebraic over $C(\mathbf{a})$. Let $m' \in C[\mathbf{x}y]$ and $m'' \in C[\mathbf{x}]$ be such that $m := m'/m'' = \chi\min(b/C(\mathbf{a}))$. Let $f \in I(\mathbf{a}b/C) \setminus I(\mathbf{a}/C)$. Then there exist $q' \in C[\mathbf{x}y] \setminus I(\mathbf{a}b/C)$, $q'' \in C[\mathbf{x}] \setminus I(\mathbf{a}/C)$, and $n < \omega$ such that*

1. *$\pi d_f = 0 \in C[\mathbf{a}y]$ (i.e. $\pi m^n \mid \pi f$ in $C(\mathbf{a})[y]$), where $d_f := m'^n q' - m''^n q'' f \in C[\mathbf{x}y]$.*

Furthermore:

2. *there exists* $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ *such that* $Z(q') \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) = \emptyset$*;*

3. *there exists* $\boldsymbol{\alpha}' \in {}^{\mathbf{x}}\tau$ *such that* $Z(m''^n q'') \cap \mathbf{B}(\boldsymbol{\alpha}'; \mathbf{a}) = \emptyset$*;*

4. $Y_f \cap \mathbf{B}(\boldsymbol{\alpha}''\beta; \mathbf{a}b) = \emptyset$*, where* $Y_f := Z(m''^n q'') \cup Z(q')$ *and* $\boldsymbol{\alpha}'' := \boldsymbol{\alpha} \cap \boldsymbol{\alpha}'$*;*

5. $Z(d_f, m') \setminus Z(m''^n q'') \subseteq Z(f)$*;*

6. $Z(d_f, f) \setminus Z(q') \subseteq Z(m')$*; and*

7. $Z(m') \cap Z_f = Z(f) \cap Z_f$*, where* $Z_f := Z(d_f) \setminus Y_f$*.*

Proof. 1. As $\pi m = \min(b/C(\mathbf{a}))$, $\pi m \mid \pi f$ in $C(\mathbf{a})[y]$. There exists $n < \omega$ such that $\pi m^n \mid \pi f$ but $\pi m^{n+1} \nmid \pi f$. Let $q \in C(\mathbf{x})[y]$ be such that $q = \chi_{\frac{\pi f}{\pi m^n}}$. Then $\pi(m^n q) = \pi f$. Let $q' \in C[\mathbf{x}y]$ be the numerator and $q'' \in C[\mathbf{x}]$ be the denominator of q. Thus $\pi(m'^n q') = \pi(m''^n q'' f)$, i.e. $\pi d_f = 0$.

2. Since $\pi q' = q'(\mathbf{a}b) \neq 0$, we may apply Proposition 5.2.1.

3. Since $\pi(m''^n q'') = m''^n q''(\mathbf{a}) \neq 0$, we may apply Proposition 5.2.1.

4. Immediate.

5.,6., and 7. follow simply from 4. and because F is an integral domain. □

5.2.3 Loci and minimal polynomials

We now prove how the locus of a tuple can be completely controlled by the minimal polynomials. For the rest of this section the tuple $\mathbf{y}$ of variables will be finite.

Proposition 5.2.4. *Let* $\mathbf{a}b \in {}^{\mathbf{x}y}F$ *and suppose that* b *is algebraic over* $C(\mathbf{a})$*. Let* $F \subseteq_{fin} I(\mathbf{a}b/C)$ *be such that* $F \not\subseteq I(\mathbf{a}/C)$*. There exists* $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ *such that*

$$\mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(m') \equiv_{\mathbf{B}} \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(F),$$

where $\mathbf{B} := \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b)$*.*

Slightly abusing notation, we write $Z'(d_f) := \{\mathbf{a}'b' \in {}^{\mathbf{x}y}F | d_f(\mathbf{a}'y) = 0\}$. For comparison, in Proposition 5.2.3 we considered $Z(d_f) = \{\mathbf{a}b \in {}^{\mathbf{x}y}F | d_f(\mathbf{a}b) = 0\}$. This is not a big difference since $Z'(d_f) \subseteq Z(d_f)$.

Proof. By Proposition 5.2.2, $\mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) = \mathrm{locus}_{\mathbf{x}}(\mathbf{a}/C) \times F \subseteq \bigcap_{f\in F} Z'(d_f)$. By Proposition 5.2.3 (part 5.), we find $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ such that $\bigcup_{f\in F} Y_f \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) = \emptyset$. Let $Z'_f := Z'(d_f) \setminus Y_f$; then $\mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) \subseteq \bigcap_{f\in F} Z'_f$.

By Proposition 5.2.3 (part 8.) and since $Z'(d_f) \subseteq Z(d_f)$; we see that for each $f \in F$ we have

$$Z(m') \cap \bigcap_{f\in F} Z'_f = Z(f) \cap \bigcap_{f\in F} Z'_f.$$

Since $Z(F) = \bigcap_{f\in F} Z(f)$, we have that

$$Z(m') \cap \bigcap_{f\in F} Z'_f = Z(F) \cap \bigcap_{f\in F} Z'_f.$$

Finally, by combining our two conclusions, we have that

$$Z(m') \cap \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) = Z(F) \cap \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b),$$

which is the required result. □

Proposition 5.2.5. *Let $\mathbf{a}b \in {}^{\mathbf{x}y}F$ and suppose that b is algebraic over $C(\mathbf{a})$. There exists $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ such that*

$$\mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(m') \equiv_{\mathbf{B}} \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}b/C),$$

where $\mathbf{B} := \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b)$.

Proof. By the Noetherianity of $C[\mathbf{x}y]$, $I(\mathbf{a}b/C)$ is finitely generated as an ideal. Let $H \subseteq_{\mathrm{fin}} C[\mathbf{x}y]$ be a finite subset such that $(H) = I(\mathbf{a}b/C)$. Thus $\mathrm{locus}_{\mathbf{x}y}(\mathbf{a}b/C) = Z(H)$. We apply Proposition 5.2.4 to find $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ such that

$$\begin{aligned} \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(m') \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) &= \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(H) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) \\ &= \mathrm{locus}_{\mathbf{x}y}(\mathbf{a}b/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b), \end{aligned}$$

which is the required result. □

Proposition 5.2.6. *Suppose that $\mathbf{y} = (y_i)_{i<n}$ is a finite tuple. Let $\mathbf{ab} \in {}^{\mathbf{xy}}F$ be such that $\mathbf{a}$ is algebraically independent over C and $\mathbf{b}$ is algebraic over $C(\mathbf{a})$. Let $\mathbf{h} = (h_i)_{i<n} \in {}^{\mathbf{y}}C[\mathbf{xy}]$ be a tuple of*

polynomials such that

$$h_i \in I(\mathbf{a}, b_j | j \leq i) \setminus I(\mathbf{a}, b_j | j < i) \subseteq C[\mathbf{x}, y_j | j \leq i],$$

for each $i < n$. *Then*

$$\operatorname{locus}(\mathbf{ab}/C) \equiv_{\mathbf{B}} Z(\mathbf{h}),$$

where $\mathbf{B}$ *is some ball around* $\mathbf{ab}$.

Proof. We proceed by induction on the length of $\mathbf{y}$. If $\mathbf{y}$ is the empty tuple, then $\operatorname{locus}(\mathbf{a}/C) = {}^{\mathtt{x}}F = Z(\emptyset)$ and the result is trivial. As an inductive hypothesis, we may suppose that there exists a ball $\mathbf{B}'''$ containing $\mathbf{ab}$ such that

$$\operatorname{locus}(\mathbf{ab}/C) \equiv_{\mathbf{B}'''} Z(\mathbf{h}).$$

Let c be algebraic over $C(\mathbf{ab})$ and let $h \in I(\mathbf{ab}c/C) \setminus I(\mathbf{ab}/C)$. By Proposition 5.2.5 there exists a ball $\mathbf{B}'$ such that

$$\operatorname{locus}(\mathbf{ab}c/C) \equiv_{\mathbf{B}'} \operatorname{locus}(\mathbf{ab}/C) \cap Z(m').$$

By Proposition 5.2.4 there exists a ball $\mathbf{B}''$ such that

$$\operatorname{locus}(\mathbf{ab}/C) \cap Z(m') \equiv_{\mathbf{B}''} \operatorname{locus}(\mathbf{ab}/C) \cap Z(h).$$

We have that

$$\begin{aligned} \operatorname{locus}(\mathbf{ab}c/C) &\equiv_{\mathbf{B}'} \operatorname{locus}(\mathbf{ab}/C) \cap Z(m') \\ &\equiv_{\mathbf{B}''} \operatorname{locus}(\mathbf{ab}/C) \cap Z(h) \\ &\equiv_{\mathbf{B}'''} Z(\mathbf{h}) \cap Z(h) \\ &= Z(\mathbf{h}h), \end{aligned}$$

as required. Letting $\mathbf{B} := \mathbf{B}' \cap \mathbf{B}'' \cap \mathbf{B}'''$, we have the required local equality. □

Chapter 6

Projections of loci

Let (F, τ) be a t-henselian field and let $C \subseteq F$ be a subfield. We will use the t-henselian topology to understand projections of loci locally, i.e. in a neighbourhood of a point.

6.1 Using the Implicit Function Theorem

Proposition 6.1.1. *Let $\mathbf{a} \in {}^{\mathbf{x}}F$ and suppose that b is separably algebraic over $C(\mathbf{a})$, and let $m' \in I(\mathbf{a}b/C), m'' \in C[\mathbf{x}] \setminus I(\mathbf{a}/C)$ be such that $m := m'/m'' = \chi\min(b/C(\mathbf{a}))$. Then there exist $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ such that*

$$Z(m') \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b)$$

is the graph of a continuous function

$$\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \longrightarrow \mathbf{B}(\beta; b).$$

Proof. First we note that D_y commutes with π. Since b is separably algebraic over $C(\mathbf{a})$, χm is irreducible and separable; thus $D_y m(\mathbf{a}b) = D_y(\pi m)(b) \neq 0$. Since m'' is constant in y, we have that $D_y m = (D_y m')/m''$. Also $m''(\mathbf{a}) \neq 0$. Thus $D_y m'(\mathbf{a}b) = (m'' \cdot D_y m)(\mathbf{a}b) = m''(\mathbf{a}) \cdot D_y m(\mathbf{a}b) \neq 0$. By Proposition 5.1.12, we may apply the Implicit Function Theorem (Property 5.1.11) in F to m' at $\mathbf{a}b$ to obtain the desired result. □

6.2 Separated projections

Let $\mathbf{ab} \in {}^{\mathbf{xy}}F$ and suppose that $C(\mathbf{ab})/C(\mathbf{a})$ is separably generated. We aim to understand the projection

$$\operatorname{locus}(\mathbf{ab}/C) \twoheadrightarrow \operatorname{locus}(\mathbf{a}/C)$$

in some neighbourhood of $\mathbf{ab}$.

6.2.1 Purely transcendental projections

Proposition 6.2.1. *Let $\mathbf{ab} \in {}^{\mathbf{xy}}F$ and suppose that $\mathbf{b}$ is purely transcendental over $C(\mathbf{a})$. Then $\operatorname{locus}(\mathbf{ab}/C) = \operatorname{locus}(\mathbf{a}/C) \times {}^{\mathbf{y}}F$.*

Proof. Since b is transcendental over $C(\mathbf{a})$, $I(b/C(\mathbf{a})) = 0$. Thus $I(\mathbf{a}/C) = I(\mathbf{ab}/C)$ and the result follows. □

6.2.2 Finite separably algebraic projections

Proposition 6.2.2. *Let $\mathbf{a}b \in {}^{\mathbf{x}y}F$ and suppose that b is separably algebraic over $C(\mathbf{a})$. There exists $\boldsymbol{\alpha}\beta \in {}^{\mathbf{x}y}\tau$ such that*

$$\operatorname{locus}_{\mathbf{x}y}(\mathbf{a}b/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b)$$

is the graph of a continuous function

$$\operatorname{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \longrightarrow \mathbf{B}(\beta; b).$$

Proof. Our aim is to find the ball $\mathbf{B}$ and a continuous function $u : \operatorname{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \longrightarrow \mathbf{B}(\beta; b)$ such that

$$\operatorname{graph}(u) \equiv_{\mathbf{B}} \operatorname{locus}_{\mathbf{x}y}(\mathbf{a}b/C).$$

Since b is algebraic over $C(\mathbf{a})$, Proposition 5.2.5 gives us that the locus over C of $\mathbf{a}b$ is locally equal to the locus over C of $\mathbf{a}$ intersected with $Z(m')$ (where m' is the numerator of the minimal polynomial of b over $C(\mathbf{a})$). Specifically: by Proposition 5.2.5 there exists a ball $\mathbf{B}' = \mathbf{B}(\boldsymbol{\alpha}'\beta'; \mathbf{a}b)$ around $\mathbf{a}b$ such that

$$\operatorname{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(m') \equiv_{\mathbf{B}'} \operatorname{locus}_{\mathbf{x}y}(\mathbf{a}b/C). \tag{6.1}$$

Since b is separably algebraic over $C(\mathbf{a})$, we may apply Proposition 6.1.1. Thus there exists a ball $\mathbf{B}'' := \mathbf{B}(\boldsymbol{\alpha}''\beta''; \mathbf{a}b)$ and a continuous function $u'' : \mathbf{B}(\boldsymbol{\alpha}''; \mathbf{a}) \longrightarrow \mathbf{B}(\beta; b)$ such that

$$\text{graph}(u'') \equiv_{\mathbf{B}''} Z(m'). \tag{6.2}$$

We now combine these two local equalities. Set $\beta := \beta' \cap \beta''$. Since the function u'' is continuous, we may choose $\boldsymbol{\alpha} \in {}^{\mathbf{x}}\tau$ to be such that $u''(\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})) \subseteq \mathbf{B}(\beta; b)$ and $\boldsymbol{\alpha} \leq \min\{\boldsymbol{\alpha}', \boldsymbol{\alpha}''\}$. Now we set $\mathbf{B} := \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b)$. Since $\mathbf{B} \subseteq \mathbf{B}' \cap \mathbf{B}''$, we have that

$$\begin{aligned} \text{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap Z(m') &\equiv_{\mathbf{B}} \text{locus}_{\mathbf{x}y}(\mathbf{a}b/C) \text{ and} \\ \text{graph}(u'') &\equiv_{\mathbf{B}} Z(m'). \end{aligned} \tag{6.3}$$

Define the function $u := u''|_{\text{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})}$ to be the restriction of u to $\text{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})$. Obviously u is continuous. Since the domain of u is a subset of $\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})$, the image of u is a subset of $\mathbf{B}(\beta; b)$. Thus u is a function $\text{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \longrightarrow \mathbf{B}(\beta; b)$, as required. For any function, restricting the domain to a set is the same as intersecting the graph with the cylinder over that set. Thus

$$\begin{aligned} \text{graph}(u) &= \text{graph}(u''|_{\text{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})}) \\ &= \text{graph}(u'') \cap \text{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap (\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \times {}^{y}F) \\ &= \text{graph}(u'') \cap \text{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b). \end{aligned}$$

The last equality is due to the fact that the image of u is a subset of $\mathbf{B}(\beta; b)$. Using Equation 6.3 we have the chain of (local) equalities

$$\begin{aligned} \text{graph}(u) &= \text{graph}(u'') \cap \text{locus}_{\mathbf{x}y}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}\beta; \mathbf{a}b) \\ &\equiv_{\mathbf{B}} \text{graph}(u'') \cap \text{locus}_{\mathbf{x}y}(\mathbf{a}/C) \\ &\equiv_{\mathbf{B}} Z(m') \cap \text{locus}_{\mathbf{x}y}(\mathbf{a}/C) \\ &\equiv_{\mathbf{B}} \text{locus}_{\mathbf{x}y}(\mathbf{a}b/C). \end{aligned}$$

□

Remark 6.2.3. Recall the tuples are not necessarily finite, unless otherwise stated. In the next proposition, we make the assumption that $\mathbf{b}$ is finite.

Proposition 6.2.4. *Let* $\mathbf{ab} \in {}^{\mathbf{xy}}F$ *and suppose that* $\mathbf{b}$ *is finite and separably algebraic over* $C(\mathbf{a})$.

There exists $\boldsymbol{\alpha\beta} \in {}^{\mathbf{xy}}\tau$ *such that*

$$\mathrm{locus}_{\mathbf{xy}}(\mathbf{ab}/C) \cap \mathbf{B}(\boldsymbol{\alpha\beta}; \mathbf{ab})$$

is the graph of a continuous function

$$\mathrm{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \longrightarrow \mathbf{B}(\boldsymbol{\beta}; \mathbf{b}).$$

Proof. We prove this by induction on the length of the tuple $\mathbf{b}$. Proposition 6.2.2 is the base case. As an inductive hypothesis, we may assume that there exists a ball $\mathbf{B}' := \mathbf{B}(\boldsymbol{\alpha}'\boldsymbol{\beta}'; \mathbf{ab})$ and a continuous function $\mathbf{u} : \mathrm{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}'; \mathbf{a}) \longrightarrow \mathbf{B}(\boldsymbol{\beta}'; \mathbf{b})$ such that

$$\mathrm{graph}(\mathbf{u}) \equiv_{\mathbf{B}'} \mathrm{locus}_{\mathbf{xy}}(\mathbf{ab}/C). \tag{6.4}$$

We must prove the inductive step. Let $b \in {}^{y}F$. We aim to find a ball $\mathbf{B} := \mathbf{B}(\boldsymbol{\alpha\beta}\beta; \mathbf{ab}b)$ and a continuous function $\mathbf{v} : \mathrm{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) \longrightarrow \mathbf{B}(\boldsymbol{\beta}\beta; \mathbf{b}b)$ such that

$$\mathrm{graph}(\mathbf{v}) \equiv_{\mathbf{B}} \mathrm{locus}_{\mathbf{xy}y}(\mathbf{ab}b/C).$$

By applying Proposition 6.2.2, we know that there exists a ball $\mathbf{B}'' := \mathbf{B}(\boldsymbol{\alpha}''\boldsymbol{\beta}''\beta''; \mathbf{ab}b)$ and a continuous function $u : \mathrm{locus}_{\mathbf{xy}}(\mathbf{ab}/C) \cap \mathbf{B}(\boldsymbol{\alpha}''\boldsymbol{\beta}''; \mathbf{ab}) \longrightarrow \mathbf{B}(\beta''; b)$ such that

$$\mathrm{graph}(u) \equiv_{\mathbf{B}''} \mathrm{locus}_{\mathbf{xy}y}(\mathbf{ab}b/C). \tag{6.5}$$

Let $\boldsymbol{\beta} := \min\{\boldsymbol{\beta}', \boldsymbol{\beta}''\}$; then we choose $\boldsymbol{\alpha} \in {}^{\mathbf{x}}\tau$ so that $\mathbf{u}(\mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})) \subseteq \mathbf{B}(\boldsymbol{\beta}; \mathbf{b})$ and $\boldsymbol{\alpha} \leq \min\{\boldsymbol{\alpha}', \boldsymbol{\alpha}''\}$. We re-write $\beta := \beta''$. Let $\mathbf{B} := \mathbf{B}(\boldsymbol{\alpha\beta}\beta; \mathbf{ab}b)$. Since $\mathbf{B} \subseteq \mathbf{B}' \cap \mathbf{B}''$, we have that

$$\begin{aligned} \mathrm{graph}(\mathbf{u}) &\equiv_{\mathbf{B}} \mathrm{locus}_{\mathbf{xy}}(\mathbf{ab}/C) \\ \mathrm{graph}(u) &\equiv_{\mathbf{B}} \mathrm{locus}_{\mathbf{xy}y}(\mathbf{ab}b/C). \end{aligned} \tag{6.6}$$

Now we define the function

$$\begin{aligned} \mathbf{v} : \mathrm{locus}_{\mathbf{x}}(\mathbf{a}/C) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a}) &\longrightarrow \mathbf{B}(\boldsymbol{\beta}\beta; \mathbf{b}b) \\ \mathbf{x} &\longmapsto (\mathbf{u} \times u)(\mathbf{x}) = (\mathbf{u}(\mathbf{x}), u(\mathbf{x}, \mathbf{u}(\mathbf{x}))). \end{aligned}$$

Since the domain of $\mathbf{v}$ is a subset of $\mathbf{B}(\boldsymbol{\alpha};\mathbf{a})$, the image of $\mathbf{v}$ is a subset of $\mathbf{B}(\boldsymbol{\beta}\beta;\mathbf{b}b)$

Using Equation 6.6 we have the chain of (local) equalities

$$\begin{aligned}
\text{graph}(\mathbf{v}) &= \{(\mathbf{x},\mathbf{u}(\mathbf{x}),u(\mathbf{x},\mathbf{u}(\mathbf{x})))|\mathbf{x}\in\text{locus}_{\mathbf{x}}(\mathbf{a}/C)\cap\mathbf{B}(\boldsymbol{\alpha};\mathbf{a})\}\\
&= \text{graph}(u)\cap(\text{graph}(\mathbf{u})\times {}^{y}F)\cap\text{locus}_{\mathbf{xy}y}(\mathbf{a}/C)\cap(\mathbf{B}(\boldsymbol{\alpha};\mathbf{a})\times {}^{\mathbf{y}y}F)\\
&\equiv_{\mathbf{B}} \text{graph}(u)\cap(\text{graph}(\mathbf{u})\times {}^{y}F)\cap\text{locus}_{\mathbf{xy}y}(\mathbf{a}/C)\\
&\equiv_{\mathbf{B}} \text{locus}_{\mathbf{xy}y}(\mathbf{ab}b/C)\cap\text{locus}_{\mathbf{xy}}(\mathbf{ab}/C)\cap\text{locus}(\mathbf{a}/C)\\
&\equiv_{\mathbf{B}} \text{locus}_{\mathbf{xy}y}(\mathbf{ab}b/C).
\end{aligned}$$

□

6.2.3 Finitely generated separable projections

Proposition 6.2.5. *Let* $\mathbf{ab}\in {}^{\mathbf{xy}}F$ *and suppose that* $\mathbf{b}$ *is finite and separable over* $C(\mathbf{a})$*. There exists a partition* $\mathbf{y}=\mathbf{y}'\mathbf{y}''$ *and* $\boldsymbol{\alpha}\boldsymbol{\beta}'\boldsymbol{\beta}''\in {}^{\mathbf{xy}'\mathbf{y}''}\tau$ *such that*

$$\text{locus}_{\mathbf{xy}}(\mathbf{ab}/C)\cap\mathbf{B}(\boldsymbol{\alpha}\boldsymbol{\beta};\mathbf{ab})$$

is the graph of a continuous function

$$\text{locus}_{\mathbf{xy}'}(\mathbf{a}/C)\cap\mathbf{B}(\boldsymbol{\alpha}\boldsymbol{\beta}';\mathbf{ab}')\longrightarrow\mathbf{B}(\boldsymbol{\beta}'';\mathbf{b}'').$$

Proof. Since $\mathbf{b}$ is separable and finite over $C(\mathbf{a})$ and by definition of separability, there exists a separating partition $\mathbf{b}'\mathbf{b}''$ of $\mathbf{b}$. This is a partition such that $\mathbf{b}'$ is purely transcendental over $C(\mathbf{a})$ and $\mathbf{b}''$ is separably algebraic $C(\mathbf{a},\mathbf{b}')$. We now simply apply Proposition 6.2.4 to find $\boldsymbol{\alpha}\boldsymbol{\beta}\in {}^{\mathbf{xy}}\tau$ such that

$$\text{locus}_{\mathbf{xy}}(\mathbf{ab}/C)\cap\mathbf{B}(\boldsymbol{\alpha}\boldsymbol{\beta};\mathbf{ab})$$

is the graph of a continuous function

$$\text{locus}_{\mathbf{xy}'}(\mathbf{ab}'/C)\cap\mathbf{B}(\boldsymbol{\alpha}\boldsymbol{\beta}';\mathbf{ab}')\longrightarrow\mathbf{B}(\boldsymbol{\beta}'';\mathbf{b}'').$$

Since $\mathbf{b}'$ is purely transcendental over $C(\mathbf{a})$ we may apply Proposition 6.2.1 to find that $\text{locus}_{\mathbf{xy}'}(\mathbf{ab}'/C)=\text{locus}_{\mathbf{xy}'}(\mathbf{a}/C)$. □

6.3 Inseparable projections

Recall the global and local alterations from chapter 4. Let $D := \boldsymbol{\Lambda}(F/C)$ denote the Λ-closure of C in F.

Proposition 6.3.1. (**Local method**) *Let* $\mathbf{ab} \in {}^{\mathbf{xy}}F$ *and suppose that* $\mathbf{b}$ *is finite. Let* $(\boldsymbol{\Sigma}, \boldsymbol{\Lambda})$ *denote the local* Λ*-alteration of* $\mathbf{b}/\mathbf{a}$ *with respect to* D*. Then there exists* $\boldsymbol{\alpha\beta} \in {}^{\mathbf{xy'y''}}\tau$ *such that*

$$\Sigma(\mathrm{pr}_{\boldsymbol{\lambda x}}(\mathrm{locus}(\boldsymbol{\lambda a}, \mathbf{b}/D) \cap \mathbf{B})) = \mathrm{pr}_{\mathbf{x}}(\mathrm{locus}_{\mathbf{xy}}(\mathbf{ab}/D) \cap \mathbf{U}),$$

for a ball $\mathbf{B}$ *and an open set* $\mathbf{U} \subseteq {}^{\mathbf{xy}}F$*.*

The tuple $\boldsymbol{\lambda x}$ is just a tuple of variables corresponding to $\boldsymbol{\lambda a}$.

Proof. The extension $D(\boldsymbol{\lambda a}, \mathbf{b})/D(\boldsymbol{\lambda a})$ is separable. By Proposition 6.2.5, there exists a partition $\mathbf{y} = \mathbf{y'y''}$ and $\boldsymbol{\lambda\alpha}, \boldsymbol{\beta'\beta''} \in {}^{\boldsymbol{\lambda x},\mathbf{y'y''}}\tau$ such that

$$\mathrm{locus}_{\boldsymbol{\lambda x},\mathbf{y}}(\boldsymbol{\lambda a}, \mathbf{b}/D) \cap \mathbf{B}(\boldsymbol{\lambda\alpha}, \boldsymbol{\beta}; \boldsymbol{\lambda a}, \mathbf{b})$$

is the graph of a continuous function

$$\mathrm{locus}_{\boldsymbol{\lambda x},\mathbf{y'}}(\boldsymbol{\lambda a}/D) \cap \mathbf{B}(\boldsymbol{\lambda\alpha}, \boldsymbol{\beta'}; \boldsymbol{\lambda a}, \mathbf{b'}) \longrightarrow \mathbf{B}(\boldsymbol{\beta''}; \mathbf{b''}).$$

We set $U^{\boldsymbol{\Sigma}} := \mathrm{locus}(\boldsymbol{\lambda a}, \mathbf{b}/D) \cap \mathbf{B}(\boldsymbol{\lambda\alpha}, \boldsymbol{\beta}; \boldsymbol{\lambda a}, \mathbf{b})$. As in Proposition 4.2.7, there exists an open subset $U^{\boldsymbol{\Lambda}} := \boldsymbol{\Sigma} U^{\boldsymbol{\Sigma}} \subseteq \mathrm{locus}(\mathbf{ab}/D)$. Applying Proposition 4.2.7, we have that $\Sigma(\mathrm{pr}_{\boldsymbol{\lambda x}}(U^{\boldsymbol{\Sigma}})) = \mathrm{pr}_{\mathbf{x}}(U^{\boldsymbol{\Lambda}})$, as required. □

Proposition 6.3.2. (**Global method**) *Let* $\mathbf{ab} \in {}^{\mathbf{xy}}F$ *and suppose that* $\mathbf{b}$ *is finite. Let* $\boldsymbol{\Sigma}$ *denote the global* Λ*-alteration of* $\mathbf{b}/\mathbf{a}$ *with respect to* E/D*. There there exists a partition* $\mathbf{y} = \mathbf{y'y''}$ *and* $\boldsymbol{\alpha\beta'\beta''} \in {}^{\mathbf{xy'y''}}\tau$ *such that*

$$\mathrm{locus}_{\mathbf{xy}}(\mathbf{ab}/D) \cap \mathbf{B}(\boldsymbol{\alpha\beta}; \mathbf{ab})$$

contains the graph of a continuous function

$$\Sigma(\mathrm{locus}_{\boldsymbol{\lambda x},\mathbf{y'}}(\boldsymbol{\lambda a}/D) \cap \mathbf{B}(\boldsymbol{\alpha}; \mathbf{a})) \times \mathbf{B}(\boldsymbol{\beta'}; \mathbf{b'}) \longrightarrow \mathbf{B}(\boldsymbol{\beta''}; \mathbf{b''}).$$

Thus $\Sigma(\mathrm{pr}_{\boldsymbol{\lambda}\mathbf{x}}(\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}, \mathbf{b}/D) \cap \mathbf{B})) \subseteq \mathrm{pr}_{\mathbf{x}}(\mathrm{locus}(\mathbf{ab}/D))$.

Proof. This proof is similar to the previous one, only this time we use Proposition 4.2.6. The main difference in this case is that the map Σ is injective; thus the image under $\boldsymbol{\Sigma}$ of a function is still a function. We do not have equality, only containment, because in the global method there are no maps $\Lambda, \boldsymbol{\Lambda}$. □

Chapter 7

Existential definability

Let (K, τ) be a non-separably closed t-henselian field, let $C \subseteq K$ be a subfield of parameters, and let $D := \boldsymbol{\Lambda}(K/C)$ be the Λ-closure of C in K.

We study existentially C-definable sets. We first present a description of their local behaviour using the techniques developed in chapter 6. We can then draw many conclusions, especially in the case of definable subsets of (the first Cartesian power of) K. We are able to classify the subfields which are generated by infinite existentially definable subsets.

Fact 7.0.3. By several well-known reductions, existentially definable sets are simply projections of algebraic sets. In fact in fields that are not algebraically closed, we may assume that the algebraic set is the zero set of a single polynomial. More precisely, let ϕ be an existential formula (with free variables $\mathbf{x}$); then there exists a polynomial f (in the variables $\mathbf{xy}$) such that ϕ is equivalent to $\exists \mathbf{y}\ f \doteq 0$ modulo the theory of K.

7.0.1 Elementary extensions

Several of our arguments will require passing to an elementary extension of K. The next proposition proves that nothing 'goes wrong' with Λ-closure in an elementary extension. For background on Model Theory, including the Compactness Theorem, see [20].

Proposition 7.0.4. *Let F/E be any field extension, let $F \preceq F^*$ be an $\mathcal{L}_{\mathrm{ring}}$-elementary extension, and suppose that E is Λ-closed in F. Then E is Λ-closed in F^*.*

Proof. Let $\mathbf{a} \in pB(E)$ be a p-base of E. Then $\mathbf{a}$ is p-independent in F and this information is encoded in the $\mathcal{L}_{\text{ring}}$-type of $\mathbf{a}$. Thus $\mathbf{a}$ is p-independent in F^* and E is Λ-closed in F^*. □

7.1 ∃-definable sets

The following technical theorem is the summary of the work so far on projections of loci. From it we will deduce all of the subsequent corollaries.

Theorem 7.1.1. *Let $Y \subseteq {}^{\mathbf{xy}}K$ be an algebraic set defined over C and let $\mathbf{ab} \in Y$. Let $\boldsymbol{\lambda}\mathbf{a}$ be the λ-splitting of $\mathbf{a}$ under either the local or the global Λ-alteration. There exists $\boldsymbol{\lambda}\boldsymbol{\alpha} \in {}^{\boldsymbol{\lambda}\mathbf{x}}\tau$ such that*

$$\Sigma\left(\operatorname{locus}_{\boldsymbol{\lambda}\mathbf{x}}(\boldsymbol{\lambda}\mathbf{a}/D) \cap \mathbf{B}(\boldsymbol{\lambda}\boldsymbol{\alpha}; \boldsymbol{\lambda}\mathbf{a})\right) \subseteq X := \operatorname{pr}_{\mathbf{x}}(Y).$$

Proof. This just summarises Proposition 6.3.2 and Proposition 6.3.1. □

The following corollary (which we will not prove) is the starting point for future investigations of possible quantifier-elimination theorems using this local approach.

Corollary 7.1.2. (to Theorem 7.1.1) *When $\mathbf{ab}$ is chosen to be sufficiently generic (with respect to the algebraic set Y) then by the proof of Proposition 6.3.1 we have that Y is locally equal to the graph of a function with domain which projects onto an open subset of* $\operatorname{locus}(\boldsymbol{\lambda}\mathbf{a}/C)$*. If we work in the context of the compact ring $\mathbb{F}_p[[t]]$ and if all of the points of Y above $\mathbf{a}$ are sufficiently generic, then the projection of Y is locally (around $\mathbf{a}$) equal to* $\operatorname{locus}(\boldsymbol{\lambda}\mathbf{a}/C)$*, which is a positive quantifier-free definable set in the ring language expanded by symbols for the component maps.*

Since none of our applications require it, in this thesis we will neither make precise nor investigate the notion of 'sufficiently generic'. However, for possible improvements to this theorem (in particular, true quantifier-elimination results) it is very important to understand precisely at which points existentially definable sets are locally positive quantifier-free (in the expanded language).

7.2 ∃-type

Applying the global method to all existentially definable sets containing $\mathbf{a}$, we obtain the following description of the realisations of the existential type of $\mathbf{a} \in {}^{\mathbf{x}}K$ (over C) in an elementary extension $K \preceq K^*$.

Definition 7.2.1. Let $\mathbf{a} \in {}^{\mathbf{x}}K$ and let $K \preceq K^*$. We define the *K-infinitesimal neighbourhood* of $\mathbf{a}$ in K^* to be

$$\mathcal{V}(\mathbf{a}/K) := \bigcap \{\mathbf{B}^*(\boldsymbol{\alpha}; \mathbf{a}) \mid \boldsymbol{\alpha} \in {}^{\mathbf{x}}\tau\},$$

where $\mathbf{B}^*$ denotes that the ball is in K^*. That is, $\mathcal{V}(\mathbf{a}/K)$ is the intersection of all the open sets (in the t-henselian topology on K^*, defined over K) which contain $\mathbf{a}$.

Theorem 7.2.2. *Let $\mathbf{a} \in {}^{\mathbf{x}}K$ and let $\boldsymbol{\lambda}\mathbf{a} := \mathbf{a}^\infty$ be the full global splitting of $\mathbf{a}$ with respect to K/C. Then*

$$\Sigma(\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/D) \cap \mathcal{V}(\boldsymbol{\lambda}\mathbf{a}/K)) \subseteq \mathrm{tp}_\exists(\mathbf{a}/C).$$

If C is dense in K then this is an equality.

Proof. If we apply the global form of Theorem 7.1.1 to each existentially C-definable set X that contains $\mathbf{a}$, then we find $\boldsymbol{\lambda}\boldsymbol{\alpha}$ such that

$$\Sigma(\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/D) \cap \mathbf{B}^*(\boldsymbol{\lambda}\boldsymbol{\alpha}; \boldsymbol{\lambda}\mathbf{a})) \subseteq X.$$

To obtain the desired result, we take the intersection over all existentially C-definable sets that contain $\mathbf{a}$.

If C is dense in K then, for each $\boldsymbol{\lambda}\boldsymbol{\alpha}$, there exists $\boldsymbol{\lambda}\boldsymbol{\gamma}$ defined over C such that $\boldsymbol{\lambda}\boldsymbol{\gamma} \subseteq \boldsymbol{\lambda}\boldsymbol{\alpha}$. Thus $\mathbf{B}^*(\boldsymbol{\lambda}\boldsymbol{\gamma}; \boldsymbol{\lambda}\mathbf{a}) \subseteq \mathbf{B}^*(\boldsymbol{\lambda}\boldsymbol{\alpha}; \boldsymbol{\lambda}\mathbf{a})$; and so

$$\Sigma(\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/D) \cap \mathbf{B}^*(\boldsymbol{\lambda}\boldsymbol{\gamma}; \boldsymbol{\lambda}\mathbf{a})) \subseteq \Sigma(\mathrm{locus}(\boldsymbol{\lambda}\mathbf{a}/D) \cap \mathbf{B}^*(\boldsymbol{\lambda}\boldsymbol{\alpha}; \boldsymbol{\lambda}\mathbf{a})) \subseteq X.$$

Since the left-hand side is already an existentially C-definable set, we have equality in the intersection.

□

7.3 ∃-definable subsets

We state three corollaries with very similar proofs. In each case we find a transcendental element and apply the global method.

7.3.1 Topology and Cardinality

Theorem 7.3.1. *Let $X \subseteq K$ be an infinite $\exists$-definable subset. Then X has an accumulation point and $|X| = |K|$.*

Proof. In an elementary extension $K \preceq K^*$ we can find an element a of X which is transcendental over D. We apply the global method in the form of Theorem 1: there exists $\boldsymbol{\lambda}\alpha \in {}^{\boldsymbol{\lambda}x}\tau$ such that

$$\Sigma(\mathrm{locus}_{\boldsymbol{\lambda}x}(\boldsymbol{\lambda}a/D) \cap \mathbf{B}(\boldsymbol{\lambda}\alpha; \boldsymbol{\lambda}a)) \subseteq X.$$

Since K/D is separable, we may partition $\boldsymbol{\lambda}a = a_i^n a_s^n$ so that a_i^n is algebraically independent over D and a_s^n is separably algebraic over $D(a_i^n)$; and, since a is transcendental over D, a_i^n is not empty. The locus of $\boldsymbol{\lambda}a$ over D is (by an application of Proposition 6.2.3) locally equal to the graph of a continuous function with domain a ball $\mathbf{B}(\boldsymbol{\alpha}_i; a_i^n)$, for some tuple of neighbourhoods $\boldsymbol{\alpha}_i \subseteq \tau^*$.

Since $K \preceq K^*$ is an elementary extension, there exist $a' \in K$ and $\boldsymbol{\alpha}_i' \subseteq \tau$ such that X contains the image under Σ of (the graph of) a continuous function (let us call it f) with domain $\mathbf{B}(\boldsymbol{\alpha}_i'; a_i'^n)$. This ball has cardinality $|K|$, as do all balls. Thus the graph of f has cardinality $|K|$. Since Σ is injective in the global method, X must also have cardinality $|K|$. Since f is continuous, we have that $\boldsymbol{\lambda}a'$ is an accumulation point of (the graph of) f. Since Σ is injective and continuous, $a' = \Sigma(\boldsymbol{\lambda}a')$ is an accumulation point of the image under Σ of f. Thus a' is an accumulation point of X. □

Theorem 7.3.2. *Let ϕ be an $\exists$-formula. Then there is uniform finiteness in each variable. I.e. let $\mathbf{x}y$ be any partition of the free-variables of ϕ so that y is a singleton: if for each $\mathbf{a} \in {}^{\mathbf{x}}K$ the set defined by $\phi(\mathbf{a}y)$ is finite; then there exists $n \in \mathbb{N}$ such that, for each $\mathbf{a} \in {}^{\mathbf{x}}K$, the set defined by $\phi(\mathbf{a}y)$ has at most n elements.*

Proof. Suppose not. Then in some elementary extension $K \preceq K^*$ there exists $\mathbf{a} \in {}^{\mathbf{x}}K^*$ such that $\phi(\mathbf{a}y)$ defines an infinite set in K^*. Infinite existentially definable sets have accumulation points, which is an elementary property, thus there must exist $\mathbf{a}' \in {}^{\mathbf{x}}K$ such that $\phi(\mathbf{a}'y)$ defines a set with an accumulation point. This set must of course be infinite, which is a contradiction. □

7.3.2 $\exists$-definable elements

In the next theorem we consider different notions of closure. We say that a is *algebraic* over C if a is the root of a non-zero polynomial over C; we say that a is *existentially-algebraic* over C if a is

contained in a finite existentially-C-definable subset of K; and we say that a is *existentially-definable* over C if $\{a\}$ is an existentially-C-definable subset of K. The operation of *closure* in K with respect to one of these definitions of dependence is defined in the usual way.

Theorem 7.3.3. *The existentially-algebraic closure of C in K equals the relative algebraic closure of $D = \boldsymbol{\Lambda}(K/C)$ in K. If C is dense in K then the existentially-definable closure of C in K also coincides with the relative algebraic closure of D and the existentially-algebraic closure of C.*

Proof. If a is transcendental over D then, by the proof of Theorem 7.3.1, any existentially C-definable set containing a must be infinite. Thus a is not existentially-algebraic over C, so the existentially-algebraic closure of C in K is contained in the relative algebraic closure of D in K.

Next we observe that each element of D is existentially-definable over C. Those elements that are algebraic over D are thus existentially-algebraic over C. This proves that the existentially-algebraic closure of C equals the relative algebraic closure of D; this is the first equality.

If in addition we have that C is dense in K, then each element algebraic over D is existentially-definable over C (this is proved by separating the roots of a polynomial over D using open sets defined over C). This proves the second equality. □

7.4 Aside: measuring algebraicity and transcendence

We aim to improve our understanding of infinite existentially definable subsets X of K from the previous section. The basic idea is that such sets X are 'big'.

In order to make this intuition precise, we define and study the notion of *algebraic exponent* of a set over a field.

7.4.1 Algebraic exponent

Let $Z \subseteq K$ be an arbitrary subset of K (not necessarily definable) and let $F \subseteq K$ be any subfield.

Definition 7.4.1. Let $a \in K$. We denote the *algebraic exponent* of a over F by $\mathrm{algex}(a/F)$; and define it to be the degree of the field extension $F(a)/F$. This takes values in $\mathbb{N} \cup \{\infty\}$ (we write ∞ instead of $\aleph_0$). We define the *algebraic exponent* of a set $Z \subseteq K$ over F to be $\mathrm{algex}(Z/F) := \sup\{\mathrm{algex}(a/F) \mid a \in Z\}$, i.e. the supremum of the algebraic exponents of the elements of Z over F. Thus $\mathrm{algex}(Z/F) \in \mathbb{N} \cup \{\infty\}$.

Remark 7.4.2. It is important to distinguish the algebraic exponent from our previous notion of inseparability exponent. Recall that the *inseparability exponent* of a over F is the least $n \in \mathbb{N}$ such that a^{p^n} is separably algebraic over F, or ∞ is no such n exists. Both notions of exponent measure a certain sort of degree of a simple extension; for a set, we measure the maximum of these degrees over the collection of simple extensions generated by each element of the set. For the algebraic exponent, we measure the degree; for the inseparability exponent we measure the inseparable degree.

Lemma 7.4.3. *Let a be algebraic over F. Then* $\operatorname{algex}(F(a)/F) = [F(a) : F]$.

Proof. Obvious since the degree of a over F is the greatest of any element of $F(a)$. □

Lemma 7.4.4. *Let G/F be a separably algebraic extension. The following are equivalent.*

1. *G/F is simple, i.e. primitive, generated by one element;*
2. *G/F is finitely generated; and*
3. *G/F is of finite algebraic exponent.*

Proof. (1) $\iff$ (2) are well known to be equivalent.

(2) $\implies$ (3) is easy since a finitely generated separably algebraic extension is simple (or primitive; i.e. generated by one element). Any simple algebraic extension is of finite degree.

$\neg(2) \implies \neg(3)$ Suppose that G/F is not finitely generated. Then there exists a strictly increasing chain $(H_i)_{i<\omega}$ of fields; each a finitely generated extension of F and contained in G. Since each of them is finitely generated and separably algebraic over F; each is simple over F. Thus $\operatorname{algex}(H_i/F) = [H_i : F]$, which is a sequence cofinal in ω. Thus $\operatorname{algex}(G/F)$ is infinite. □

7.4.2 Describing algebraicity and transcendence elementarily

Now we suppose that Z is definable. We briefly describe formulas and types that describe the algebraicity or transcendence of elements over the field (Z) generated by Z. Note that (Z) is not *a priori* definable.

Definition 7.4.5. We say a is of degree $n \in \mathbb{N}$ over Z if there exists a polynomial of degree n with coefficients from Z of which a is a root. If a has degree n over Z, for some $n \in \mathbb{N}$, then we say that a is *algebraic* over Z. If no such $n \in \mathbb{N}$ exists then we say that a is of infinite degree over Z, and that a is *transcendental* over Z.

Note that if a is of degree n over (Z) it does not necessarily follow that a is of degree n over Z.

Definition 7.4.6. Let ϕ be a formula with one free variable which defines Z in K, so that $Z = \phi(K)$. We let

$$\delta_Z^n := \forall y_0 ... \forall y_n \left(\left(\bigwedge_{i \leq n} \phi(y_i) \wedge \neg \bigwedge_{i \leq n} y_i \doteq 0 \right) \longrightarrow \neg \sum_{i \leq n} y_i \cdot x^i = 0 \right)$$

be the formula defining those elements of degree greater than n over Z.

For example, ϕ might define a set Z which is the intersection of a ball with the projection of the complement of an algebraic set.

Lemma 7.4.7. *Let $Y \subseteq Z$ be two definable subsets and let $n' \leq n$. Then $K \models \forall x(\delta_Z^n \longrightarrow \delta_Y^{n'})$.*

Proof. If a is of degree greater than $n \geq n'$ over $Z \supseteq Y$, then certainly a cannot be a root of any polynomial of degree n' over Y. □

Definition 7.4.8. Let $tr_Z := \{\delta_Z^n | n \in \mathbb{N}\}$ be the 1-type defining those elements transcendental over the set Z.

7.4.3 Generating fields

Definition 7.4.9. Let Z be any set. Then $\langle Z \rangle$ and (Z) are the closure of Z under the ring and field operations, respectively. These can be constructed by a simple recursion over the ordinal ω. For example, we could set $(Z)_0, \langle Z \rangle_0 := Z$ and let

$$\langle Z \rangle_{n+1} := \langle Z \rangle_n \cup (\langle Z \rangle_n + \langle Z \rangle_n) \cup \ -\langle Z \rangle_n \cup (\langle Z \rangle_n \cdot \langle Z \rangle_n)$$

and

$$(Z)_{n+1} := (Z)_n \cup ((Z)_n + (Z)_n) \cup \ -(Z)_n \cup ((Z)_n \cdot (Z)_n) \cup ((Z)_n \setminus \{0\})^{-1}.$$

There are of course many ways of organising these recursions, but we fix one particular way and denote by $\langle Z \rangle_n, (Z)_n$ the n-th set constructed in each recursion. We simply require that $(\langle Z \rangle_n)_{n<\omega}$ and $((Z)_n)_{n<\omega}$ are increasing chains of sets with unions equal to the ring and field generated by Z, respectively.

Fact 7.4.10. We may choose the sequences $(\langle Z \rangle_n)_{n<\omega}$ and $((Z)_n)_{n<\omega}$ such that if Z is definable/existentially definable then both $\langle Z \rangle_n$ and $(Z)_n$ are definable/existentially definable, respectively and for each $n < \omega$.

Definition 7.4.11. Suppose that Z is a definable subset. We let

$$tr_{(Z)} := \bigcup_{m<\omega} tr_{(Z)_m} = \{\delta^n_{(Z)_m} | m, n \in \mathbb{N}\}$$

be the 1-type which defines elements transcendental over (Z). Note that if (Z) happened to be already definable, this notation would coincide with the previous meaning of $tr_{(Z)}$.

Definition 7.4.12. We recursively define types defining algebraic independence over (Z). Let x be the (single) free variable in the definition of Z and let $\mathbf{y}$ be an n-tuple of different variables. We write $tr^1_{(Z)}(x) := tr_{(Z)}(x)$ and let

$$tr^{n+1}_{(Z)}(x\mathbf{y}) := tr^n_{(Z)}(\mathbf{y}) \cup tr_{(Z\cup\mathbf{y})}(x)$$

be the $n+1$-type defining the $n+1$-tuples algebraically independent over (Z).

Remark 7.4.13. Strictly $Z \cup \mathbf{y}$ is not a definable subset of K since it has the $n+1$-tuple $x\mathbf{y}$ as free variables. Of course, this is not a problem: we think of $Z \cup \mathbf{y}$ as a definable subset of K in the variable x with $\mathbf{y}$ a tuple of *parameters*.

7.5 Aside: big subfields and subsets

7.5.1 Big subfields

In this section we briefly step out of the context of t-henselian fields and consider any arbitrary field extension G/E. We assume that E is a 'big' subfield of G (in a precise sense) and find that this puts a very great restriction on E. Although definability does not matter in this section, in the next section we will apply the idea of big subfields to subfields of t-henselian fields that are generated by infinite $\exists$-definable subsets.

Definition 7.5.1. We say that E is *big* (in G) if $\operatorname{algex}(G/E) < \infty$.

We suppose that E is big in G.

Let $F := E^{\text{sep}} \cap EG^{p^\infty}$ be the relative separable algebraic closure of E in EG^{p^∞}, i.e. the compositum of all separably algebraic extensions of E coming from elements of G^{p^∞}.

We have the following situation.

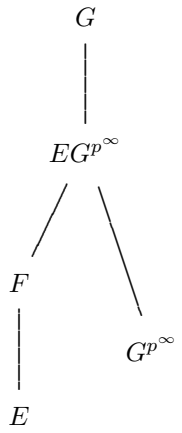

Lemma 7.5.2. *F/E is a finite separable extension of degree equal to* $\text{algex}(F/E) \leq \text{algex}(G/E)$.

Proof. Since $\text{algex}(G/E) < \infty$, then also $\text{algex}(F/E) < \infty$. By definition of F, F/E is separably algebraic; thus, by Lemma 7.4.4, F/E is a finite extension. □

Lemma 7.5.3. $G^{p^\infty} \subseteq F$.

Proof. Otherwise, let $a \in G^{p^\infty} \setminus F$. We already know that EG^{p^∞}/F is a purely inseparable extension. Thus $F(a)/F$ is purely inseparable and $[F(a^{p^{-n}}) : F] \geq [F(a^{p^{-n}}) : F(a)] \geq p^n$ and $a^{p^{-n}} \in G^{p^\infty}$, for every $n \in \mathbb{N}$. Therefore $\text{algex}(G^{p^\infty}/F) = \infty$; so that $\text{algex}(G/E) = \infty$, which is a contradiction to the assumption that E is big in G. □

Proposition 7.5.4. $F = EG^{p^\infty}$.

Proof. Immediate from the definition of F and Lemma 7.5.3. □

Lemma 7.5.5. *EG^{p^∞}/E is a finite separable extension of degree* $\leq \text{algex}(G/E)$.

Proof. From Lemma 7.5.2 and Proposition 7.5.4. □

Fact 7.5.6. Let F/E be a field extension. Let a be of degree n over E and over F. Then the minimal polynomial of a over F is an element of $E[x]$.

Proof. Let $m_1 \in F[x]$ and $m_2 \in E[x]$ be the minimal polynomials. Then $m_1 | m_2$ in the ring $F[x]$. Since they are both monic and of the same degree, they are equal. □

Lemma 7.5.7. *Let $a \in G^{p^\infty}$. Then a is separably algebraic over $G^{p^\infty} \cap E$ of degree $[E(a) : E]$.*

Proof. By Lemma 7.5.5, the degrees of each element of $\{a^{p^n} | n \in \mathbb{Z}\}$ over E are bounded by $\text{algex}(G/E)$. Let $N := \max\{\text{algex}(a^{p^n}/E) | n \in \mathbb{Z}\}$ be the maximum of these degrees and choose $b \in \{a^{p^n} | n \in \mathbb{Z}\}$ to be such that $\text{algex}(b/E) = N$.

Let $m \in E[x]$ be the minimal polynomial of b over E and let $n \in \mathbb{N}$. By taking p^n-th roots, $b^{p^{-n}}$ has degree N over $E^{p^{-n}}$ and has minimal polynomial $m^{(p^{-n})}$, i.e. the polynomial m with the coefficients replaced by their p^n-th roots.

By maximality of the degree of b over E, $b^{p^{-n}}$ also has degree N over E. By applying Fact 7.5.6, $m^{(p^{-n})} \in E[x]$, i.e. the coefficients of m have p^n-roots in E, for all $n \in \mathbb{N}$. Thus $m \in (G^{p^\infty} \cap E)[x]$. Consequently, every element of $\{a^{p^n} | n \in \mathbb{Z}\}$ has the same minimal polynomial over E as over $G^{p^\infty} \cap E$. □

Proposition 7.5.8. *$G^{p^\infty}/(G^{p^\infty} \cap E)$ is a finite separable extension and* $\text{algex}(G^{p^\infty}/(G^{p^\infty} \cap E)) \leq \text{algex}(EG^{p^\infty}/E)$.

Proof. Let $a \in G^{p^\infty}$. By Lemma 7.5.7, the degree of a over $G^{p^\infty} \cap E$ is the same as its degree over E. □

7.5.2 Big subsets

Let $Z \subseteq G$ be any subset of G, not necessarily definable.

Definition 7.5.9. We say that Z is *big* if there exists $n \in \mathbb{N}$ such that $G \models \forall x\ \neg\delta_Z^n$. That is to say that all elements of G have degree less than or equal to n over Z. Thus Z is not big if $G \models \exists x\ \delta_Z^n$ for each $n \in \mathbb{N}$.

This is compatible with the previous definition of 'big': in the case that Z is already a field, Z is big as a set if and only if Z is big as a field.

Fact 7.5.10. Since Z is assumed to be definable in G, 'Z is big in G' is an elementary property (in the theory of G).

Lemma 7.5.11. *Suppose that $(Z)_n$ is big, for some $n \in \mathbb{N}$, Then (Z) is big in G.*

Proof. By assumption, all elements of G have degree less than or equal to n over $(Z)_n$. Thus *a fortiori*, all elements of G have degree less than or equal to n over $(Z) \supseteq (Z)_n$. Thus $\text{algex}(G/(Z)) \leq n$. □

However, it is not true that if (Z) is big then $(Z)_n$ is big, for some $n \in \mathbb{N}$.

Example 7.5.12. Let $G := \mathbb{F}_p(t)$ and let $Z := \{t\}$. Then $(Z) = G$. For each $m, n \in \mathbb{N}$, $(Z)_m$ is finite, so there are only finitely many elements which are roots of non-trivial polynomials of degree $\leq n$ with coefficients from $(Z)_m$. Thus $(Z)_m$ is not big, for any $m \in \mathbb{N}$; but $(Z) = G$ is big.

7.5.3 Uniformly big subfields

We strengthen the notion of big, to take into account elementary extensions. Accordingly, we now suppose that Z is definable; though (Z) need not be definable (in any model of G). For an elementary extension $G \preceq G^*$, we let Z^* denote the set defined in G^* by the formula that defines Z in G.

Definition 7.5.13. If (Z^*) is big in G^* for each elementary extension $G \preceq G^*$ then we say that (Z) is *uniformly* big in (the theory of) G.

In the rest of this section, realisation of types and uniform properties should always be thought of in the theory of G.

Lemma 7.5.14. *If (Z) is uniformly big then $tr_{(Z)}$ is not realisable.*

Proof. If (Z) is uniformly big, then (Z^*) is big in G^*, for each elementary extension $G \preceq G^*$. Big subfields are, in particular, co-algebraic. Thus $G^*/(Z^*)$ is algebraic, for each $G \preceq G^*$. Certainly the type $tr_{(Z)}$ is not realisable. □

Lemma 7.5.15. *Suppose that $(Z)_n$ is not big, for any $n \in \mathbb{N}$. Then $tr_{(Z)}$ is realisable.*

Proof. For each $m, n \in \mathbb{N}$, we have that $G \models \exists x\ \delta^n_{(Z)_m}$. By Lemma 7.4.7, the formulas $\delta^n_{(Z)_m}$ are finitely consistent. Thus the type $tr_{(Z)} = \{\delta^n_{(Z)_m} | m, n \in \mathbb{N}\}$ is realisable. Thus, in some elementary extension $G \preceq G^*$, there exists an element transcendental over (Z^*). This is a contradiction since big subfields are, in particular, co-algebraic; i.e. $G^*/(Z^*)$ is algebraic. □

Proposition 7.5.16. *The following are equivalent.*

1. *(Z) is uniformly big;*
2. *$(Z)_n$ is big, for some $n \in \mathbb{N}$; and*
3. *$tr_{(Z)}$ is not realisable.*

Proof. Lemma 7.5.14 and Lemma 7.5.15 provide (1) $\implies$ (3) $\implies$ (2).

It remains to show that (2) $\implies$ (1). Suppose that $(Z)_n$ is big in G and let $G \preceq G^*$ be an elementary extension. Bigness of definable subsets is an elementary property. Thus $(Z^*)_n$ is big in G^*. By Lemma 7.5.11, (Z^*) is big in G^*. Thus (Z) is uniformly big, as required. □

Thus, if $tr_{(Z)}$ is not realisable then (Z) is not uniformly big. We now investigate the weaker assumption that $tr^n_{(Z)}$ is not realisable, for some $n \in \mathbb{N}$.

Lemma 7.5.17. *Suppose that $tr_{(Z)}$ is realisable. Then $tr^n_{(Z)}$ is realisable, for each $n \in \mathbb{N}$.*

Proof. We proceed by induction: the assumption forms the base case. Suppose that $tr^n_{(Z)}$ is realisable in the theory of G. Let $G \preceq G^*$ be an elementary extension in which there exists an n-tuple $\mathbf{a}$ such that $\mathbf{a}$ is algebraically independent over (Z^*). Let $M \in \mathbb{N}$ and choose $m \in \mathbb{N}$ so that $p^m > M$. Observe that $\mathbf{a}^{p^m}$ is algebraically independent over (Z^*). For any $a \in \mathbf{a}$, a is purely inseparable of degree $p^m \geq M$ over $(Z^*, \mathbf{a}^{p^m})$. Consequently a is of degree at least p^m over the set $(Z^*, \mathbf{a}^{p^m})_M$. Thus $a\mathbf{a}^{p^m}$ is an $x\mathbf{y}$-tuple which satisfies the partial type $tr^n_{(Z)}(\mathbf{y}) \cup \{\delta^M_{(Z,\mathbf{y})_M}(x)\}$. Therefore $tr^{n+1}_{(Z)}$ is finitely consistent, and so consistent, with the theory of G. □

Definition 7.5.18. We say that $\mathrm{trdeg}(G/(Z))$ is of *uniformly finite* (in the theory of G) if there exists $n \in \mathbb{N}$ such that $\mathrm{trdeg}(G^*/(Z^*)) \leq n$, for all elementary extensions $G \preceq G^*$.

Proposition 7.5.19. *Suppose that* $\mathrm{trdeg}(G/(Z))$ *is uniformly finite. Then (Z) is uniformly big.*

Proof. Suppose that $tr_{(Z)}$ is realisable. By Lemma 7.5.17, $tr^n_{(Z)}$ is realisable for each $n \in \mathbb{N}$. This contradicts the uniform finiteness of $\mathrm{trdeg}(G/(Z))$. Thus $tr_{(Z)}$ is not realisable. By Proposition 7.5.16, (Z) is uniformly big. □

Note that since (Z) need not be definable, '(Z) is big in G' does not seem *a priori* to be an elementary property.

7.6 ∃-definable subsets are big

We turn our attention to infinite existentially definable subsets X of K. In this section we show that $\mathrm{trdeg}(K/(X))$ is uniformly finite. By Proposition 7.5.19 we can conclude that (X) is uniformly big in K. In later sections we will thus be able to apply Proposition 7.5.8.

Let $Y \subseteq {}^{x\mathbf{y}}K$ be a C-definable algebraic set and let $X := \mathrm{pr}_x(Y)$ be the projection of Y onto the

x-coordinate. Suppose that X is infinite. Fix an elementary extension $K \preceq K^*$ in which we realise the type $tr_D \cup \{x \in X\}$: let $a \in K^*$ be such that $a \in X$ and a is transcendental over D. Let $\mathbf{b} \in {}^{\mathbf{y}}K^*$ be such that $a\mathbf{b} \in Y$. Also set $F := D(a\mathbf{b})$.

We apply the global version of Theorem 7.1.1 to $a\mathbf{b}$ over D: there exists $\boldsymbol{\lambda}\alpha, \boldsymbol{\beta} \in {}^{\boldsymbol{\lambda}x,\mathbf{y}}\tau$ such that X contains the set $\Sigma(\mathrm{locus}(\boldsymbol{\lambda}a/D) \cap \mathbf{B}(\boldsymbol{\lambda}\alpha; \boldsymbol{\lambda}a))$.

Recall that $\boldsymbol{\lambda}a$ has a separating partition $a_p^n a_s^n$, for some $n \in \mathbb{N}$.

Lemma 7.6.1. *Let $a_p^n = \mathbf{e}e$ be any partition of a_p^n into a singleton e and a tuple $\mathbf{e}$. (Any partition will do: our convention is that tuples are unordered.) Then a is transcendental over $D(\mathbf{e})$.*

Proof. Since $a_p^n = \mathbf{e}e$ is a p-base of F over D, $D(\mathbf{e})$ is Λ-closed in F and $\mathbf{e}e$ is algebraically independent over D.

Let E be the relative algebraic closure of $D(\mathbf{e})$ in F. Then $E/D(\mathbf{e})$ is separably algebraic and F/E is Λ-closed. Suppose that a is algebraic over E. Then $a \in E$ and so E contains the Λ-closure of $D(a)$ in F, i.e. $\boldsymbol{\Lambda}(F/D(a)) \subseteq E$. But $e \in \boldsymbol{\Lambda}(F/D(a))$ is transcendental over $D(\mathbf{e})$, which is a contradiction. Therefore a is transcendental over E. □

Fact 7.6.2. Let M/L be an extension of transcendence degree one and suppose that $a \in M$ is transcendental over L. Then M is algebraic over $L(a)$.

Proof. This is the exchange lemma for algebraic dependence. □

Lemma 7.6.3. *Let $a_p^n = \mathbf{e}e$ be any partition of a_p^n into a singleton e and a tuple $\mathbf{e}$. Then e is algebraic over $D(\mathbf{e}a)$.*

Proof. Recall that $a^n = a_p^n a_s^n$. The tuple a_s^n is algebraic over $D(a_p^n)$, thus $D(a^n)/D(\mathbf{e})$ has transcendence degree 1. By Lemma 7.6.1, $a \in D(a^n)$ is transcendental over $D(\mathbf{e})$. We apply Fact 7.6.2 to find that $D(a^n)$ is algebraic over $D(\mathbf{e}a)$, as required. □

Now we look more closely at $\mathrm{locus}(\boldsymbol{\lambda}a/D)$ locally around $\boldsymbol{\lambda}a$. We noted about that $\boldsymbol{\lambda}a = a_p^n a_s^n$ and a_s^n is separably algebraic over $D(a_p^n)$. By Proposition 6.2.3 there exists $\boldsymbol{\gamma\delta} \in {}^{\mathbf{vw}}\tau$ such that $\mathrm{locus}(\boldsymbol{\lambda}a/D) \cap \mathbf{B}(\boldsymbol{\gamma\delta}; \boldsymbol{\lambda}a)$ is the graph of a continuous function $\mathbf{f}$ with domain $\mathbf{B}(\boldsymbol{\gamma}; a_p^n)$. (Note that a_p^n is purely transcendental over D so $\mathrm{locus}(a_p^n/D) = {}^{\mathbf{v}}K$.)

First, and as an aside, we prove two quick facts about 'transcendence genericity' (although we will not explore this property any further in this work). Let F/E be any field extension.

Fact 7.6.4. Let $\mathbf{a}b, \mathbf{c}d \in {}^{\mathbf{x}y}F$ be such that $E(\mathbf{a}) \cong E(\mathbf{c})$, that b is algebraic over $E(\mathbf{a})$, and that $\mathbf{c}d \in \mathrm{locus}(\mathbf{a}b/E)$. Then $E(\mathbf{a}b) \cong E(\mathbf{c}d)$.

Proof. Let $m \in E(\mathbf{x})[y]$ be such that $m(\mathbf{a}y)$ is the minimal polynomial of b over $E(\mathbf{a})$. We may write $m = m'/m''$ where $m' \in E[\mathbf{x}y]$ and $m'' \in E[\mathbf{x}]$. Note that $m'(\mathbf{a}b) = 0$ and $m''(\mathbf{a}) \neq 0$. Since $E(\mathbf{a}) \cong E(\mathbf{c})$, we have that $m''(\mathbf{c}) \neq 0$. Since $\mathbf{c}d \in \mathrm{locus}(\mathbf{a}b/E)$, we have that $m'(\mathbf{c}d) = 0$. Thus $m(\mathbf{c}d) = 0$ and we may extend the isomorphism so that $E(\mathbf{a}b) \cong E(\mathbf{c}d)$. □

Fact 7.6.5. Let $\mathbf{a}, \mathbf{c} \in {}^{\mathbf{x}}F$ be such that $\mathrm{trdeg}(\mathbf{a}/E) \leq \mathrm{trdeg}(\mathbf{c}/E)$ and suppose that $\mathbf{c} \in \mathrm{locus}(\mathbf{a}/E)$. Then $E(\mathbf{a}) \cong E(\mathbf{c})$.

Proof. Let $\mathbf{a}' \subseteq \mathbf{a}$ be subtuple of $\mathbf{a}$ which is also a transcendence base of $\mathbf{a}$ over E. We may choose an algebraically independent subtuple $\mathbf{c}' \subseteq \mathbf{c}$ of the same cardinality as $\mathbf{a}'$. Thus $\mathbf{a}'$ and $\mathbf{c}'$ generate isomorphic extensions of E. We proceed by induction on the length of $\mathbf{a} \backslash \mathbf{a}'$; then we use Fact 7.6.4. □

Now we return to main argument.

Lemma 7.6.6. *Let $\mathbf{e} \in \mathbf{B}(\boldsymbol{\gamma}; a_p^n) \subseteq {}^{\mathbf{v}}K^*$ be algebraically independent over D. Then $D(a_p^n a_s^n) \cong D(\mathbf{e}, \mathbf{f}(\mathbf{e}))$ is an isomorphism over D.*

Proof. Note that a_p^n and $\mathbf{e}$ are both algebraically independent and that a_s^n is (separably) algebraic over $D(a_p^n)$. Thus $\mathrm{trdeg}(a_p^n a_s^n/D) \leq \mathrm{trdeg}(\mathbf{e}, \mathbf{f}(\mathbf{e})/D)$. By definition of $\mathbf{f}$ (which is well-defined on $\mathbf{e}$ by our choice of $\boldsymbol{\gamma}$), we have that $(\mathbf{e}, \mathbf{f}(\mathbf{e})) \in \mathrm{locus}(a_p^n a_s^n/D)$. We apply Fact 7.6.5 to find that $D(a_p^n a_s^n) \cong D(\mathbf{e}, \mathbf{f}(\mathbf{e}))$, as required. □

Note that in the next lemma, $\mathbf{e}$ is no longer a $\mathbf{v}$-tuple.

Lemma 7.6.7. *Let $\mathbf{e}e \in \mathbf{B}(\boldsymbol{\gamma}; a_p^n) \subseteq {}^{\mathbf{v}}K^*$ and let $\sigma := \Sigma(\mathbf{e}e, \mathbf{f}(\mathbf{e}e))$. Then e is algebraic over $D(\mathbf{e}\sigma)$.*

Proof. The isomorphism from Lemma 7.6.6 sends $a \longmapsto \sigma$ because Σ is a polynomial and $a = \Sigma(a_p^n a_s^n)$. Thus we may apply Lemma 7.6.3. □

Let $\mathbf{d}$ be a transcendence base of D and let $\mathbf{e}e = a_p^n$ be any (unordered) partition. Note that $\mathbf{e}$ is algebraically independent over D.

Proposition 7.6.8. *K^* is algebraic over $(X^* \cup \mathbf{de})$.*

Proof. Let T be a transcendence base of $K^*/(\mathbf{de})$ such that $\{\mathbf{e}\} \times T \subseteq \mathbf{B}(\boldsymbol{\gamma}; a_p^n)$. (We may always choose a transcendence base in any ball.) Let $t \in T$. By Lemma 7.6.7, t is algebraic over $D(\mathbf{e}\sigma)$. Since $\sigma \in X^*$, t is algebraic over $(X \cup \mathbf{de})$. □

Thus $\mathrm{trdeg}(K^*/X^*) \leq |\mathbf{de}|$.

Remark 7.6.9. The tuple $\mathbf{e}$ is of finite length, but $\mathbf{d}$ could be infinite. We will use the fact that Proposition 7.6.8 and this observation are true in all elementary extensions of K. This can be seen by applying Proposition 7.6.8 in K^* to see that K^* is algebraic over $X^* \cup \mathbf{de}$.

Proposition 7.6.10. $\mathrm{trdeg}(K/(X))$ *is uniformly finite.*

Proof. The field D needs to contain the finitely many parameters from the definition of X and the coefficients of the finitely many separable minimal polynomials of a_s^n over $D(a_p^n)$ and of $\mathbf{b}$ over $D(a^n)$. Thus we may assume that the transcendence degree of D is finite! By Proposition 7.6.8, $\mathrm{trdeg}(K^*/X^*) \leq |\mathbf{de}| < \infty$. This result will also hold for any other elementary extension of K which realises $tr_D \cup \{x \in X\}$ since this was our only assumption on K^*. It is not yet clear that this holds for every elementary extension of K.

Let $K \preceq K'$ be an elementary extension (possibly trivial) and suppose that $\mathrm{trdeg}(K'/(X')) > |\mathbf{de}|$. Let $n := |\mathbf{de}|+1$. Then K' realises $tr^n_{(X)}$; so any elementary extension $K'^* \succeq K'$ also realises $tr^n_{(X)}$. But as discussed above, we may choose an elementary extension $K' \preceq K'^*$ which realises $tr_D \cup \{x \in X\}$. Thus we may apply Proposition 7.6.8 in K'^* to find that $\mathrm{trdeg}(K'^*/(X'^*)) \leq |\mathbf{de}| = n$, which is a contradiction.

Thus, for any elementary extension $K \preceq K^*$, we have that $\mathrm{trdeg}(K^*/(X^*)) \leq |\mathbf{de}| < \infty$. This proves that $\mathrm{trdeg}(K/(X))$ is uniformly finite, as required. □

Corollary 7.6.11. (X) *is uniformly big in* K.

Proof. We apply Proposition 7.5.19 to Proposition 7.6.10. □

7.7 Aside: special cases

We deal with two special cases. These special cases will be very important when we come to study ∃-definable substructures.

7.7.1 K^{p^∞}-points

The first case we need to understand is when X contains a point a in the maximal perfect subfield $K^{p^\infty} := \bigcap_{n<\omega} K^{p^n}$ of K which is transcendental over D. Let $a^{p^{-\infty}}$ denote the set $\{a^{p^{-n}} | n \in \mathbb{N}\}$.

Proposition 7.7.1. *Let F/E be any separable extension of fields and let $a \in F^{p^\infty}$. Then $E(a^{p^{-\infty}})$ is Λ-closed in F.*

Proof. Let $E' := E(a^{p^{-\infty}})$. Observe that $E'^p = E^p(a^{p^{-\infty}})$. Thus $E' = E'^p E$ and E p-spans E'. Let $\mathbf{a} \in pB(E)$ be a p-base of E. Since F/E is separable, so is E'/E. Thus $\mathbf{a} \in pI(E')$. Since E p-spans E', $\mathbf{a} \in pB(E')$. Again since F/E is separable, $\mathbf{a} \in pI(F)$. Thus a p-base of E' is p-independent in F. By Proposition 2.2.8, E' is Λ-closed in F. □

Proposition 7.7.2. *Let $X \subseteq K$ be a $\exists$-C-definable subset and let $a \in X \cap K^{p^\infty}$ be transcendental over D. Then there exists $n \in \mathbb{N}$ and $\alpha \in \tau$ such that $B^n(\alpha; a) \subseteq X$.*

Proof. We apply the global method. By Proposition 7.7.1, $D(a^{p^{-\infty}}) \subseteq K$ is Λ-closed in K. Thus (since we are applying the global method) $\boldsymbol{\lambda} a = a^{p^{-n}}$, for some n. Thus $\Sigma = x^{p^n}$. Since a is transcendental over D, $\mathrm{locus}(a^{p^{-n}}/D) = K$. By Theorem 7.1.1, there exists $\alpha \in {}^v\tau$ such that $\Sigma(\mathrm{locus}(\boldsymbol{\lambda} a/D) \cap B(\alpha; \boldsymbol{\lambda} a)) = B^n(\alpha; a) \subseteq X$. □

7.7.2 A 'nice' field of constants: $C \subseteq K^{p^\infty}$

Another special case is when the field C generated by the parameters in the definition is contained in the maximal perfect subfield K^{p^∞} of K.

Proposition 7.7.3. *Suppose that $C \subseteq K^{p^\infty}$. Let $X \subseteq K$ be a $\exists$-C-definable subset and let $a \in X \cap (K \setminus K^p)$ be transcendental over C. Then there exists $\alpha \in \tau$ such that $B(\alpha; a) \subseteq X$.*

Proof. Since $C \subseteq K^{p^\infty}$, $D = C^{p^{-\infty}}$. Since $a \notin K^p$, $D(a)$ is Λ-closed in K. By Theorem 7.1.1, there exists $\alpha \in \tau$ such that $B(\alpha; a) \subseteq X$. □

7.8 $\exists$-definable substructures

We now prove that there are not very many existentially definable subrings and subfields. Although the main result of this section will be superseded in the next section, the proofs are simpler in this special case.

7.8.1 Subrings

Recall the definition of balls

Theorem 7.8.1. *Let $R \subseteq K$ be an infinite $\exists$-C-definable subring. There exists $n \in \mathbb{N}$, $a \in R$, and $\alpha \in \tau$ such that $B^n(\alpha; a) \subseteq R$.*

Proof. Since R is infinite, we may find an elementary extension $K \preceq K^*$ in which there is an element $a^* \in R^* \cap K^{*p^\infty}$ and which is transcendental over D. We apply Proposition 7.7.2 to find $n \in \mathbb{N}$ and $\alpha^* \in \tau^*$ such that $B^n(\alpha^*; a^*) \subseteq R^*$. This is an elementary statement in the language of filtered fields: thus, for the same natural number n, there exists $a \in K$ and $\alpha \in \tau$ such that $B^n(\alpha; a) \subseteq R$. □

7.8.2 Subfields

First we see that balls generate the whole field.

Proposition 7.8.2. *Let $\alpha \in \tau$ and let $(\alpha)_2$ denote the second stage in the generation of a field, as above. Then $F = (\alpha)_2$.*

Proof. Let $a \in F$. There exists $\beta \in \tau$ such that $a\beta \subseteq \alpha$. We might assume that $\beta \subseteq \alpha$. Choose any $b \in \beta$ and observe that $a = ab/b \in \alpha/\alpha$. □

Theorem 7.8.3. *Let $F \subseteq K$ be an infinite $\exists$-C-definable subfield. There exists $n \in \mathbb{N}$ such that $K^{p^n} \subseteq F$. If $C \subseteq K^{p^\infty}$ or if $\mathrm{impdeg}(K) = 1$ then we may choose n such that $F = K^{p^n}$.*

Proof. Subfields are *a fortiori* subrings. By Theorem 6, there exists $n \in \mathbb{N}$, $a \in F$, and $\alpha \in \tau$ such that $B^n(\alpha; a) \subseteq F$. Since balls generate fields we have that $K^{p^n} \subseteq F$.

We turn to the first special case: suppose that $C \subseteq K^{p^\infty}$. Let $n \in \mathbb{N}$ be the least such that $F \cap (K^{p^n} \setminus K^{p^{n+1}}) \neq \emptyset$; then $F \subseteq K^{p^n}$. Let $a \in F \cap (K^{p^n} \setminus K^{p^{n+1}})$. Since K/K^{p^∞} is regular, a is transcendental over D. Thus $a^{p^{-n}}$ is p-independent over D and, by Proposition 7.7.3, there exists $\alpha \in \tau$ such that $B(\alpha; a) \subseteq F^{p^{-n}}$. Since balls generate fields, we must have that $K = F^{p^{-n}}$. Thus $K^{p^n} = F$, as required.

In the second special case we assume that $\mathrm{impdeg}(K) = 1$, and in this case the only proper subfield of K^{p^n} which contains $K^{p^{n+1}}$ is $K^{p^{n+1}}$. □

7.9 Subfields generated by ∃-definable subsets

We want to study the subfields of K that are generated by an infinite existentially C-definable subset X.

Recall that $(X)_n$ denotes the n-th stage of the field generated by X.

7.9.1 Finding a K^{p^∞}-point

Proposition 7.9.1. *There exist $m \in \mathbb{N}$, an elementary extension $K \preceq K^*$, and $a \in (X^*)_m \cap K^{*p^\infty}$ which is transcendental over D.*

Proof. Let $K \preceq K^*$ be an elementary extension such that $\mathrm{trdeg}(K^{*p^\infty}) > \mathrm{trdeg}(D)$. By Corollary 7.6.11, (X) is uniformly big in K; thus (X^*) is big in K^*. By Proposition 7.5.8, $K^{*p^\infty}/(K^{*p^\infty} \cap (X^*))$ is a finite separable (thus algebraic) extension. Therefore $\mathrm{trdeg}(K^{*p^\infty} \cap (X^*)) = \mathrm{trdeg}(K^{*p^\infty}) > \mathrm{trdeg}(D)$. So there exists $a \in K^{*p^\infty} \cap (X^*)$ which is transcendental over D. □

7.9.2 Subfields

Proposition 7.9.2. *There exist $m, n \in \mathbb{N}$ and $\alpha \in \tau$ such that $B^n(\alpha; 0) \subseteq (X)_m$.*

Proof. Since $a^* \in (X^*)$, there exists $m \in \mathbb{N}$ such that $a^* \in (X^*)_m$. Note that $(X^*)_m$ is an $\exists$-C-definable subset of K. Combining Proposition 7.7.2 with Proposition 7.9.1, there exists $n \in \mathbb{N}$ and $\alpha^* \in \tau^*$ such that $B^n(\alpha^*; a^*) \subseteq (X^*)_m$; this is an elementary property so there exists $a \in (X)_m$ and $\alpha \in \tau$ such that $B^n(\alpha; a) \subseteq (X)_m$. Simply by subtracting a, we have that $B^n(\alpha; 0) \subseteq (X)_m$, as required. □

Theorem 7.9.3. *Let $X \subseteq K$ be an infinite $\exists$-C-definable subset. There exist $m, n \in \mathbb{N}$ such that $K^{p^n} \subseteq (X)_m$.*

Proof. The ball $B^n(\alpha; 0)$ generates the field K^{p^n} and, by Proposition 7.9.2, $B^n(\alpha; 0) \subseteq (X)_m$. By definition of the subsets $(X)_i$ (for $i \in \mathbb{N}$), $K^{p^n} \subseteq (X)_{m+2} \subseteq (X)$. □

Chapter 8

Existentially closed embeddings

We apply the results from previous chapters to address questions of existential closedness of extensions.

8.1 Dense embeddings

Theorem 8.1.1. *Let F/K be a regular extension, that F is t-henselian, and that K is dense in F. Then $K \preceq_\exists F$.*

Note that F/K is regular if and only if K is relatively algebraically closed and relatively inseparably closed in F.

Proof. We are required to show that existential sentences with parameters from K that are modelled by F are also modelled by K. Let $\phi \in \mathcal{L}_{\text{ring}}$ be an existential sentence in the language of rings with parameters from K. By the remark in section 1.2, there exists a polynomial $f \in K[\mathbf{x}]$ such that ϕ is equivalent to $\exists \mathbf{x}\, f(\mathbf{x}) \dot{=} 0$ in all fields and under all interpretations of the parameters.

Suppose that $F \models \phi$. Then there exists $\mathbf{a} \in {}^{\mathbf{x}}F$ such that $F \models f(\mathbf{a}) \dot{=} 0$. We must find $\mathbf{a}' \in {}^{\mathbf{x}}K$ such that $K \models f(\mathbf{a}') \dot{=} 0$. The extension F/K is Λ-closed; thus $\mathbf{a}$ is separable over K. Let $\mathbf{a}_i\mathbf{a}_s$ be a separating partition of $\mathbf{a}$ over K. By Proposition 6.2.5 (applied in L), there exist open sets $\boldsymbol{\alpha}_i\boldsymbol{\alpha}_s \in {}^{\mathbf{v}_i\mathbf{v}_s}\tau$ such that $\operatorname{locus}(\mathbf{a}/K) \cap \mathbf{B}(\boldsymbol{\alpha}_i\boldsymbol{\alpha}_s; \mathbf{a}_i\mathbf{a}_s)$ is the graph of a continuous algebraic function $\mathbf{f} : \mathbf{B}(\boldsymbol{\alpha}_i; \mathbf{a}_i) \longrightarrow \mathbf{B}(\boldsymbol{\alpha}_s; \mathbf{a}_s)$.

We make sure that we choose $\boldsymbol{\alpha}_i$ to be small enough that $\mathbf{B}(\boldsymbol{\alpha}_i; \mathbf{a}_i)$ is disjoint from the zero-sets of the leading coefficients of the polynomials which algebraise $\mathbf{a}_s$ over $K(\mathbf{a}_i)$.

Since K is dense in F, we may choose a tuple $\mathbf{a}'_i \in {}^{\mathbf{v}_i}K \cap \mathbf{B}(\boldsymbol{\alpha}_i; \mathbf{a}_i)$; then $\mathbf{f}(\mathbf{a}'_i) \in \mathbf{B}(\boldsymbol{\alpha}_s; \mathbf{a}_s) \subseteq {}^{\mathbf{v}_s}F$. Let $\mathbf{a}' := (\mathbf{a}'_i, \mathbf{f}(\mathbf{a}'_i))$. We have that $\mathbf{a}' \in \text{locus}(\mathbf{a}/K)$. Since f is a polynomial over K, we have that $f(\mathbf{a}') = 0$. Furthermore, we note that (by our choice of $\boldsymbol{\alpha}_i$) $\mathbf{f}(\mathbf{a}'_i)$ is algebraic over $K(\mathbf{a}'_i) \subseteq K$. Therefore $\mathbf{a}'$ is algebraic over K, and is thus contained in K.

This proves that $K \models \exists \mathbf{x}\ f \doteq 0$; thus $K \models \phi$, and so $K \preceq_\exists F$. □

As a corollary, we deduce a well-known result of Ershov.

Corollary 8.1.2. $\mathbb{F}_p(t)^h \preceq_\exists \mathbb{F}_p((t))$.

Proof. $\mathbb{F}_p((t))$ is the completion of $\mathbb{F}_p(t)^h$, so the extension is dense. It is well known that this extension is relatively algebraically closed (alternatively, use the fact from [25] that a t-henselian field is relatively separably algebraically closed in its completion). Since t is a p-base of both fields, the extension is separable. Therefore the extension is regular and the hypotheses of Theorem 8.1.1 are satisfied. □

8.2 Extensions of $\mathbb{F}_q((t))$

Using ideas similar to those in the previous section, we investigate extensions of t-henselian fields which are not necessarily dense. We consider the valued field of $\mathbb{F}_q((t))$ with the t-adic valuation. This valuation is discrete, the value group is $\mathbb{Z}$, and the residue field is $\mathbb{F}_q$.

Proposition 8.2.1. *Let $K/\mathbb{F}_q((t))$ be an extension of valued fields such that t is a uniformiser in K and $\mathbb{F}_q$ is the residue field of K. Let $\mathcal{O}_K$ be the valuation ring of K. Let $\pi : \mathcal{O}_K \longrightarrow \mathbb{F}_q[[t]]$ be the 'standard part' specialisation. Then π is a ring homomorphism and $\mathbb{F}_q[[t]] \preceq_{\exists+} \mathcal{O}_K$.*

The standard part map is well-defined because the valuation on $\mathbb{F}_q((t))$ is complete under a valuation which is discrete of rank 1.

Proof. Denote $\mathcal{O} := \mathbb{F}_q[[t]]$. Note that vK has a convex subgroup isomorphic to $\mathbb{Z}$. the '$\mathcal{O}$-infinitesimal neighbourhood' of 0 in $\mathcal{O}_K$ is in fact an ideal. The map π is just the quotient map $\mathcal{O}_K \longrightarrow \mathbb{F}_q[[t]]$.

Let $f \in \mathbb{F}_q[[t]][\mathbf{x}]$ and let $\mathbf{a} \in {}^{\mathbf{x}}\mathcal{O}_K$ be such that $\mathcal{O}_K \models f(\mathbf{a}) \doteq 0$. Applying π, we have that $\mathcal{O} \models f(\pi\mathbf{a}) \doteq 0$, as required. □

Let $\mathbf{T}$ be the elementary theory of $\mathbb{F}_q((t))$ in the language of rings with the constant t, which we denote $\mathcal{L}_{\text{ring}}(t)$.

Proposition 8.2.2. *Let $K/\mathbb{F}_q((t))$ be an extension of models $\mathbf{T}$. Then $\mathbb{F}_q[[t]] \preceq_{\exists+} \mathcal{O}_K$.*

Note that when we speak of *the* valuation ring of K we really mean the one defined by the same formulas (with no parameters) which defines $\mathbb{F}_q[[t]]$ in $\mathbb{F}_q((t))$.

Proof. First we observe that $K/\mathbb{F}_q((t))$ is an extension of valued fields, t is a uniformiser in K, and $\mathbb{F}_q$ is the residue field of K. We apply Proposition 8.2.1. □

Part III

F-definability in $F((t))$

Introduction to Part III

We study definability in the power series field $F((t))$ in the language of rings allowing parameters from F. We assume that F is perfect. Our method is to analyse the orbits of elements of $\mathcal{M} = tF[[t]]$ under F-automorphisms of $F((t))$. By old work of Schilling, [27], these automorphisms are given by composition with power series of value 1. These orbits have a very simple description, and we exploit this to give several properties of F-definable subsets of F. After proving a 'Hensel-like Lemma', we show how to choose, for each $b \in \mathcal{M} \setminus \{0\}$, a polynomial $f_b \in xF[x]$ such that $b = f(a)$ for some a of value 1. Thus b is interdefinable (over F) with the element a. Let $\mathcal{P}$ be the set of elements of value 1. Likewise, every element of $F((t)) \setminus F$ is interdefinable (over F) with an element of $\mathcal{P}$. Moreoever, the interdefinability is witnessed by a polynomial. Using Schilling's representation of F-automorphisms as substitution by elements of $\mathcal{P}$, we see that the orbit of $b \in \mathcal{M} \setminus \{0\}$ is $f_b(\mathcal{P})$, which is an infinite $\exists$-t-definable subset of $F((t))$. Thus any F-definable subset not contained in F contains an infinite $\exists$-t-definable subset. Combining this with our earlier work on existential definability in t-henselian fields, we are able to draw several conclusions about F-definable subsets of $F((t))$. For example, we prove that the only subfields of $F((t))$ which are F-definable are those of the form $F((t^{p^n}))$, for some $n \in \mathbb{N}$.

Of course, any expansion of $F((t))$ which has the same F-automorphisms will have the same theory of F-definability. In [27], Schilling proved that all ring automorphisms of $F((t))$ are continuous, that is are already valuation-preserving. Thus all the above results apply also to F-definability in the expansion of $\mathcal{L}_{\text{ring}}$ by a predicate for the valuation ring.

In Appendix A we will use the existential definition of the t-henselian topology in $F((t))$ from Prestel and Ziegler's paper, [25], and the above analysis of F-orbits of elements of $\mathcal{M} \setminus \{0\}$ to find an $\exists$-$\emptyset$-definition of the valuation ring in the field $\mathbb{F}_q((t))$ of formal power series over a finite field.

Chapter 9

A lemma like Hensel's

Let $K := F((t))$ be a power series field with perfect constant field F of characteristic exponent p. We begin our study of F-definability in K by proving a Hensel-like lemma for K (Theorem 9.4.2). In chapter 10, we will use this to prove that every element of K is interdefinable (over F) with an element of $\mathcal{P}$, the set of elements of value 1, and then to derive consequences for F-definability in K. The elements of $\mathcal{P}$ are important because, as we show in section 9.1, there is only one 1-type over F of elements of value 1 .

Recall that $F((t)) := \{\sum_{i \geq n} a_i t^i | a_i \in F, n \in \mathbb{Z}\}$ along with the t-adic valuation $\sum a_i t^i \longmapsto \min\{i \in \mathbb{Z} | a_i \neq 0\}$ is an henselian valued field. Let $\mathcal{O} := F[[t]]$ be the valuation ring, and let $\mathcal{M} := tF[[t]]$ be the maximal ideal of $\mathcal{O}$. Let $\mathcal{P} := \{a \in K | va = 1\}$ be the set of elements of value 1; these are the *uniformisers* of K.

9.0.1 The goal of 'A lemma like Hensel's'

The theme of this chapter is solving 'substitution equations' of power series. More precisely, our aim is to study how approximate solutions to the equation $f(y) = b$ may be refined to exact solutions, in a similar way to Hensel's Lemma, for a polynomial $f \in xF[x]$ and $b \in \mathcal{M}$.

Definition 9.0.3. We define the *p-th root depth* of $a \in F((t))$ to be $\text{dep}_p(a) := \max\{l \in \omega \cup \{\infty\} | a \in F^{p^l}\}$.

For each $f \in xF[x]$, we associate two natural numbers H_f and H_f^* such that the following Hensel-like lemma holds.

Proposition 9.0.4. *(*Theorem 9.4.2*) Let $f \in xF[x]$. Let $a' \in \mathcal{P}$ and $b \in \mathcal{M}$ be such that $f(a') \in b + \mathcal{M}^{H_f}$ and $\mathrm{dep}_p(b) \geq \mathrm{dep}_p(f)$. Then there exists (a unique) $a \in a' + \mathcal{M}^{H_f^*+1}$ such that $f(a) = b$.*

We will define the function dep_p more precisely later on; roughly speaking it counts how many p^n-th roots an element has in the field $F((t))$. We find H_f, which we call the *Hensel degree* of f, by studying how the coefficients of y contribute to the coefficients of $f(y)$

Satisfying Theorem 9.4.2 is not the only property we require of the Hensel degree. We must have enough control over our choice of H_f that, given $b \in \mathcal{M}$, we can choose $f \in xF[x]$ and $a' \in \mathcal{P}$ such that $f(a') \in b + \mathcal{M}^{H_f}$; thus satisfying the hypothesis of Theorem 9.4.2. This last step we explore in chapter 10.

9.0.2 Basic facts, conventions, and notation

Let a, b, c, y denote elements of K. We will often write $a = \sum_{i=va}^{\infty} a_i t^i$ and $b = \sum_{i=vb}^{\infty} b_i t^i$ without mentioning it further. For a power series a, we let $S_a := \{i \in \mathbb{Z} | a_i \neq 0\}$ be the *support* of a. By definition of $F((t))$, S_a is a well-ordered subset of $\mathbb{Z}$ (using the convention that the empty set is well-ordered).

Fact 9.0.5. Let $a, b \in \mathcal{M}$. Then $a = b$ if and only if $a_i = b_i$, for all $i \in \mathbb{N}$.

Proof. This is immediate from the definition of the field of *formal* power series. □

For $n \in \mathbb{Z}$ and $a \in F((t))$, we let $a_i^{(n)}$ denote the i-th coefficient of a^n so that

$$a^n = \sum_{i=nva}^{\infty} a_i^{(n)} t^i.$$

We denote by $P(i; {}^n S_a)$ the set of all *ordered n-partitions* of i from the set S_a, i.e. n-tuples from the set S_a which sum to i.

Remark 9.0.6. Note that we are concerned with *ordered* partitions of k, i.e. combinatorial not number theoretic partitions. For example $(1, 2)$ and $(2, 1)$ are two different ordered partitions of 3.

Fact 9.0.7. Let $n \in \mathbb{N}$ and $i \in \mathbb{Z}$. Then $a_i^{(n)} = \sum \left\{ \prod_{j \in \mathbf{j}} a_j | \mathbf{j} \in P(i; {}^n S_a) \right\} \in \mathbb{F}[a_j | j \leq i]$. This observation is proved directly by using the definition of multiplication in $F((t))$.

This fact shows that the greatest index of a coefficient appearing in $a_i^{(n)}$ is less than or equal to i. The lemma only applies to coefficients $a_i^{(n)}$ of powers; but refining this result to find the greatest index appearing in the i-th coefficient of $f(a)$ will be crucial to proving Proposition 9.4.1.

Lemma 9.0.8. *Let $n \in \mathbb{N}$ and $i \in \mathbb{Z}$. If $i < nva$ then $a_i^{(n)} = 0$. Thus if $i < n$ and $0 < va$ then $a_i^{(n)} = 0$.*

Proof. We apply Fact 9.0.7 directly: if $i < nva$ then there are no n-element partitions of i from S_a, the support of a. For the second part, we observe that $i < n \leq nva = va^n$ (since $1 \leq va$). Applying the first part, $a_i^{(n)} = 0$. □

We will let $f \in xF[x]$ denote a polynomial over F (often with zero constant term). We write $f = \sum_{n=0}^{N} c_n x^n$. Let $\mathbb{F}_f := [c_n | n \leq N]$ be the *ring* generated by the coefficients of f (over the empty set).

9.1 Useful motivation: automorphisms of fields of formal power series

In this first section we recall the theory of F-automorphisms of the power series field $K = F((t))$. In [27], Schilling proved that F-automorphisms of $F((t))$ correspond to composition (also called substitution) by power series of value 1. Everything in this section can be found in his paper, although some proofs differ slightly in order to motivate methods used later in the chapter.

Although we will not use automorphisms in the other sections of this chapter, the method of 'reverting power series' (in which we recursively define a solution to a power series equation) is useful motivation for the Hensel-like lemma (Theorem 9.4.2), which is the main result in this chapter. In any case, the study of F-automorphisms of K will be crucial in chapter 10.

By an F-automorphism we mean a field (ring) automorphism of K which fixes F pointwise.

9.1.1 Composition of polynomials with power series

Definition 9.1.1. Let $a \in F[[t]]$. By the universal property of the polynomial ring $F[t]$ (as the free ring generated over F on 1 generator) there exists a unique ring homomorphism $S_a : F[t] \longrightarrow F[[t]]$ such that $t \longmapsto a$. We call this map *composition* by a. For $b \in F[t]$, we denote $b \circ a := S_a(b)$.

Proposition 9.1.2. *Let $a \in \mathcal{M}$ and let $b = \sum_{j=vb}^{n} b_j t^j \in F[t]$. Then $S_a(b) = \sum_{i=0}^{\infty} (b \circ a)_i t^i$, where $(b \circ a)_i := \sum_{j=vb}^{i} b_j a_i^{(j)}$.*

Note that each sum $(b \circ a)_i = \sum_{j=vb}^{i} b_j a_i^{(j)}$ is finite, and thus a well-defined element of F.

Proof. Since b is a polynomial, $b(a) = b \circ a$ is a finite sum of power series and we may exchange the order of summation:

$$\begin{aligned} b \circ a = S_a(b) &= S_a \textstyle\sum_{j=vb}^{n} b_j t^j \\ &= \textstyle\sum_{j=vb}^{n} b_j S_a t^j \\ &= \textstyle\sum_{j=vb}^{n} b_j a^j \\ &= \textstyle\sum_{j=vb}^{n} b_j \sum_{i=jva}^{\infty} a_i^{(j)} t^i \\ &= \textstyle\sum_{i=0}^{\infty} \big(\sum_{j=vb}^{n} b_j a_i^{(j)} \big) t^i \\ &= \textstyle\sum_{i=0}^{\infty} \big(\sum_{j=vb}^{i} b_j a_i^{(j)} \big) t^i \\ &= \textstyle\sum_{i=0}^{\infty} (b \circ a)_i t^i, \end{aligned}$$

where the penultimate equality follows from Lemma 9.0.8. □

9.1.2 Composition of power series with power series

Using Proposition 9.1.2, we extend S_a to be a function $F[[t]] \longrightarrow F[[t]]$. Let $(b \circ a)_i$ be defined in exactly the same way as above.

Definition 9.1.3. Let $a \in \mathcal{M}$ and let $b \in F[[t]]$. We define $S_a(b) := \sum_{i=0}^{\infty} (b \circ a)_i t^i$. Then $S_a : F[[t]] \longrightarrow F[[t]]$ is called *composition* by a.

Remark 9.1.4. Note that S_a agrees with our previous definition when restricted to $F[t]$ by Proposition 9.1.2.

Lemma 9.1.5. *Let $a \in \mathcal{M}$ and let $b, b' \in F[[t]]$. If $b' \in b + \mathcal{M}^i$ then $S_a(b') \in S_a(b) + \mathcal{M}^i$.*

Proof. Let $I < i$. By definition, the I-th coefficient of $S_a(b)$ is $(b \circ a)_I = \sum_{j=vb}^{I} b_j a_I^{(j)}$, which clearly only includes the coefficients b_j of b, for $j \leq I < i$. Thus if the j-th coefficients of b and b' are equal, for all $j \leq I$, then the I-th coefficients of $S_a(b)$ and $S_a(b')$ are equal. Therefore, if b and b' are in the same coset of $\mathcal{M}^i$ then their I-th coefficients are equal, for all $I < i$. Consequently, the I-th coefficients of $S_a(b)$ and $S_a(b')$ are equal, for all $I < i$, and $S_a(b') \in S_a(b) + \mathcal{M}^i$. □

Proposition 9.1.6. *For each $a \in \mathcal{M}$, $S_a : F[[t]] \longrightarrow F[[t]]$ is the unique ring homomorphism which fixes F and sends $t \longmapsto a$.*

Proof. First we argue that S_a is a ring homomorphism. Let $b, c \in F[[t]]$. We must show that

$$S_a(b+c) = S_a(b) + S_a(c)$$

and

$$S_a(b \cdot c) = S_a(b) \cdot S_a(c).$$

It suffices to show that these equalities hold up to addition by the ideal $\mathcal{M}^i$, for each $i \in \mathbb{N}$. This is really the same as insisting that the i-th coefficients are equal, for each $i \in \mathbb{N} \cup \{0\}$.

We choose 'truncations' $b', c' \in F[t]$ of b, c such that $b' \in b + \mathcal{M}^i, c' \in c + \mathcal{M}^i$. By Lemma 9.1.5, $S_a(b') \in S_a(b) + \mathcal{M}^i$ and $S_a(c') \in S_a(c) + \mathcal{M}^i$. Since $b' + c' \in b + c + \mathcal{M}^i$ and $b'c' \in bc + \mathcal{M}^i$, another application of Lemma 9.1.5 gives that $S_a(b' + c') \in S_a(b + c) + \mathcal{M}^i$ and $S_a(b' \cdot c') \in S_a(b \cdot c) \in \mathcal{M}^i$.

By definition, S_a is a ring homomorphism when restricted to $F[t]$. Thus

$$\begin{aligned} S_a(b+c) &\in S_a(b'+c') + \mathcal{M}^i \\ &= S_a(b') + S_a(c') + \mathcal{M}^i \\ &= S_a(b) + S_a(c) + \mathcal{M}^i. \end{aligned}$$

and similarly

$$\begin{aligned} S_a(b \cdot c) &\in S_a(b' \cdot c') + \mathcal{M}^i \\ &= S_a(b') \cdot S_a(c') + \mathcal{M}^i \\ &= S_a(b) \cdot S_a(c) + \mathcal{M}^i. \end{aligned}$$

Since this holds for each $i \in \mathbb{N}$, $S_a(b+c) = S_a(b) + S_a(c)$ and $S_a(b \cdot c) = S_a(b) \cdot S_a(c)$. Thus S_a is a ring homomorphism.

The same argument proves that S_a is unique: the restriction of S_a to $F[t]$ is unique by the universal property of $F[t]$ (as mentioned in Definition 9.1.1), and the image of b under S_a is determined by the images of the polynomial truncations of b (i.e. elements $b' \in F[t]$ lying in the same coset of $\mathcal{M}^i$ as b, for successively higher powers $i \in \mathbb{N}$). □

Definition 9.1.7. For each $a \in \mathcal{M}$, let $S_a : F((t)) \longrightarrow F((t))$ be the ring homomorphism which uniquely extends $S_a : F[[t]] \longrightarrow F[[t]]$.

Remark 9.1.8. Such an extension always exists by the universal property of fields of fractions.

9.1.3 The representation S

Definition 9.1.9. We can view $\circ : F((t)) \times \mathcal{P} \longrightarrow F((t))$ as a function $(a, b) \longmapsto a \circ b$. Since $v(b \circ a) = vb \cdot va$, it is clear that we may restrict $\circ$ to a binary operation $\mathcal{M} \times \mathcal{M} \longrightarrow \mathcal{M}$.

Lemma 9.1.10. *S_t is the identity map $F((t)) \longrightarrow F((t))$.*

Proof. By Proposition 9.1.6, the identity map is the unique homomorphism $F[[t]] \longrightarrow F[[t]]$ which fixes F and t. Since S_t also fixes F and t, S_t must be the identity map. □

Lemma 9.1.11. *For all $a \in \mathcal{M}$, $t \circ a = a$.*

Proof. Immediately from the definition of S_a, $t \circ a = S_a(t) = a$. □

Note that we write composition of functions in the usual way 'on the left': thus $f \circ g$ means 'do g first then f.

Lemma 9.1.12. *For all $a, b \in \mathcal{M}$, $S_{b \circ a} = S_a \circ S_b$.*

Proof. Note that $b \circ a \in \mathcal{M}$; thus $S_{b \circ a}$ is well-defined. The composition (of functions) $S_a \circ S_b$ is a homomorphism $F((t)) \longrightarrow F((t))$ which sends

$$t \longmapsto b \longmapsto S_a(b) = b \circ a.$$

Since $S_{b \circ a}$ also sends $t \longmapsto b \circ a$, and by the uniqueness property in Proposition 9.1.6, $S_a \circ S_b = S_{b \circ a}$. □

Lemma 9.1.13. *Let $a, b \in \mathcal{M}$ and let $c \in F((t))$. Then $c \circ (b \circ a) = (c \circ b) \circ a$.*

Proof. By Lemma 9.1.12, $S_{b \circ a} = S_a \circ S_b$. Thus $c \circ (b \circ a) = S_{b \circ a}(c) = S_a(S_b(c)) = (c \circ b) \circ a$. □

Proposition 9.1.14. *$(\mathcal{M}, \circ)$ is a monoid with identity t.*

Proof. By Lemma 9.1.13, $(\mathcal{M}, \circ)$ is associative; and by Lemma 9.1.10 and Lemma 9.1.11 t is an identity element in $(\mathcal{M}, \circ)$. Thus $(\mathcal{M}, \circ)$ is a monoid. □

Let $(E, \circ) := \mathrm{End}(F((t))/F)$ denote the monoid of ring endomorphisms of $F((t))$ which fix F pointwise and let $(E, \circ)^{\mathrm{op}}$ denote the *opposite monoid* of $(E, \circ)$ (see, for example, chapter 7 of [16]).

Proposition 9.1.15. *The map $S : (\mathcal{M}, \circ) \longrightarrow (E, \circ)^{\mathrm{op}}$ which sends $a \longmapsto S_a$ is a monomorphism of monoids.*

Proof. By Lemma 9.1.12, Lemma 9.1.10, and Lemma 9.1.11; S is a homomorphism of monoids. Let $a, b \in \mathcal{M}$ and suppose that $S_a = S_b$. Then $a = S_a(t) = S_b(t) = b$, by definition. Thus S is injective. □

9.1.4 Reversion of power series

We now identify the set of homomorphisms of the form $S_a : F((t)) \longrightarrow F((t))$ which are in fact automorphisms.

Definition 9.1.16. Let $\mathcal{P} := \{a \in F((t)) | va = 1\}$ be the set of uniformisers in $F((t))$.

Note that $(\mathcal{P}, \circ)$ is a submonoid of $(\mathcal{M}, \circ)$.

Lemma 9.1.17. *Let $a \in \mathcal{P}$. Define $b = \sum_{j=1}^{\infty} b_j t^j$ where the coefficients are defined recursively by:*

$$b_1 := 1 \text{ and } b_j := -(a_j^{(j)})^{-1} \sum_{i=1}^{j-1} b_i a_j^{(i)}, \text{ for } j > 1.$$

Then $b \circ a = t$, i.e. b is a left-inverse for a in $(\mathcal{P}, \circ)$.

Note that $b \in \mathcal{P}$. Also note that this recursion makes sense because $a_j^{(j)} = a_1^j \neq 0$.

Proof. Observe that $(b \circ a)_0 = 0$ and $(b \circ a)_1 = 1$. Let $j \in \mathbb{N}$, $j \geq 2$. Then

$$\begin{aligned}
(b \circ a)_j &= \textstyle\sum_{i=1}^{j} b_i a_j^{(i)} \\
&= \textstyle\sum_{i=1}^{j-1} b_i a_j^{(i)} + b_j a_j^{(j)} \\
&= \textstyle\sum_{i=1}^{j-1} b_i a_j^{(i)} + -(a_j^{(j)})^{-1} \sum_{i=1}^{j-1} b_i a_j^{(i)} a_j^{(j)} \\
&= 0.
\end{aligned}$$

□

9.1.5 F-automorphisms of $F((t))$

Proposition 9.1.18. *$(\mathcal{P}, \circ)$ is a group; in fact it is the group of invertible elements of $(\mathcal{M}, \circ)$.*

The proof requires this easy fact from the theory of monoids.

Fact 9.1.19. A monoid in which every element has a left inverse is a group.

Proof. (of proposition) As remarked above, $(\mathcal{P}, \circ)$ is a submonoid of $(\mathcal{M}, \circ)$. By Lemma 9.1.17, each element of $\mathcal{P}$ is left invertible. By Fact 9.1.19, $(\mathcal{P}, \circ)$ is in fact a group. Note that $a \in \mathcal{M} \setminus \mathcal{P}$ implies that a is not invertible since $v(a \circ b) = va \cdot vb > 1$, for each $b \in \mathcal{M}$. □

The following proposition appears as Lemma 1 in Schilling's paper [27].

Proposition 9.1.20. *Each field automorphism of $F((t))$ is already an automorphism of valued fields.*

Proof. Let $\phi : F((t)) \longrightarrow F((t))$ be a field automorphism. Then $\phi(\mathcal{O})$ is a valuation ring of $F((t))$ (corresponding to the $\phi(t)$-adic valuation), with respect to which $F((t))$ is complete. By the famous theorem of F K Schmidt (from [28]), any field which is henselian with respect to two independent valuations must be separably closed. Note that $\mathcal{O} \neq \phi(\mathcal{O})$ implies that $\mathcal{O}$ and $\phi(\mathcal{O})$ are independent. Therefore, and since $F((t))$ is not separably closed, we have that $\phi(\mathcal{O}) = \mathcal{O}$. □

The above argument applies to any field automorphism of $F((t))$. For an alternative proof in the simpler case of F-automorphisms, we could observe that the valuation ring is F-definable (using Ax's definition from [1]). Thus any F-automorphism must preserve $\mathcal{O}$ set-wise.

Remark 9.1.21. Note that we cannot use Theorem A.1.10 for an F-definition of $\mathcal{O}$ because Theorem A.1.10 applies to $K = \mathbb{F}_q((t))$, i.e. $F = \mathbb{F}_q$ is a finite field. This does not matter though; any definition will do, it does not have to be existential.

Let $(A, \circ)$ be the group $\mathrm{Aut}(F((t))/F)$ of field automorphisms of $F((t))$ which fix F pointwise, under composition of functions.

Proposition 9.1.22. *$S : (\mathcal{P}, \circ) \longrightarrow (A, \circ)$ is an isomorphism of groups.*

Proof. From Proposition 9.1.15 and Proposition 9.1.18, we know that (the restriction of) S is an embedding of $(\mathcal{P}, \circ)$ into $(A, \circ)$. The image is contained in $(A, \circ)$ since every element must be invertible. Now let ϕ be any F-automorphism of $F((t))$. By Proposition 9.1.20, $\phi(t) \in \mathcal{M}$. Thus (by the usual uniqueness argument from Proposition 9.1.6) $S_{\phi(t)} = \phi$. Thus ϕ is in the image of $(\mathcal{P}, \circ)$ under S. □

Proposition 9.1.23. *Let $f \in F[t]$. Then* $\mathrm{Orb}(f/F) = f(\mathcal{P})$.

Proof. It is clear that the orbit of f is simply the set of images of f under all automorphisms. By Proposition 9.1.22, this is equal to the image of $\mathcal{P}$ under the polynomial map f. □

9.2 Solving an equation of coefficients

We now turn to our main task of understanding the equation $f(y) = b$. We will think of b as fixed, y as a variable, and a as a solution or approximate solution to the equation $f(y) = b$.

Let $b \in \mathcal{M}$ and let $f = \sum_{n=1}^{N} c_n x^n \in xF[x]$ be a polynomial over F with zero constant coefficient. Recall that $\mathbb{F}_f$ denotes the subring of F generated by the coefficients of f.

9.2.1 The coefficients of $f(y)$

Let $k \in \mathbb{N}$. We set $C_{f,k} := \sum_{n=1}^{N} c_n y_k^{(n)} \in \mathbb{F}_f[y_l | l \leq k]$. We abuse notation a little and write $C_{f,k}(a)$ to mean $C_{f,k}(a_i | i \in \mathbb{N}) = \sum_{n=1}^{N} c_n a_k^{(n)}$, i.e. the evaluation of the polynomial $C_{f,k}$ at the tuple of coefficients of a.

Lemma 9.2.1. *$C_{f,k}$ is the k-th coefficient of $f(y)$.*

Proof. We observe that

$$\begin{aligned} f(y) &= \sum_{n=1}^{N} c_n y^n \\ &= \sum_{n=1}^{N} c_n \left(\sum_{k=1}^{\infty} y_k t^k \right)^n \\ &= \sum_{n=1}^{N} c_n \sum_{k=1}^{\infty} y_k^{(n)} t^k \\ &= \sum_{k=1}^{\infty} \left(\sum_{n=1}^{N} c_n y_k^{(n)} \right) t^k. \end{aligned}$$

In fact, since $C_{f,k} = (f \circ y)_k$ we could have just applied Proposition 9.1.2. □

Lemma 9.2.2. *$f(a) = b$ if and only if $C_{f,k}(a) = b_k$, for each $k \in \mathbb{N}$.*

Proof. Immediate from Lemma 9.2.1 and Fact 9.0.5. □

By Lemma 9.0.8, the coefficient $C_{f,k}$ is a polynomial over $\mathbb{F}_f$ in at most the first k-many of the variables $(y_l | l \in \mathbb{N})$. As remarked above, it is important for us to know precisely what is the greatest index of a variable y_i which appears non-trivially in $C_{f,k}$. This motivates the following definition which we use to label the 'highest index' that occurs.

Definition 9.2.3. Define the function $gi_f : \mathbb{N} \longrightarrow \mathbb{N} \cup \{0\}$ by

$$gi_f(k) := \min\{n \in \mathbb{N} \cup \{0\} | C_{f,k} \in \mathbb{F}_f[y_l | l \leq n]\}.$$

We say that the function gi_f is the *greatest-index* of f, and that $gi_f(k)$ is the *k-th greatest-index* of f.

Remark 9.2.4. Note that (since there is no variable y_0) $\mathbb{F}_f[y_l | l \leq 0] = \mathbb{F}_f$. Thus $gi_f(k) = 0$ if and only if the k-th coefficient of $f(y)$ is a constant polynomial.

Our immediate aim is to compute $gi_f(k)$, for each f and each $k \in \mathbb{N}$.

9.2.2 The greatest-index of x^n

Before we try to understand gi_f in general, we look at the special case where f is a monomial.

Fix $n \in \mathbb{N}$ and consider the monomial x^n. We aim to understand the function gi_{x^n}. Write $n = p^j i$, where $p \nmid i$.

Lemma 9.2.5. *We have*

$$y_k^{(n)} = \begin{cases} y_{kp^{-j}}^{(i)p^j} & \text{if } p^j \mid k, \\ 0 & \text{else.} \end{cases}$$

Proof. $\sum_{k=1}^{\infty} y_k^{(n)} t^k = y^n = \left(\sum_{l=1}^{\infty} y_l^{(i)} t^l\right)^{p^j} = \sum_{l=1}^{\infty} y_l^{(i)p^j} t^{lp^j}$. The result then follows by equating coefficients. □

Lemma 9.2.6. *There exists $c \in \mathbb{F}[y_m : m \leq k-n]$ such that*

$$y_k^{(n)} = \begin{cases} c + n y_1^{n-1} y_{k-n+1} & \text{if } n \leq k, \\ 0 & \text{else.} \end{cases}$$

Proof. By Lemma 9.0.8 $y_k^{(n)} = \sum \left\{ \prod_{j \in \mathbf{j}} y_j \middle| \mathbf{j} \in P(i; {}^n\mathbb{N}) \right\}$. Thus, we consider ordered n-partitions of k into natural numbers. No such partition can contain any number greater than $k-n+1$ as a summand, since then the sum would exceed k. On the other hand, there are exactly n partitions containing $k-n+1$; these are all permutations of the order of $(1, ..., 1, k-n+1)$. □

Lemma 9.2.7. *Denote $k^* := kp^{-j} - i + 1$. There exists $c \in \mathbb{F}[y_m : m < k^*]$ such that*

$$y_k^{(n)} = \begin{cases} c + \left(i y_1^{i-1} y_{k^*}\right)^{p^j} & \text{if } p^j \mid k \text{ and } n \leq k, \\ 0 & \text{else.} \end{cases}$$

Proof. By Lemma 9.2.5 we have that $y_k^{(n)} = y_{kp^{-j}}^{(i)p^j}$. Then we apply Lemma 9.2.6 to $y_{kp^{-j}}^{(i)}$. We note that $i \leq kp^{-j}$ if and only if $n \leq k$; and that $l_n(k) = kp^{-j} - i + 1$. □

Lemma 9.2.8.

$$gi_{x^n}(k) = \begin{cases} kp^{-j} - i + 1 & \text{if } p^j \mid k \text{ and } n \leq k, \\ 0 & \text{else.} \end{cases}$$

Proof. Again we denote $k^* := kp^{-j} - i + 1$. We apply Lemma 9.2.7. Simply note that $p \nmid i$ so y_{k^*} appears non-trivially in $y_k^{(n)}$ (of course, in the case that $p^j \mid k$ and $n \leq k$). □

Remark 9.2.9. Note that the function gi_{x^n} is equal to the linear function $k \longmapsto kp^{-j} - i + 1$ everywhere that it does not take the value 0.

Lemma 9.2.10. $y_k^{(n)} \in \mathbb{F}[y_m : m \leq kp^{-j} - i + 1]$.

Proof. This follows *a fortiori* from Lemma 9.2.7 and Lemma 9.2.8. □

9.2.3 The greatest-index of f

We use our understanding of the greatest-indices of monomials to study the greatest-index function of a polynomial. Recall that $S_f := \{n | c_n \neq 0\}$ is the *support* of f.

Lemma 9.2.11. *Let $k \in \mathbb{N}$. Then $C_{f,k} = \sum_{n \in S_f} c_n C_{x^n,k}$.*

Proof. Clear from the writing of f as $\sum_{n=1}^{N} c_n x^n$. □

Definition 9.2.12. Let $\bigvee gi_f := \max\{gi_{x^n} | n \in S_f\}$ be the function defined to be the point-wise maximum of the greatest-indices of the monomials of f. We call $\bigvee gi_f$ the *greatest-posible-index* of f.

Lemma 9.2.11 shows us that the variables y_m of $C_{f,k}$ are contributed by the summands $C_{x^n,k}$, for $n \in S_f$. Since we are aiming to find the greatest-index of a variable occuring non-trivially in $C_{f,k}$, we naturally wonder whether this is simply the maximum of the greatest-indices of the variables occuring non-trivially in the coefficients $C_{x^n,k}$, for $x \in S_f$. Lemma 9.2.13 answers the easy direction of this question.

Lemma 9.2.13. *Let $k \in \mathbb{N}$. Then $gi_f \leq \bigvee gi_f$.*

Proof. Obvious from Lemma 9.2.11. Any variable y_i appearing non-trivially in $C_{f,k}$ must come from one of the summands $C_{x^n,k}$, by Lemma 9.2.11. Thus is contributed by one of the monomials x^n, for $n \in S_f$. □

9.2.4 Avoiding coincidences

It need not be the case that equality holds in Lemma 9.2.13. The contributions of the different monomials of f can 'cancel out' to reduce the greatest-index of f below the maximum of the greatest-indices of the monomials of f. We call this *interference*. The cause of interference is the *coincidence* (i.e. equality) of the greatest-indices at k of two different monomials of f.

Our focus now switches to finding out when we may avoid interference. Next we make precise these notions of interference and coincidence. Note that coincidence does not imply interference!

Definition 9.2.14. We say that k is an *interference* point for f if $gi_f(k) < \bigvee gi_f(k)$. Let $M_f := \{k \in \mathbb{N} | gi_f(k) < \bigvee gi_f(k)\}$ be the set of interference points for f.

We say that k is a *coincidence* point for f if $gi_m(k) = gi_n(k)$ for distinct $m, n \in S_f$. Let $N_f := \{k \in \mathbb{N} | gi_{x^m}(k) = gi_{x^n}(k) = \bigvee gi_f(k)$ for distinct $m, n \in S_f\}$ be the set of coincidence points for f.

Remark 9.2.15. Note that N_f may be an infinite set because the greatest-index of a monomial may take the value 0 infinitely often.

Lemma 9.2.16. $M_f \subseteq N_f$.

Proof. Interference implies coincidence: for a variable y_i to occur in the contribution of a monomial but not to occur in the polynomial means that it must have been cancelled out by the contribution of another monomial. Basically, this is linear independence of monomials. □

The next lemma, although not used in the main argument, is illustrative of the purpose of the notions of coincidence and interference.

Lemma 9.2.17. *Suppose that $k \notin N_f$. There exists a unique $h \in S_f$ such that $gi_{x^h}(k) = \bigvee gi_f(k)$. Consequently $gi_f(k) = \bigvee gi_f(k)$.*

Proof. The existence and uniqueness of h are immediate from the definition of N_f. We set $k^* := kp^{-j} - i + 1$, where $h = p^j i$ and $p \nmid i$. By Lemma 9.2.8, the variable y_{k^*} occurs non-trivially in $C_{x^h,k}$ and not at all in $C_{x^n,k}$, for $n \in S_f \setminus \{h\}$. By Lemma 9.2.11, y_{k^*} occurs non-trivially in $C_{f,k}$. Thus $gi_f(k) \geq \bigvee gi_f(k)$. □

9.2.5 Solving the equation $C_{f,k} = b_k$

We avoid coincidences (thus avoiding interference) to solve the equation $C_{f,k} = b_k$. As before, we write $h = p^j i$, where $p \nmid i$.

Proposition 9.2.18. *Let $k \in \mathbb{N} \setminus N_f$. Denote $k^* := gi_f(k)$. There exists $d \in \mathbb{F}_f[y_m | m < k^*]$ such that*

$$C_{f,k} = d + c_h \left(i y_1^{i-1} y_{k^*}\right)^{p^j}.$$

Proof. Since $k \notin N_f$ and by the definition of N_f, there exists $h \in S_f$ such that $k^* := gi_{x^h}(k) > gi_{x^n}(k)$, for all $n \in S_f \setminus \{h\}$. By Lemma 9.2.10,

$$C_{x^n,k} \in \mathbb{F}[y_m | m \leq gi_{x^n}(k)] \subseteq \mathbb{F}[y_m | m < k^*],$$

for $n \in S_f \setminus \{h\}$. By Lemma 9.2.7, there exists $c \in \mathbb{F}[y_m | m < k^*]$ such that

$$C_{x^h,k} = c + (iy_1^{i-1} y_{k^*})^{p^j}.$$

Let $d := \sum_{n \in S_f \setminus \{h\}} c_n C_{x^n,k} + c$. Then $d \in \mathbb{F}_f[y_m | m < k^*]$ and

$$C_{f,k} = d + c_h (iy_1^{i-1} y_{k^*})^{p^j},$$

as required. □

Proposition 9.2.19. *Let $k \in \mathbb{N} \setminus N_f$. Denote $k^* := gi_f(k)$. Let $a := (a_m | m < k^*)$ be any tuple with $a_1 \neq 0$ and let $b_k \in K$. Then there exists (a unique) $a_{k^*} \in \mathbb{F}_f[a_1^{1-i}][a_m, b_k, c_h^{-1} | m < k^*]^{p^{-j}}$ such that $C_{f,k}(a) = b_k$.*

Proof. By Proposition 9.2.18 there exists that $d \in \mathbb{F}_f[y_m | m < k^*]$ such that $C_{f,k} = d + c_h \left(iy_1^{i-1} y_{k^*}\right)^{p^j}$. We set

$$Y := \left(\frac{b_k - d}{c_h}\right)^{p^{-j}} \cdot \frac{1}{iy_1^{i-1}} \in \mathbb{F}_f[y_1^{1-i}, b_k, c_h^{-1}, y_m | m < k^*]^{p^{-j}}.$$

Note that $p \nmid i$ and that $c_h \in S_f$ (thus $c_h \neq 0$); thus Y is the p^j-th root of a well-defined rational function in the variables $(y_m | m < k^*)$. We choose

$$a_{k^*} := Y(a) = \left(\frac{b_k - d(a)}{c_h}\right)^{p^{-j}} \cdot \frac{1}{ia_1^{i-1}} \in \mathbb{F}_f[a_1^{1-i}, b_k, c_h^{-1}, a_m | m < k^*]^{p^{-j}}.$$

This is well-defined since $a_1 \neq 0$. Using a_{k^*} to extend the tuple a (by redefining $a := (a_m | m \leq k^*)$), we see that the equation $C_{f,k}(a) = b_k$ is solved. Note also that a_{k^*} is unique since we all we have done to find it is to rearrange the equation we require it to solve. □

9.3 The Hensel degree

In Proposition 9.2.18 we saw that we may solve the equation $C_{f,k} = b_k$, if $k \notin N_f$. Our greater aim is to prove a Hensel-like lemma (namely Theorem 9.4.2) that we may refine an approximate solution a' of the polynomial equation $f = b$ to an exact solution a. By Lemma 9.2.2, this involves simultaneously solving the equations

$$C_{f,k} = b_k, k \in \mathbb{N}.$$

We aim to proceed by an induction on k, using Proposition 9.2.19 as an inductive step. There are two principal difficulties:

1. we must solve the equation $C_{f,k} = b_k$ without interfering with the solution of previous equations $C_{f,l} = b_l$, for $l < k$; and
2. k may be an interference point for f (see Definition 9.2.14).

In Lemma 9.2.16, we saw that the may avoid interference by avoiding coincidences; however we also remarked that there may well be infinitely many coincidences.

We seek a natural number H_f, which we will call the *Hensel degree*, such that we may solve $C_{f,k} = b_k$, for all $k \geq H_f$. We also would like gi_f to be strictly monotonically increasing after H_f, otherwise solving the equation $C_{f,k} = b_k$ might interfere with solutions of equations from earlier in the induction.

9.3.1 The well-order $\preceq$ on $\mathbb{N}$

Definition 9.3.1. For $m, n \in \mathbb{N}$, write

$$m \preceq_p n \text{ iff } v_p(m) \leq v_p(n).$$

Note that $\preceq_p$ is certainly not asymmetric! We define a (weak) linear ordering $\preceq$ on $\mathbb{N}$ to be the lexicographic product

$$\preceq := \preceq_p \underset{\text{lex}}{\otimes} \leq .$$

Thus

$$m \preceq n \text{ iff } (v_p(m) < v_p(n) \text{ or } (v_p(m) = v_p(n) \text{ and } m \leq n)).$$

We also write

$$m \trianglelefteq n \text{ iff } (v_p(m) = v_p(n) \text{ and } m \leq n).$$

These orderings will help us to understand the functions gi_{x^n}, for $n \in S_f$.

Lemma 9.3.2. *$\preceq$ is a well-order on $\mathbb{N}$.*

9.3.2 Comparing 'linearised' greatest-indices

As we saw in Remark 9.2.9, the greatest-index of a monomial is equal to a linear function everywhere it is not zero. This motivates the following definition. For the rest of this section, we fix $m, n \in \mathbb{N}$ and write $m = ip^j$ and $n = Ip^J$, where $p \nmid i, I$.

Definition 9.3.3. We define the *linearised greatest-index* of x^m to be the function

$$\begin{aligned} l_m : \mathbb{N} &\longrightarrow \mathbb{N} \cup \{0\} \\ k &\longmapsto kp^{-j} - i + 1. \end{aligned}$$

These functions define lines in the real plane. Understanding the intersection points of these lines will help us understand when equality holds in Lemma 9.2.13.

Definition 9.3.4. Set

$$x_{m,n} := \begin{cases} p^{j+J}(i-I)/(p^J - p^j) & \text{if } m \prec n \text{ and } j \neq J \text{ (whence } j \leq J); \\ -\infty & \text{if } m \prec n, j = J, \text{ and } i < I; \text{ and} \\ \text{(undefined)} & \text{if } n \preceq m. \end{cases}$$

Lemma 9.3.5. *Suppose that $m \prec n$. Let $k \in \mathbb{R}$. Then*

$$l_m(k) \left\{ \begin{array}{c} < \\ = \\ > \end{array} \right\} l_n(k) \text{ if and only if } k \left\{ \begin{array}{c} < \\ = \\ > \end{array} \right\} x_{m,n}.$$

Proof. Suppose that $j \neq J$. Then

$$\begin{aligned} l_m(k) < l_n(k) \quad &\text{iff} \quad kp^{-j} - i + 1 < kp^{-J} - I + 1 \\ &\text{iff} \quad k(p^{-j} - p^{-J}) < i - I \\ &\text{iff} \quad k < (i-I)/(p^{-j} - p^{-J}) \\ &\text{iff} \quad k < p^{j+J}(i-I)/(p^J - p^j) = x_{m,n}. \end{aligned}$$

If $j = J$ then $i < I$ and $l_m(k) = kp^{-j} - i + 1 > kp^{-J} - I + 1 = l_n(k)$; but also $x_{m,n} = -\infty$ and so certainly $k > x_{m,n}$. □

Lemma 9.3.6. *If $m \lhd n$ then $l_m > l_n$.*

Proof. In this case $x_{m,n} = -\infty$. For all $k \in \mathbb{R}$, $k > -\infty$ and so $l_m(k) > l_n(k)$. □

Lemma 9.3.7. *If $m \preceq n$ and $k > x_{m,n}$ then $l_m(k) > l_n(k)$.*

9.3.3 Comparing greatest-indices

Lemma 9.3.8. *If $m \leq k$ then $0 \leq gi_{x^m}(k) \leq l_m(k) \geq 1$.*

Proof. Since $m \leq k$, we have that $l_m(k) = kp^{-j} - i + 1 \geq 1$. Either $gi_{x^m}(k) = l_m(k)$ or $gi_{x^m}(k) = 0 < 1 \leq l_m(k)$. □

Lemma 9.3.9. *If $m \preceq n$ then*

1. *if $p^j \nmid k$ then $gi_{x^m}(k) = gi_{x^n}(k) = 0$, and*
2. *if $p^j \mid k$ and $k \geq m$ and $k > x_{m,n}$ then $gi_{x^m}(k) > gi_{x^n}(k)$.*

Proof. Suppose that $p^j \nmid k$; then $p^J \nmid k$. So, by Lemma 9.2.8, $gi_{x^m}(k) = gi_{x^n}(k) = 0$. If instead $p^j \mid k$ and $k \geq m$ and $k > x_{m,n}$ then $gi_{x^m}(k) = l_m(k) > l_n(k)$. If additionally $k \geq n$ then $l_n(k) = gi_{x^n}(k)$, and we are done. Otherwise $k < n$ and thus $g_{x^n} = 0$, so certainly $g_{x^m}(k) > g_{x^n} = 0$. □

9.3.4 Crossing points

Definition 9.3.10. We define the set of *crossing points* to be

$$\chi_f := \{k \in \mathbb{R} | l_m(k) = l_n(k) \text{ for distinct } m, n \in S_f\},$$

which is the projection onto the first co-ordinate of the set of real intersection points of distinct pairs of these lines.

Lemma 9.3.11. *χ_f is finite.*

Proof. The polynomials $l_m = kp^{-j} - i + 1$ are linear. Thus the graphs of any distinct pair of them can intersect at most once. The set χ_f is the projection of this set of intersections. □

Lemma 9.3.12. *$\chi_f = \{x_{m,n} | m, n \in S_f$ such that $m \prec n\}$.*

Proof. Obvious from Lemma 9.3.5. □

9.3.5 The monomial-in-chief

Using $\preceq$ we are able to understand the behaviour of gi_f in terms of the greatest-index of one of the monomials of f.

Definition 9.3.13. Set $n_f := \min_{\preceq} S_f$. This is the *monomial-in-chief* of f.

Proposition 9.3.14. *Let $k > \max\chi_f$ and let $m \in S_f \setminus \{n_f\}$. Then $l_m(k) < l_{n_f}(k)$.*

Proof. Since $n_f \prec m$, $x_{n_f,m} \in \chi_f$ is well-defined. By Lemma 9.3.7 and since $k > x_{n_f,m}$, we have that $l_m(k) < l_{n_f}(k)$. □

9.3.6 The Hensel degree

We have just seen that the linearised greatest-index of x^{n_f} will be eventually greater than the linearised greatest-indices of the other monomials of f. Our next goal is to prove that the greatest-index of x^{n_f} is eventually greater than the other greatest-indices and to identify more precisely what is meant by 'eventually'.

Definition 9.3.15. We define the *Hensel degree* of f to be

$$H_f := \max\{\min\{h \in \mathbb{N} | \forall k \geq h \ \forall m \in S_f \ (m \neq n_f \implies l_m(k) < l_{n_f}(k))\}, n_f\}.$$

Let $H_f^* := gi_{x^{n_f}}(H_f)$.

Lemma 9.3.16. *H_f is well-defined.*

Proof. We just need to check that $X := \{h \in \mathbb{N} | \forall k \geq h \ \forall m \in S_f \ (m \neq n_f \implies l_m(k) < l_{n_f}(k))\}$ is non-empty. We note that χ_f is a finite set; we choose $k > \max\chi_f$. Then $l_m(k) < l_{n_f}(k)$ by Proposition 9.3.14. Thus $k \in X$ and so H_f is well-defined. □

Recall the definition of p-th root depth from Definition 9.0.3. Note that $\mathrm{dep}_p(f) = \log_p(n_f)$. Write $n_f = ip^j$, where $p \nmid i$.

Lemma 9.3.17. *Let $k \geq H_f$. Then*

1. *if $p^j \mid k$ then for all $m \in S_f \setminus \{n_f\}$ we have that $gi_{x^m}(k) < gi_{x^{n_f}}(k)$; or*
2. *if $p^j \nmid k$ then for all $m \in S_f$ we have that $gi_{x^m}(k) = 0$.*

Proof. Since $k \geq H_f \geq n_f$, we have that $gi_{x^{n_f}}(k) = 0$ if and only if $p^j \nmid k$. Now suppose that $gi_{x^{n_f}}(k) = 0$ and let $m \in S_f \setminus \{n_f\}$. Then $gi_m(k) = 0$.

Suppose instead that $gi_{n_f}(k) \neq 0$ and again let $m \in S_f \setminus \{n_f\}$. Then $l_m(k) < l_{n_f}(k) = gi_{n_f}(k)$, since $k \geq H_f$. Either $gi_{x^m}(k) = l_m(k)$, in which case we are done; or $gi_{x^m}(k) = 0$, in which case we are also done since $gi_{x^{n_f}}(k) > 0$. □

Proposition 9.3.18. *Let $k \geq H_f$. Then $k \in N_f$ if and only if $p^j \nmid k$. Also gi_f is strictly monotonically increasing on the set of $k \in \mathbb{N}$ such that $k \geq H_f$ and $p^j \mid k$.*

Proof. Immediate from the definition of the set N_f of coincidence points and from Lemma 9.3.17. For the second part, note that for such k, $gi_f(k) = gi_{x^{n_f}}(k) > gi_{x^m}(k)$, for all $m \in \mathcal{S}_f \setminus \{n_f\}$. □

9.4 The Hensel-like lemma

We can finally prove the Hensel-like lemma by induction. This first proposition will form the basis of the inductive step.

Proposition 9.4.1. *Let $f \in xF[x]$ and let $k \geq H_f$ and suppose that $p^j \mid k$. Let $a' \in \mathcal{P}$ and $b \in \mathcal{M}$ be such that $f(a') \in b + \mathcal{M}^k$ and $\mathrm{dep}_p(b) \geq \mathrm{dep}_p(f)$. Then there exists $a \in a' + \mathcal{M}^{k^*}$ such that every $a'' \in a + \mathcal{M}^{k^*+1}$ satisfies $f(a'') \in b + \mathcal{M}^{k+1}$. In particular $f(a) \in b + \mathcal{M}^{k+1}$.*

Proof. Let $(a_m | m < k^*)$ be the tuple of the first $k^* - 1$-many coefficients of a'. (This is a slight abuse of notation since we decreed above that such elements would stand for coefficients of a). As above, b_k will denote the k-th coefficient of b. Note that $a_m, b_k, c_n \in F$, for each $m < k^*$ and $n \in S_f$. We need to solve the equation $C_{f,k}(a_m, y_{k^*} | m < k^*) = b_k$. By Proposition 9.2.19, there exists a unique $a_{k^*} = \mathbb{F}_f[a_1^{1-i}][b_k, c_h^{-1}, a_m | m < k^*]^{p^{-j}}$ such that $C_{f,k}(a_m | m \leq k^*) = b_k$. Note that

$$\mathbb{F}_f[a_1^{1-i}][a_m, b_k, c_h^{-1} | m < k^*]^{p^{-j}} \subseteq F^{p^{-j}} = F.$$

Now we let $a := \sum_{m<k^*} a_k t^k + a_{k^*} t^{k^*}$. Then $a \in a' + \mathcal{M}^{k^*}$ and, by our choice of a_{k^*}, $C_{f,k}(a) = b_k$.

We now turn to the question of whether a'' solves all the equations

$$C_{f,k'} = b_{k'},$$

for each $k' \leq k$. By Proposition 9.3.18, the left hand side of this equation is necessarily zero if $p^j \nmid k$;

but since $\mathrm{dep}_p(b) \geq \mathrm{dep}_p(f)$ we have that $p^j \nmid k \implies b_k = 0$. Thus if the left hand side is necessarily zero, the right hand side is already zero. This leaves us only to worry about $k \notin N_f$, on which set gi_f is strictly monotonically increasing. Thus the equations $C_{f,k'} = b_{k'}$ (with $k' \leq k$) only mention variables with indices $\leq k^* = gi_f(k)$. So any a'' in the $\mathcal{M}^{k^*+1}$-coset of a will solve all these equations, as required. □

Theorem 9.4.2. *Let $f \in xF[x]$. Let $a' \in \mathcal{P}$ and $b \in \mathcal{M}$ be such that $f(a') \in b + \mathcal{M}^{H_f}$ and $\mathrm{dep}_p(b) \geq \mathrm{dep}_p(f)$. Then there exists (a unique) $a \in a' + \mathcal{M}^{H_f^*}$ such that $f(a) = b$.*

Proof. We define a sequence of elements $a_{[k]}$, for $k \geq H_f - 1$. Let $a_{[H_f-1]} := a'$. Then $f(a_{[H_f-1]}) \in b + \mathcal{M}^{H_f}$. We proceed by induction using Proposition 9.4.1. Let $k \geq H_f$ and suppose we have found $a_{[k-1]} \in a' + \mathcal{M}^{H_f^*}$ such that $f(a_{[k-1]}) \in b + \mathcal{M}^k$ (and that this holds for any element of $a_{[k-1]} + \mathcal{M}^{gi_f(k-1)+1}$). If $p^j \nmid k$ then both sides of the equation $C_{f,k} = b_k$ are zero (since $\mathrm{dep}_p(f) = \mathrm{dep}_p(b)$, by assumption). There is nothing to do in this case; we automatically have that $f(a_{[k-1]}) \in b + \mathcal{M}^{k+1}$.

Alternatively $p^j \mid k$, in which case we simply apply Proposition 9.4.1. Thus, we find $a_{[k]} \in a_{[k-1]} + \mathcal{M}^{k^*} = a_{[k-1]} + \mathcal{M}^{gi_f(k)} \subseteq a_{[k-1]} + \mathcal{M}^{gi_f(k-1)+1}$ (since $gi_f(k-1) < gi_f(k)$) such that $C_{f,k}(a_{[k]}) = b_k$. Since also $a_{[k]} \in a_{[k-1]} + \mathcal{M}^{gi_f(k-1)+1}$ we have that $f(a_{[k]}) \in b + \mathcal{M}^k$. Thus $f(a_{[k]}) \in b + \mathcal{M}^{k+1}$. Again, by Proposition 9.4.1, we find that this will hold for any element of $a_{[k]} + \mathcal{M}^{gi_f(k)+1}$. This proves the inductive step.

Since $F((t))$ is complete, we may let $a \in a' + \mathcal{M}^{H_f^*}$ be the unique limit of the sequence in $F((t))$. This element certainly solves all the equations. Thus $f(a) = b$. □

Remark 9.4.3. The above theorem is obviously 'Hensel-like' in character: it describes how we may refine an approximate solution of a polynomial equation to an exact one. In order to apply the usual version of Hensel's Lemma, one must have an approximate solution to a polynomial equation $f(x) = 0$: that is an element a such that $vDf(a) = 0$ and $vf(a) > 0$. If f is a p-polynomial (i.e. each power of x which occurs is a multiple of p) then there is no hope of finding such an element a because the formal derivative Df is the zero polynomial. Theorem 9.4.2 enables us to refine approximate solutions (under certain assumptions) to equations of the form $f(x) = b$.

In the next chapter, we will use this understanding to study the orbits of elements of $\mathcal{M}$ under all F-automorphisms of $F((t))$.

Chapter 10

F-definability

In this chapter we are able to apply the Hensel-like lemma (Theorem 9.4.2) to the question of F-definability in fields of formal power series $F((t))$. It turns out that F-orbits are $\exists$-t-definable and this enables us to transfer some of the results from chapter 7 about $\exists$-definability to F-definability.

Throughout we make still make the assumption that F **is perfect**.

10.1 F-orbits

Let $b \in F((t))$. Our aim is to understand the F-orbit of b, which we denote by $\mathrm{Orb}(b/F)$. We first tackle the special case $b \in \mathcal{M} \setminus \{0\}$ by defining a polynomial $f_b \in xF[x]$ (which we call the auxilliary polynomial of b) such that $f_b(t) \in b + \mathcal{M}^{H_{f_b}}$ (using the notation of chapter 9). In the second section, we apply Theorem 9.4.2 to the polynomial f_b to find some $s \in \mathcal{P}$ such that $f_b(s) = b$. In this way we see that every element is 0-interdefinable with an element of $\mathcal{P}$. In section 9.1, we proved that there was only one F-orbit of elements of $\mathcal{P}$.

10.1.1 The sequence of p^n-prime t-adic 'values'

Let $b \in \mathcal{M} \setminus \{0\}$. As stated above, our goal is to identify a polynomial $f_b \in xF[x]$ such that $f(t) \in b + \mathcal{M}^{H_f}$, where H_f denotes the *Hensel degree* of f as defined in the previous chapter. Then we can apply Theorem 9.4.2 in the next section. First we associate to b a sequence $h_j, ..., h_0$ which will help us understand which polynomials have the same Hensel degree. Let $v_p : \mathbb{Q} \longrightarrow \mathbb{Z}$ be the p-adic valuation.

Definition 10.1.1. Let $j \in \mathbb{N}$. We define the *p^j-prime t-adic valuation* to be

$$\begin{aligned} w_j : \quad F((t)) &\longrightarrow \mathbb{N} \cup \{\infty\} \\ b &\longmapsto \min\{l \in S_b | p^j \nmid l\}, \end{aligned}$$

where, as usual, $S_b := \{l \in \mathbb{Z} | b_l \neq 0\}$ denotes the *support* of b. We may also define $w_\infty := v$ to be the usual t-adic valuation. Recall that we defined the *p-th root depth* of $b \in F((t))$ to be $\mathrm{dep}_p(b) := \max\{l \in \omega \cup \{\infty\} | b \in F((t^{p^l}))\}$.

Remark 10.1.2. Beware that w_m is not a valuation and note that w_0 is the constant function with value ∞.

Write $vb = Ip^J$, where $p \nmid I$.

Lemma 10.1.3. *We have that* $v(b) = w_\infty(b) = ... = w_{J+1}(b) \leq ... \leq w_1(b) \leq w_0(b) = \infty$.

Proof. For each $l \in \mathbb{N}$, we have that $(p^{j+1} \mid l \implies p^j \mid l)$. Thus $(p^j \nmid l \implies p^{j+1} \nmid l)$. Since $p^{J+1} \nmid vb$, $vb = w_{J+1}(b)$. □

Lemma 10.1.4. *$b \in F((t))^{p^j}$ if and only if $w_j(b) = \infty$.*

Proof. Write $b = \sum b_l t^l$. Then $b \in F((t))^{p^j}$ if and only if $(p^j \nmid l \implies b_l = 0)$, because F is perfect. The latter holds if and only if $w_j(b) = \infty$. □

Lemma 10.1.5. $\mathrm{dep}_p(b) = \max\{j \in \mathbb{N} | w_j(b) = \infty\}$.

Proof. Clear from the definitions. □

10.1.2 The auxilliary polynomial

Let $f, g \in xF[x]$ be two polynomials over F with zero constant terms.

Definition 10.1.6. We write $f \preceq g$ if $S_f \subseteq S_g$ and, for each $n \in S_g$, there exists $m \in S_f$ such that $m \trianglelefteq n$.

Proposition 10.1.7. *Suppose that $f \preceq g$. Then $n_f = n_g$ and $H_f = H_g$ and* $\mathrm{dep}_p(f) = \mathrm{dep}_p(g)$.

Proof. Let $m \in S_g$. Then by definition there exists $m' \in S_f$ such that $m' \trianglelefteq m$. By definition of $n_f := \min_\preceq S_f$, $n_f \preceq m' \trianglelefteq m$. Thus $n_f = \min_\preceq S_g =: n_g$.

Now set $h_f := \min\{h \in \mathbb{N} | \forall k \geq h \ \forall m \in S_f \ (m \neq n_f \implies l_m(k) < l_{n_f}(k))\}$. Let $k \geq h_f$ and let $m \in S_g$. Then there exists $m' \in S_f$ such that $m' \trianglelefteq m$; thus $l_m \leq l_{m'}$. By definition of h_f, $l_{m'}(k) < l_{n_f}(k)$. Combining these two inequalities, we have that $l_m(k) \leq l_{m'}(k) < l_{n_f}(k)$. Thus $h_f = h_g$ and so $H_f = H_g$. The depth of an element is entirely determined by its monomial-in-chief. □

Definition 10.1.8. Let $b \in \mathcal{M} \setminus \{0\}$. Let H_b be defined as it would for a polynomial. We define the *auxilliary polynomial* of b to be the truncation $f_b := \sum_{h \leq H_b} b_h x^h$ of b to the degree H_b.

Proposition 10.1.9. *We have that $f_b \preceq b$. Thus $n_{f_b} = n_b$ and $H_{f_b} = H_b$ and $\operatorname{dep}_p(f_b) = \operatorname{dep}_p(b)$.*

Proof. Since f_b is a truncation of b, $S_f \subseteq S_b$. Let $m \in S_b$ and write $m = ip^j$, where $p \nmid i$. Then $w_j(b) \leq w_1(b) = n_b \leq H_b$. Also $b_{w_j}(b) \neq 0$, thus $w_j(b) \in S_f \subseteq S_b$. Since $w_j(b) \trianglelefteq m$, we have that $f_b \preceq b$, as required. By Proposition 10.1.7 $n_{f_b} = n_b$ and $H_{f_b} = H_b$ and $\operatorname{dep}_p(f_b) = \operatorname{dep}_p(b)$. □

Proposition 10.1.10. *We have that $f_b(t) \in b + \mathcal{M}^{H_{f_b}+1}$.*

Proof. Clear from the definition of f_b. □

Proposition 10.1.11. *Let $b \in \mathcal{M} \setminus \{0\}$. There exists a unique $a \in t + \mathcal{M}^{H_f^*+1} \subseteq \mathcal{P}$ such that $f_b(a) = b$.*

Proof. It follows from Proposition 10.1.10 that $f_b(t) \in b + \mathcal{M}^{H_{f_b}}$. Thus we may apply the Hensel-like lemma (Theorem 9.4.2) to the auxiliary polynomial f_b. □

Now we extend the notion of auxilliary polynomial to cover all elements of $\mathcal{O}$. Note that bv denotes the residue of b under the t-adic valuation.

Definition 10.1.12. Let $b \in \mathcal{O}$. We define the *auxilliary polynomial* of b to be

$$f_b := \begin{cases} f_{b-bv} + bv & \text{if } b - bv \neq 0 \text{ and} \\ bv & \text{else,} \end{cases}$$

where f_{b-bv} denotes the auxilliary polynomial (in the sense of Definition 10.1.8) of $b - bv \in \mathcal{M}$.

Thus $H_{f_b} = H_{f_{b-bv}}$ and $H_{f_b}^* = H_{f_{b-bv}}^*$, for $b \in \mathcal{O}$ such that $b \neq bv$.

Theorem 10.1.13. *Let $b \in \mathcal{O}$. There exists $a \in t + \mathcal{M}^{H_{f_b}^*}$ such that $f_b(a) = b$. If $b \notin F$ then a is unique.*

Proof. If $b = bv$ then $f_b = bv$ is a constant and any a satisfies $f_b(a) = f_b = bv = b$. Otherwise we apply Proposition 10.1.11 to $b - bv$ and the polynoimal f_{b-bv}. We find a unique $a \in t + \mathcal{M}^{H_{f_b}^*}$ such that $f_{b-bv}(a) = b - bv$ and $f_b(a) = f_{b-bv}(a) + bv = b - bv + bv = b$. □

10.1.3 F-orbits are $\exists$-definable

We now aim to fully characterise F-orbits of elements $b \in F((t))$ using Theorem 10.1.13. We start with elements of $\mathcal{O}$.

Proposition 10.1.14. *Let $b \in \mathcal{O}$. Let $f_b \in F[x]$ be the auxilliary polynomial of b. Then* $\mathrm{Orb}(b/F) = f_b(\mathcal{P})$.

Proof. By Theorem 10.1.13 there exists $a \in \mathcal{P}$ such that $f_b(a) = b$. We apply a F-automorphism ϕ to $f_b(a) = b$ to obtain

$$f_b(\mathcal{P}) \ni f_b(a\phi) = f_b(a)\phi = b\phi \in \mathrm{Orb}(b/F).$$

Let $s \in \mathcal{P}$ and let $\phi \in \mathrm{Aut}(F((t))/F)$ be a F-automorphism such that $\phi : a \longmapsto s$. Then $f_b(s) = f_b(a\phi) = f_b(a)\phi = b\phi \in \mathrm{Orb}(b/F)$. Conversely, let $\psi \in \mathrm{Aut}(F((t))/F)$ and let $s := a\phi$. Then $f_b(s) = f_b(a\phi) = f_b(a)\phi = b\phi \in \mathrm{Orb}(b/F)$. Thus $b\phi \in f(\mathcal{P})$. Now it is clear that $\mathrm{Orb}(b/F) = f_b(\mathcal{P})$. □

Proposition 10.1.15. *Let $b \notin \mathcal{O}$. Then there exists $f \in xF[x]$ such that* $\mathrm{Orb}(b/F) = 1/f(\mathcal{P})$.

Proof. We have that $1/b \in \mathcal{M} \setminus \{0\}$, so we set $f := f_{1/b}$. By Proposition 10.1.14 $\mathrm{Orb}(\frac{1}{b}/F) = f(\mathcal{P})$; thus $\mathrm{Orb}(b/F) = 1/f(\mathcal{P})$. □

Proposition 10.1.16. *Let $b \in F((t))$. Then* $\mathrm{Orb}(b/F)$ *is $\exists$-t-definable and $\exists$-$\mathcal{P}$-definable.*

Proof. Suppose that $b \in \mathcal{O}$ and let f be as in Proposition 10.1.14. Then $\mathrm{Orb}(b/F)$ is defined by the formula

$$\phi_b := \exists y \ (y \in \mathcal{P} \wedge f(y) = x).$$

Suppose instead that $b \notin \mathcal{O}$ and let f be an in Proposition 10.1.15. Then $\mathrm{Orb}(b/F)$ is defined by the formula

$$\phi_b := \exists y \ (y \in \mathcal{P} \wedge 1 = f(y)x).$$

Since $\mathcal{P}$ is $\exists$-t-definable, these formulas can also be written as $\exists$-t-formulas. □

Although orbits need not be open sets (for each the orbit of t^p is just $\mathcal{P}^p$), this is the only thing that can go wrong.

Proposition 10.1.17. *Let $b \in F((t))$. Then* $\mathrm{Orb}(b/F)$ *is the image of $\mathcal{P}^{p^n}$ under an open map, for some $n \in \omega$.*

Proof. All polynomials can be written as the composition of a purely inseparable polynomial with a polynomial with non-zero formal derivative. □

10.2 F-definable subsets

We use Proposition 10.1.16 to prove that F-definable subsets of $F((t))$ are large in various senses.

Theorem 10.2.1. *Let $X \subseteq F((t))$ be an F-definable subset not contained in F. Then $X \setminus F$ contains an infinite $\exists$-t-definable subset.*

Proof. This is a simple consequence of Proposition 10.1.16. Let $b \in X \setminus F$. By Proposition 10.1.16, $\mathrm{Orb}(b/F)$ is $\exists$-t-definable. The orbit of any element of $\mathcal{M} \setminus \{0\}$ is equal to $f(\mathcal{P})$ for some non-constant polynomial f. Any such set is clearly infinite, and any other orbit of an element not in F is in bijection with one of this form. □

Proposition 10.2.2. *Let $X \subseteq F((t))$ be an F-definable subset not contained in F. Then $X \setminus F$ has the same cardinality as $F((t))$.*

Proof. Immediate from Theorem 10.2.1 and Theorem 7.3.1. □

Corollary 10.2.3. *F is model-theoretically algebraically closed in $F((t))$.*

Proof. Let $a \in F((t)) \setminus F$. Then the orbit of a over F is infinite. Thus a is not model-theoretically algebraic over F. □

10.2.1 F-definable subrings

Theorem 10.2.4. *Let $R \subseteq F((t))$ be an F-definable subring not contained in F. There exists $n \in \mathbb{N}$, $a \in R$, and $\alpha \in \tau$ such that $B^n(\alpha; a) \subseteq R$.*

Proof. Again, this is clear from Theorem 10.2.1 and Theorem 7.8.1. □

Let (X) denote the subfield of K generated by X.

10.2.2 Subfields generated by an F-definable subset

Theorem 10.2.5. *Let $X \subseteq F((t))$ be a F-definable subset not contained in F. There exist $m, n \in \mathbb{N}$ such that $(F((t)))^{p^n} = (X)_m = (X)$. Consequently, (X) is $\exists$-t-definable (and $\exists$-$\mathcal{P}$-definable).*

Proof. By Theorem 10.2.1, there exists an infinite $Y \subseteq X \setminus F$ which is $\exists$-t-definable. By Theorem 7.9.3, there exists $m, n \in \mathbb{N}$ such that $F((t))^{p^n} \subseteq (Y)_m \subseteq (X)_m$. Let n be minimal with this property. Note that, since F is perfect, the only subfields of $F((t))$ which contain $(F((t)))^{p^n} = F((t^{p^n}))$ are of the same form, i.e. $F((t^{p^m}))$ for $m \leq n$. $\square$

Chapter 11

Further work

In the introduction, we stated that our wider goal was to understand the model theory of $\mathbb{F}_p((t))$ and, of course, that problem still lies before us. In this thesis we have made some progress towards a description of the sets definable in the $\mathcal{L}_{\text{ring}}$-theory of $\mathbb{F}_p((t))$; we now understand the local behaviour of existentially definable sets and subsets definable with no parameters (equivalently, parameters from $\mathbb{F}_p$). This brings us to our first goal: to extend these results to an understanding of *global* behaviour.

Extending Part III

For subsets definable with few parameters, an understanding of global behaviour would mean finding a description of which combinations of orbits could be in a definable set. Suppose that $X \subseteq \mathbb{F}_p((t))$ is $\mathbb{F}_p$-definable. We may assume that all the elements of X have the same angular component because there are only finitely many possible angular components and the angular component map is definable with just a predicate for the set $t(1 + \mathcal{M})$. The automorphisms preserving this extra predicate are precisely those which act by substitution by elements of $t(1 + \mathcal{M})$. Adapting our arguments from Part III, we can still obtain a good theory of orbits under this action. Therefore the most important thing to understand are the possible values taken by elements of X:

Question What are the possible sets of values of elements of X?

Another extension is to study definable sets in higher Cartesian powers.

Question What are the orbits of n-tuples? What are the $\mathbb{F}_p$-definable sets?

We might as well study orbits of $p^n + 1$-tuples. Let $(x_0, ..., x_{p^n})$ be a tuple from $\mathcal{M}$ and suppose that x_0 is p-independent. Then we may 'code' the other elements $x_1, ..., x_{p^n}$ into a new element, x say, by letting $x = \sum_{i<p^n} x_0^{i-1} x_{i-1}^{p^n}$. The information of the orbits of $(x_0, ..., x_{p^n})$ is coded into the orbit of (x_0, x). Therefore it suffices to undestand the orbits of 2-tuples.

We observe that each orbit of 1-tuples in $\mathbb{F}_p((t))$ includes elements from $\mathbb{F}_p(t)^h$. One of our goals is to prove that $\mathbb{F}_p(t)^h \preceq \mathbb{F}_p((t))$ or, weaker, that $\mathbb{F}_p(t)^h \equiv \mathbb{F}_p((t))$. An understanding of orbits of n-tuples might allow us to prove these statements by a simple application of the Tarski-Vaught test.

Extending Part II

Theorem 7.1.1 can be seen as a kind of 'local quantifier elimination' because it says that an existentially definable set is locally described by λ-algebraic equations, i.e. algebraic equations in the components of the variables. For sufficiently generic points, this description is exact in the sense that the existentially definable set is locally equal to the λ-algebraic set (rather simply containing it). Using a compactness argument, we can prove that there are in fact only finitely many λ-algebraic sets which describe this local behaviour. This gives a kind of uniformity to the use of Λ-alterations. Using these 'global' results, it might be possible to prove a true quantifier-elimination theorem.

Quite another line of research seeks to improve upon chapter 8. If we could remove the hypothesis of density from Theorem 8.1.1 then, by Robinson's test, we would have proved the model-completeness of $\mathbb{F}_p((t))$. This leads to the following question.

Question Given an extension L/K of models of $\mathbb{F}_p((t))$, is it possible to find compatible extensions $L \subseteq L'$ and $K \preceq_\exists K'$ such that L'/K' is dense and regular and K' is t-henselian?

Here, compatible means that the diagram should commute. Unfortunately the characteristic zero analogy is useless here because every extension of p-adically closed fields is elementary by model completeness of that theory.

Yet another line of investigation is to seek to replicate the Prestel-Roquette approach of [24]. This would mean finding an expansion of the ring language in which we could prove embedding theorems. In their work, in the context of p-adically closed fields, the relevant expansion is the Macintyre language. The obvious expansion would be to combine the Macintyre language with function symbols for the component maps (with respect to the p-base t). But inseparable extensions are not quite the only difference between characteristic zero and positive characteristic: Artin-Schreier extensions have an important rôle to play. It will probably be necessary to add in predicates for the Artin-Schreier subgroup and perhaps also the subgroups corresponding to composed Artin-Schreier polynomials. In attempting to prove embedding theorems in this context, it becomes clear that one of the main problems is the (valuation theoretical) defect.

Which brings us once again to the urgent realisation of just how much there is still left to know.[1]

[1] With apologies to Monty Python.

Appendices

Appendix A

Definability of the valuation ring

Let $\mathbb{F}_q((t))$ be the field of formal power series over the finite field $\mathbb{F}_q$. In this chapter we aim to prove that the valuation ring $\mathbb{F}_q[[t]]$ is $\exists$-$\emptyset$-definable in $\mathbb{F}_q((t))$. This is useful because it reduces questions of existential definability in $\mathcal{L}_{\text{val}}$ to existential definability in $\mathcal{L}_{\text{ring}}$, conservatively in parameters; i.e. without needing more parameters. First we briefly recall the most simple definition of $\mathbb{F}_q[[t]]$ which uses the parameters t.

The following definition is very well-known.

Proposition A.0.6. *$\mathcal{O}$ is defined in $\mathbb{F}_q((t))$ by the existential formula $\exists y\ 1 + x^l t = y^l$, for any prime l such that $l \nmid q$.*

Proof. Suppose that $x \in \mathcal{O}$. Clearly $1 + x^l t \in 1 + \mathcal{M} = (1 + \mathcal{M})^l$ by henselianity. Conversely, suppose x is such that $vx < 0$. Then $vx^l \leq -l$ and $v(x^l t) \leq 1 - l < 0$. Thus $v(1 + x^l t) = v(x^l t) = 1 + lvx$ cannot be divisible by l and there can exist no y such that $1 + x^l t = y^l$. □

Other definitions are also well-known. One example is an existential-universal definition with no parameters due to Ax, from [1], which applied to all power series fields; another definition, in greater generality, which uses no parameters is due to Koenigsmann and is Lemma 3.6 from [10].

A.1 The $\exists$-$\emptyset$-definability of $\mathbb{F}_q[[t]]$ in $\mathbb{F}_q((t))$

We prove that $\mathbb{F}_q[[t]]$ is existentially definable in $\mathbb{F}_q((t))$ using no parameters (equivalently, only parameters from $\mathbb{F}_p$). The method is to use the definition of the henselian topology given by Prestel-Ziegler

in [25] and Koenigsmann in [10] to give us an $\exists$-$\mathbb{F}_q$-definable bounded neighbouhood of 0. Then we 'tweak' this set by subtracting, taking roots, and applying our understanding of $\mathbb{F}_q$-orbits in $\mathbb{F}_q((t))$ (obtained in Proposition 10.1.14) in order to find an $\exists$-$\mathbb{F}_q$-definable subset of $\mathcal{O}$ which contains $\mathcal{M}$. Finally, we use the fact that $\mathbb{F}_q$ is the algebraic set defined by the formula $x^q - 1 \doteq 0$ to 'tweak' our set once more until it equals the valuation ring, as required.

A.1.1 Spheres and balls in valued fields

We briefly recap a few definitions and basic facts about valued fields. Let $(K, \mathcal{O})$ be a valued field and let vK denote the value group.

Definition A.1.1. For $\alpha \in vK$ and $a \in K$, we let

1. $S(\alpha) := v^{-1}(\{\alpha\})$ be the set of elements of value α,
2. $B(\alpha; a) := a + v^{-1}((\alpha, \infty))$ be the *open ball* of radius α around a, and
3. $\bar{B}(\alpha; a) := a + v^{-1}([\alpha, \infty))$ be the *closed ball* of radius α around a.

As before, let $\mathcal{P} := S(1)$ be the set of uniformisers.

Lemma A.1.2. *Let $\alpha \in vK$. Then*

1. $B(\alpha; 0) \subseteq S(\alpha) - S(\alpha)$,
2. $\bar{B}(\alpha; 0) = S(\alpha) \sqcup B(\alpha; 0)$, *and*
3. $\bar{B}(\alpha; 0) - \bar{B}(\alpha; 0) = \bar{B}(\alpha; 0)$.

Proof. 1. Let $x \in B(\alpha; 0)$ and let $y \in S(\alpha)$. Then $v(y) = \alpha < v(x)$, so that $v(x - y) = \alpha$ (by an elementary consequence of the ultrametric inequality) and $x - y \in S(\alpha)$. Thus $x = x - y + y \in S(n) - S(n)$.

2. Let $x \in \bar{B}(\alpha; 0)$. Then either $v(x) = \alpha$ or $v(x) > \alpha$.

3. Let $x, y \in \bar{B}(\alpha; 0)$. By the ultrametric inequality $v(x - y) \geq \alpha$. Thus $x - y \in \bar{B}(\alpha; 0)$.

□

A.1.2 An ∃-definable filter for the neighbourhood filter of zero

We now suppose that $\mathcal{O}$ induces a t-henselian topology on K. The following definition of the t-henselian topology is from [25] and [23].

Let $D := D_x$ denote the formal derivative with respect to the variable x.

Lemma A.1.3. *Suppose that K is not separably closed. Let $f \in K[x]$ be a separable irreducible polynomial without a zero in K. Let $a \in K \setminus Z(Df)$ be any element which is not a zero of the formal derivative of f. Let $U_{f,a} := f(K)^{-1} - f(a)^{-1}$. Then $\mathcal{U} := \{c \cdot U_{f,a} | c \in K^\times\}$ is a base for the open neighbourhoods around zero in the (unique) t-henselian topology.*

We already stated and proved this in Proposition 5.1.13.

Proposition A.1.4. *Suppose that $C \subseteq K$ is a relatively algebraically closed subfield of K which is not separably closed. There exists $V \subseteq K$ which is an $\exists$-C-definable bounded neighbourhood of 0 in the t-henselian topology.*

Proof. We choose $f \in C[x]$ to be non-linear, irreducible, and separable. Let $n := \deg(f)$; thus $\deg(Df) \leq n - 1$. If $|C| > n - 1$ then we may choose $a \in C \setminus Z(Df)$. On the other hand, if C is a finite field, then C is not square-closed. So we may choose f to be of degree 2; then Df is of degree 1 and again there exists $a \in C$ which is not a root of Df. Let $V := U_{f,a} = f(K)^{-1} - f(a)^{-1}$. Clearly V is $\exists$-C-definable. As discussed in Lemma A.1.3, V is a bounded neighbourhood of 0. □

A.1.3 An ∃-definable set between $\mathcal{O}$ and $\mathcal{M}$ in $F((t))$

Let $K := F((t))$ be the field of formal power series over the field F. Let $\mathcal{O} := F[[t]]$ be the valuation ring of the t-adic valuation, let $\mathcal{M} := t\mathcal{O}$ be its maximal ideal, and let vK be its value group. Note that $(K, \mathcal{O})$ is henselian and that $vK = \mathbb{Z}$. Let $C \subseteq K$ be any subfield.

In our analysis of F-definability in $F((t))$, we characterised the F-orbits of elements of $K \setminus F$: from Proposition 10.1.14, the orbit of $b \in \mathcal{M} \setminus \{0\}$ is $f_b(\mathcal{P})$. The next lemma makes use of this characterisation.

Proposition A.1.5. *Suppose that $V \subseteq K$ is an $\exists$-C-definable bounded neighbourhood of 0.*

1. *There exists $W \subseteq K$ which is bounded, $\exists$-$(C \cup F)$-definable, and is such that $\mathcal{P} \subseteq W$.*
2. *There exists $X \subseteq K$ which is bounded, $\exists$-$(C \cup F)$ definable, and is such that $\mathcal{M} \subseteq X$.*

3. *There exists $Y \subseteq K$ which is bounded by $\mathcal{O}$, $\exists$-$(C \cup F)$-definable, and is such that $\mathcal{M} \subseteq Y$.*

Proof. 1. V is a neighbourhood of 0. Let $\alpha \in \mathbb{Z}$ be such that $B(\alpha; 0) \subseteq V$. The set $B(\alpha; 0) \cap (\mathcal{M} \setminus \{0\})$ is infinite, thus we may choose any $a \in B(\alpha; 0) \cap (\mathcal{M} \setminus \{0\})$. We now apply our results about F-orbits of elements of $\mathcal{M} \setminus \{0\}$. By Proposition 10.1.14, there exists a polynomial $f \in F[x]$ such that f has no constant term and $f(\mathcal{P}) = \mathrm{Orb}(a/F) \subseteq V$. Let $\phi(x) := f(x) \in V$ and let $W := \phi(K)$ be the set defined by ϕ in K. Note that W is $\exists$-$(C \cup F)$-definable, and $\mathcal{P} \subseteq W$.

It remains to show that W is bounded. Let $\beta \in \mathbb{Z}$ be such that $V \subseteq B(\beta; 0)$. Let $\beta' := \min\{\beta, -1\}$, and write $f := \sum_{i=1}^{m} c_i x^i$. Let $b \notin B(\beta'; 0)$. Since $vb \leq \beta' \leq -1 < 0$, we have that $i > j$ implies that $v(b^i) < v(b^j)$. Thus $v(f(b)) = v(c_m b^m) = mv(b) \leq v(b) \leq \beta' \leq n$. Since $V \subseteq B(n; 0)$, we have that $f(b) \notin V$. Thus

$$(f(b) \in V \implies b \in B(\beta'; 0)).$$

So $W \subseteq B(\beta'; 0)$.

2. Let $W' := W \cup \{0\}$ and set $X := W - W'$. Clearly X is bounded and $\exists$-$(C \cup F)$-definable. By Lemma A.1.2, we see that $B(1; 0) \subseteq S(1) - S(1) = \mathcal{P} - \mathcal{P} \subseteq W - W \subseteq X$. Also $\mathcal{P} \subseteq W - 0 \subseteq X$. Thus $\mathcal{M} = \bar{B}(1; 0) = \mathcal{P} \sqcup B(1; 0) \subseteq X$.

3. X is bounded but contains $\mathcal{M}$, so there exists $n \in \mathbb{N}$ such that $X \subseteq B(-n; 0)$. Let $\psi(x) := x^n \in X$ and set $Y := \psi(K) - \psi(K)$. Observe that Y is $\exists$-$(C \cup F)$-definable. It remains to show that Y is bounded by $\mathcal{O}$ and that $\mathcal{M} \subseteq Y$.

   If $v(a) \leq -1$ then $v(a^n) = nv(a) \leq -n$. Thus if $v(a) \leq -1$, then $a^n \notin B(-1, 0) \supseteq X$ and $a \notin \psi(K)$. Therefore $\psi(K) \subseteq \mathcal{O}$. By Lemma A.1.2, $Y = \psi(K) - \psi(K) \subseteq \mathcal{O} - \mathcal{O} = \mathcal{O}$.

   Since $\mathcal{P}^n \subseteq S(n)$ (where $\mathcal{P}^n$ is the set of n-th powers of elements of $\mathcal{P}$) and $S(n) \subseteq \mathcal{M} \subseteq X$; we have that $\mathcal{P} \subseteq \psi(K)$. Thus $\mathcal{P} - \mathcal{P} \subseteq \psi(K) - \psi(K)$. By Lemma A.1.2, $B(1; 0) \subseteq \mathcal{P} - \mathcal{P}$; thus $B(1; 0) \subseteq \psi(K) - \psi(K)$. Since $0^n = 0 \in \mathcal{M} \subseteq X$, $0 \in \psi(K)$ and $\mathcal{P} - 0 \subseteq \psi(K) - \psi(K)$. By another application of Lemma A.1.2, this means that $\mathcal{M} = \mathcal{P} \sqcup B(1; 0) \subseteq \psi(K) - \psi(K) = Y$, as required.

□

A.1.4 An $\exists$-$\mathbb{F}_q$-definition of $\mathbb{F}_q[[t]]$ in $\mathbb{F}_q((t))$

Finally, we consider the special case where F is a finite field. Let q be a prime power. We fix $K := \mathbb{F}_q((t))$. Thus $\mathcal{O} = \mathbb{F}_p[[t]]$.

Proposition A.1.6. *There exists an $\exists$-$\mathbb{F}_q$-definable bounded neighbourhood of 0.*

Proof. $\mathbb{F}_q \subseteq K$ is relatively algebraically closed in K and is not separably closed. By Proposition A.1.4 there exists V with the required properties. □

Proposition A.1.7. *$\mathcal{O}$ is $\exists$-$\mathbb{F}_q$-definable in K.*

Proof. We combine Proposition A.1.6 and Proposition A.1.5 to obtain an $\exists$-$\mathbb{F}_q$-definable set Y which contains $\mathcal{M}$ and is bounded by $\mathcal{O}$. Note that $\mathbb{F}_q$ is an algebraic set defined by the formula $x^q - 1 \doteq 0$ in K. Let $\chi(x) := \exists y(y^q - 1 \doteq 0 \wedge x \in y + Y)$. This is obviously an $\exists$-formula with parameters from $\mathbb{F}_q$. Since $\mathcal{O} = \mathbb{F}_q + \mathcal{M}$ and $\mathcal{M} \subseteq Y \subseteq \mathcal{O}$, it is clear that $\chi(K) = \mathcal{O}$. □

A.1.5 An $\exists$-$\emptyset$-definition of $\mathbb{F}_q[[t]]$ in $\mathbb{F}_q((t))$

We can very marginally improve Proposition A.1.7 by removing the parameters. In the definition of the set $U_{f,a}$ we used a and the coefficients of f as parameters. All of these come from $\mathbb{F}_q$, but not necessarily from $\mathbb{F}_p$. Although elements of $\mathbb{F}_q$ are not closed terms, they are algebraic over $\mathbb{F}_p$. We use this algebraicity and a few simple tricks to find an existential formula with no parameters which defines $\mathcal{O}$.

Fact A.1.8. We state a simple consequence of Euclid's famous argument about the infinitude of the primes. Let $\{p_i | i \in I\}$ be a finite set of primes. There exists another prime $p' \leq \prod_{i \in I} p_i + 1$ which is not in the set $\{p_i | i \in I\}$.

Let $k \in \mathbb{N}$ and let P be the set of primes that divide k. Of course $\prod_{p \in P} p \leq k$. By the previous remark, there exists another prime $p' \notin P$ such that $p' \leq \prod_{p \in P} p + 1$. If $p' > k$ then $k = \prod_{p \in P} p$ and $p' = k + 1$. Thus $p' \leq k + 1$. Thus the least prime p' not dividing a natural number k is no greater than $k + 1$. Of course $k + 1$ is a very bad upper bound for p' in general, although if $k = 1, 2$ then $p' = k + 1$.

Proposition A.1.9. *There exists an $\exists$-$\emptyset$-definable bounded neighbourhood of 0.*

Proof. We seek a polynomial $f \in \mathbb{F}_p[x]$ which is separable, irredicible in $\mathbb{F}_q[x]$, and is such that $x^q - x \nmid Df$, i.e. not all elements of $\mathbb{F}_q$ are roots of Df.

Write $q = p^k$ and let l be the least prime not dividing k. By Fact A.1.6, $l \leq k+1$; consequently $l \leq p^k = q$. Since $\mathbb{F}_{p^l}/\mathbb{F}_p$ is a finite separably algebraic extension, there exists $b \in \mathbb{F}_{p^l}$ such that $\mathbb{F}_{p^l} = \mathbb{F}_p(b)$. Let $f \in \mathbb{F}_p[x]$ be the minimal polynomial of b over $\mathbb{F}_p$. Then f is separable, irreducible in $\mathbb{F}_p[x]$, and of degree l.

Since l and k are coprime, $\mathbb{F}_{q^l} = \mathbb{F}_{p^{kl}}$ is the least field containing $\mathbb{F}_q$ and $\mathbb{F}_{p^l}$. Thus $\mathbb{F}_q(b) = \mathbb{F}_{q^l}$, so the degree of b over $\mathbb{F}_q$ is still l. Therefore f is irreducible in the ring $\mathbb{F}_q[x]$.

Next we observe that the degree of Df is $\leq l-1 < q$. Since there are q elements of $\mathbb{F}_q$, not all of them can be roots of Df. Let $a \in \mathbb{F}_q$ be any element which is not a zero of Df. Then $U_{f,a}$ is a bounded neighbourhood of zero, however it would use the parameter a in its definition. The intersection of finitely many bounded neighbourhoods is still a bounded neighbourhood. Using this simple trick, we can avoid the use of parameters.

Let $A \subseteq \mathbb{F}_q$ be the set of elements of $\mathbb{F}_q$ which are not zeros of Df. Then A is $\emptyset$-definable using the formula $(x^q - x \doteq 0 \wedge \neg Df(x) \doteq 0)$. Using this definition of A we see that the set $V := \bigcap_{a \in A} U_{f,a}$ is an $\exists$-$\mathbb{F}_p$-definable bounded neighbourhood of zero, as required. We can replace elements of $\mathbb{F}_p$ by closed terms, thus eliminating all parameters from the definition of V. □

Theorem A.1.10. *$\mathcal{O}$ is $\exists$-$\emptyset$-definable in K.*

Proof. From Proposition A.1.9 we obtain an $\exists$-$\emptyset$-definable bounded neighbourhood of 0. If we look into the proof of Proposition A.1.5 (part 1), we see that the only remaining parameters come from the polynomial f. But in fact we could have chosen f without parameters: V is a neighbourhood of zero, so there exists $m \in \mathbb{N}$ such that $\mathcal{P}^m \subseteq V$ and we may set $f := x^m$. Thus we obtain an $\exists$-$\emptyset$-definable set Y which contains $\mathcal{M}$ and is bounded by $\mathcal{O}$. We define χ as before: let $\chi(x) := \exists y(y^q - 1 \doteq 0 \wedge x \in y + Y)$. This is an $\exists$-formula with no parameters and it defines $\mathcal{O}$. □

Nevertheless the formula still depends on $\mathbb{F}_q$ in several ways: our choices of f, a, m in the preceeding two proofs depend on $\mathbb{F}_q$, as does the choice of n in part 3 of Proposition A.1.5, q also appears directly in the definition of χ. All these factors tell us that χ is highly non-uniform in q. In fact, in recent as-yet-unpublished joint work of Cluckers, Derakhshan, Leenknegt, and Macintyre ([4]) it is shown that no definition exists which is uniform in p or in k (where $q = p^k$).

A.1.6 An $\exists$-$\emptyset$-definition of $\mathbb{F}_q[[t]]^{\mathrm{perf}}$ in $\mathbb{F}_q((t))^{\mathrm{perf}}$

We still denote $K := \mathbb{F}_q((t))$. We can use Theorem A.1.10 to existentially define the valuation ring $\mathcal{O}^{\mathrm{perf}} = \bigcup_{n<\omega} \mathcal{O}^{p^{-n}}$ in K^{perf}.

Theorem A.1.11. *$\mathcal{O}^{\mathrm{perf}}$ is $\exists$-$\emptyset$-definable in K^{perf}.*

Proof. Let χ be an $\exists$-formula (with no parameters) which defines $\mathcal{O}$ in K. The Frobenius map is an isomorphism, thus in each of the fields $K^{p^{-n}}$ the formula χ defines $\mathcal{O}^{p^{-n}}$. This now just follows from basic considerations of existential formulas: in any union of structures an existential formula defines the unions of sets that it defines in each of the structures. Thus χ defines $\mathcal{O}^{\mathrm{perf}}$ in K^{perf}. □

A.2 Consequences for $\exists$-definability in $\mathcal{L}_{\mathrm{val}}$

We return to imperfect fields and focus on the case that $q = p$ is a prime. Let $K := \mathbb{F}_p((t))$. In this case, the elements of $\mathbb{F}_p$ are the images of closed terms. Thus the $\exists$-$\mathbb{F}_p$-definition of $\mathcal{O}$ obtained in Proposition A.1.7 is already an $\exists$-$\emptyset$-definition.

The most important consequence of Proposition A.1.7 is that questions of existential definability in $\mathcal{L}_{\mathrm{val}}$ reduce to questions of existential definability in $\mathcal{L}_{\mathrm{ring}}$. Let $C \subseteq \mathbb{F}_p((t))$ be any subfield of parameters.

Proposition A.2.1. *Let $\phi \in \mathcal{L}_{\mathrm{val}}$ be an existential formula with parameters from C. Then there exists $\psi \in \mathcal{L}_{\mathrm{ring}}$ with parameters in C such that ϕ and ψ are equivalent modulo the theory of $\mathbb{F}_p((t))$.*

Proof. By Proposition A.1.7, $\mathcal{O}$ is $\exists$-$\emptyset$-definable. Set $\mathcal{O}^* := \mathcal{O} \setminus \{0\}$. Then $\mathcal{O}^{*-1}$ is also $\exists$-$\emptyset$-definable; consequently $\mathcal{M} = \mathbb{F}_q((t)) \setminus \mathcal{O}^{*-1}$ is $\forall$-$\emptyset$-definable. Finally $\mathcal{O}$ is $\forall$-$\emptyset$-defined by the formula $\bigvee_{a<p} a + x \in \mathcal{M}$ (which is made first-order by replacing the symbols a with the sum $1 + ... + 1$ of a-many of the constant symbol 1).

Since $\mathcal{O}$ is both $\forall$-$\emptyset$-definable and $\exists$-$\emptyset$-definable, we may convert any $\exists$-C-formula of $\mathcal{L}_{\mathrm{val}}$ into an $\exists$-C-formula of $\mathcal{L}_{\mathrm{ring}}$. □

Appendix B

Existential definability in large fields

A field L is *large* if every smooth curve defined over L either has no or infinitely many L-rational points. Fehm studied existential definability in perfect large fields in [9]. Among other things, he proved that infinite existentially definable subsets are not contained in any proper subfield. Another way of putting this is to say that any infinite existentially definable subset generates the whole field. This obvious fails in imperfect large fields (since the p-th powers are always existentially definable), but in this chapter we prove the best-possible alternative: infinite existentially definable subsets generate subfields which contain the p^n-th powers, for some $n \in \mathbb{N}$.

Many thanks to Arno Fehm for many very useful conversations about this and other topics.

In [22] it was shown that a field L is large if and only if $L \preceq_\exists L((t))$. Let L be a large field of exponential characteristic p. We will use this existential closure property to 'pull down' the results of chapter 7 from the henselian field $L((t))$ to L. Although we lose the topological information (since there is not necessarily any nice topology on L), we still have the results about cardinality and subfields.

Theorem B.0.2. *Let $X \subseteq L$ be an infinite $\exists$-definable subset. There exist $m, n \in \mathbb{N}$ such that $L^{p^n} \subseteq (X)_m$.*

Proof. Let ϕ be the existential formula defining X in L. Let X' be the set defined by the formula ϕ

in the henselian field $L((t))$. Since existential formulas go up, X' is infinite. Applying Theorem 7.9.3, there exist $m, n \in \mathbb{N}$ such that $(L((t)))^{p^n} \subseteq (X')_m$. Note that we may existentially define $(X')_m$ in $L((t))$ by a formula ϕ, which will also define $(X)_m$ in L. Now let $a \in L^{p^n}$; then $a \in (L((t)))^{p^n}$ so that $a \in (X')_m$. Since $L \preceq_\exists L((t))$, $a \in (X)_m$. Thus $L^{p^n} \subseteq (X)_m$, as required. □

Corollary B.0.3. *Let $X \subseteq L$ be an infinite $\exists$-definable subset. Then $|X| = |L|$.*

Proof. By Theorem B.0.2, $L^{p^n} \subseteq (X)_m$, for some $m, n \in \mathbb{N}$. Thus $|L| = |L^{p^n}| = |(X)_m|$, but it should be clear that $|(X)_m| = |X|$, since each step in the generation of a field is obtained by adjoining the image of a finite number of rational functions. □

Corollary B.0.4. *Let $C \subseteq L$ be any subfield and let $D := \boldsymbol{\Lambda}(L/C)$ be the Λ-closure of C in L. The existentially-definable and existentially-algebraic closures of C in L contain D and are contained in the relative algebraic closure of D in L.*

Proof. It is clear that elements in D are existentially definable over C. Suppose that a is transcendental over D and contained in the set existentially C-defined by the formula ϕ. Since L is Λ-closed in $L((t))$ (since the extension is regular, or by direct calculation), D is Λ-closed in $L((t))$. Applying Theorem 7.3.3, we have that ϕ must define (in $L((t))$) an infinite set. Since $L \preceq_\exists L((t))$, ϕ must also define an infinite set in L. Thus a is not existentially-algebraic over C. □

Bibliography

[1] James Ax. On the undecidability of power series fields. *Proc. Amer. Math. Soc.*, 16:846, 1965.

[2] James Ax and Simon Kochen. Diophantine problems over local fields i, ii. *Amer. J. Math.*, 87:605–648, 1965.

[3] Gregory L. Cherlin. Undecidability of rational function fields in nonzero characteristic. In *Logic colloquium '82 (Florence, 1982)*, volume 112 of *Stud. Logic Found. Math.*, pages 85–95, 1984.

[4] Raf Cluckers, Jamshid Derakshan, Eva Leenknegt, and Angus Macintyre. Uniformly defining valuation rings in henselian valued fields with finite and pseudo-finite residue field. To appear.

[5] Paul Cohen. Decision procedures for real and p-adic fields. *Comm. on Pure and Applied Math.*, 22:131–151, 1969.

[6] Jan Denef. p-adic semi-algebraic sets and cell decomposition. *J. Reine Angew. Math.*, 369:154–166, 1986.

[7] James K. Deveney and John N. Mordeson. Subfields and invariants of inseparable field extensions. *Can. J. Math.*, 29(6):1304–1311, 1977.

[8] Antonio J. Engler and Alexander Prestel. *Valued Fields*. Springer Monographs in Mathematics. Springer, 2005.

[9] Arno Fehm. Subfields of ample fields. rational maps and definability. *J. of Algebra*, 323:1738–1744, 2010.

[10] Jochen Koenigsmann. Elementary characterization of fields by their absolute galois group. *Siberian Adv. Math.*, 14:16–42, 2004.

[11] Hanspeter Kraft. Inseparable Körpererweiterungen. *Comment. Math. Helv.*, 45:110–118, 1970.

[12] Franz-Viktor Kuhlmann. Elementary properties of power series fields over finite fields. *J. Symb. Logic*, 66:771–791, 2001.

[13] Saunders Mac Lane. Modular fields. I. Separating transcendence bases. *Duke Math. J.*, 5:372–393, 1939.

[14] Saunders Mac Lane. Steinitz field towers for modular fields. *Trans. of the Amer. Math. Soc.*, 46:23–45, 1939.

[15] Saunders Mac Lane. Modular fields. *The Amer. Math. Monthly*, 47:259–274, 1940.

[16] Saunders Mac Lane. *Categories for the Working Mathematician*. Graduate Texts in Mathematics. Springer, second edition, 1998.

[17] Serge Lang. *Algebra*. Springer, third edition, 1987.

[18] Angus Macintyre. On definable subsets of p-adic fields. *J. Symbolic Logic*, 41:605–610, 1976.

[19] Angus Macintyre. Twenty years of p-adic model theory. *Stud. Logic Found. Math.*, 120:121–153, 1986.

[20] David Marker. *Model Theory: An Introduction*. Graduate Texts in Mathematics. Springer, 2000.

[21] J. N. Mordeson and B. Vinograde. Separating p-bases and transcendental extension fields. *Proc. Amer. Math. Soc.*, 31:417–422, 1972.

[22] Florian Pop. Embedding problems over large fields. *Ann. of Math.*, 144:1–34, 1996.

[23] Alexander Prestel. Algebraic number fields elementarily determined by their absolute Galois group. *Israel J. Math.*, 73(2):199–205, 1991.

[24] Alexander Prestel and Peter Roquette. *Lectures on formally p-adic fields*, volume 1050 of *Lecture Notes in Mathematics*. Springer, 1984.

[25] Alexander Prestel and Martin Ziegler. Model-theoretic methods in the theory of topological fields. *J. Reine Angew. Math.*, 299 (300):318–341, 1978.

[26] Thomas Rohwer. *Valued Difference Fields as Modules over Twisted Polynomial Rings*. PhD thesis, 2003.

[27] O. F. G. Schilling. Automorphisms of fields of formal power series. *Bull. Amer. Math. Soc.*, 50:892–901, 1944.

[28] F. K. Schmidt. Mehrfach perfekte Körper. *Math. Ann.*, 108:1–25, 1933.

[29] Alfred Tarski. Arithmetical classes and types of algebraically closed and real-closed fields. *Bull. Amer. Math. Soc.*, 55:63–64, 1949.

[30] O. Teichmüller. p-Algebren. *Deutsche Mathematik*, pages 362–388, 1936.