

Lecture Notes

Algebraic Difference Equations



Michael Wibmer
February 6, 2013

Contents

1	Basic notions of difference algebra	7
1.1	Difference rings	7
1.1.1	Quotients	9
1.1.2	Localization	9
1.1.3	Tensor products	10
1.1.4	Difference polynomials	10
1.1.5	Adjunction of elements	11
1.1.6	Inversive closure	11
1.1.7	From algebras to difference algebras	13
1.2	Difference ideals	16
1.2.1	Properties of difference ideals	16
1.2.2	Perfect difference ideals	18
1.2.3	Well-mixed difference ideals	20
1.2.4	The Cohn topology	22
2	Difference varieties	26
2.1	Difference varieties and their coordinate rings	26
2.2	The topological space of a difference variety	30
2.3	Morphisms of difference varieties	32
3	Around the basis theorem	34
3.1	Characteristic sets	35
3.2	Difference algebras of finite presentation	39
3.3	Ritt difference rings and the basis theorem	40
3.4	Application to difference varieties	42
4	Extensions of difference fields	44
4.1	Difference algebraic extensions	44
4.1.1	Application to difference varieties	45
4.2	Difference transcendence bases	47
4.2.1	Application to difference varieties	50
4.3	The limit degree	51
4.4	Finitely generated extensions of difference fields	55
4.5	The core and Babbitt's decomposition	56
4.6	Compatibility	65
5	Difference kernels	69
5.1	The difference degree	69
5.2	Prolongations and realizations of difference kernels	73
5.2.1	Application to difference varieties	76

Motivation

So what is an algebraic difference equation? Surely you have already seen some examples: E.g., the functional equation of the gamma function

$$\Gamma(z + 1) = z\Gamma(z), \quad (1)$$

or the recurrence formula for the Fibonacci numbers

$$a_{n+2} = a_{n+1} + a_n. \quad (2)$$

These are very simple examples of algebraic difference equations, in particular these are linear difference equations. The second one even has constant coefficients. Equation (1) has order one while equation (2) has order two. An example of a first order non-linear algebraic difference equation would be the so-called logistic map

$$a_{n+1} = ra_n(1 - a_n) \quad (3)$$

(r a positive real number) which is a popular population model. Another non-linear example in the spirit of (1) would be

$$(f(z + 1) - f(z))^2 - 2(f(z + 1) + f(z)) + 1 = 0 \quad (4)$$

where f , the function to be found, is a function of a complex variable z . An example of a difference equation which is not algebraic would be something like

$$a_{n+1} = \sin(a_n).$$

Solving difference equations in explicit terms (closed form solutions) is very hard and in general impossible. For example, even for the relatively simple equation (3) it seems that only for $r = 2$ and $r = 4$ explicit solutions are known. (For $r = 2$, $a_n = \frac{1}{2} - \frac{1}{2}(1 - 2a_0)^{2^n}$ and for $r = 4$, $a_n = \sin(2^n \theta \pi)^2$ with $\theta = \frac{1}{\pi} \arcsin(\sqrt{a_0})$.)

In this course we are not really interested in “solving”. For example, starting with equation (1) we are not interested in finding analytic representations of the gamma function, like the product formula. Our perspective is rather to study the solutions and their properties from an abstract, purely algebraic point of view. The resulting theory, usually referred to as “difference algebra” enjoys many analogies with commutative algebra and algebraic geometry and is essentially devoid of analytic considerations.

The basic algebraic concept which allows us to treat equations of such different types as (1) and (2) simultaneously is the notion of a difference ring. By definition, a difference ring is simply a commutative ring R equipped with a ring endomorphism $\sigma: R \rightarrow R$. If R is a field, we speak of a difference field.

For equations (2) and (3) the underlying difference ring is the ring R of sequences in \mathbb{C} , (or \mathbb{R} if you prefer) with componentwise addition and multiplication and $\sigma: R \rightarrow R$ the shift, defined by $\sigma((a_n)_{n \in \mathbb{N}}) = (a_{n+1})_{n \in \mathbb{N}}$. We can rewrite equations (2) and (3) as

$$\sigma^2(y) - \sigma(y) - y = 0$$

and

$$\sigma(y) - ry(1 - y) = 0$$

respectively.

For equations (1) and (4) we can consider the field \mathcal{M} of meromorphic functions on \mathbb{C} equipped with the endomorphism σ given by $\sigma(f(z)) = f(z + 1)$ and rewrite (1) and (4) as

$$\sigma(y) - zy = 0$$

and

$$(\sigma(y) - y)^2 - 2(\sigma(y) + y) + 1 = 0$$

respectively. Here y is considered as a formal variable and an expression like the above is called a difference polynomial. The expression $\sigma(y) - zy$ can be considered as a difference polynomial with coefficients in the difference field $\mathbb{C}(z)$, the field of rational functions in z , with action of σ also given by $\sigma(f(z)) = f(z + 1)$. Then $\Gamma \in \mathcal{M}$ is a solution of $\sigma(y) - zy$ in the difference field extension \mathcal{M} of $\mathbb{C}(x)$. This is somewhat analogous to saying that $\sqrt{2} \in \mathbb{R}$ is a solution of $y^2 - 2$ in the field extension \mathbb{R} of \mathbb{Q} .

Other examples of classical interest are also covered by the difference algebra formalism: For example q -difference equations, where σ acts by $\sigma(f(z)) = f(qz)$ with q a non-zero complex number; or Mahler difference equations where σ acts by $\sigma(f(z)) = f(z^d)$ for some integer $d \geq 2$.

Note that the difference polynomial $(\sigma(y) - y)^2 - 2(\sigma(y) + y) + 1$ is irreducible when considered as polynomial in the two variables $y, \sigma(y)$ (say, over \mathbb{C}). Nevertheless $(\sigma(y) - y)^2 - 2(\sigma(y) - y) + 1$ has two families of solutions inside \mathcal{M} : One given by $f_1(z) = (z + c)^2$ and another one given by $f_2(z) = (ce^{i\pi z} + \frac{1}{2})^2$ where $c \in \mathcal{M}$ denotes an arbitrary 1-periodic function, i.e., $\sigma(c) = c$. Note that f_1 satisfies the difference polynomial $\sigma^2(y) - 2\sigma(y) + y - 2$ while f_2 satisfies the difference polynomial $\sigma^2(y) - y$. These difference polynomials can be obtained from $p := (\sigma(y) - y)^2 - 2(\sigma(y) - y) + 1$ (without “solving”!) by factoring $\sigma(p) - p$. Indeed,

$$\begin{aligned} \sigma(p) - p &= (\sigma^2(y) - \sigma(y))^2 - 2(\sigma^2(y) + \sigma(y)) + 1 - ((\sigma(y) - y)^2 - 2(\sigma(y) + y) + 1) = \\ &= \sigma^2(y)^2 - 2\sigma(y)\sigma^2(y) - 2\sigma^2(y) + 2y\sigma(y) - y^2 - 2y = \\ &= (\sigma^2(y) - 2\sigma(y) + y - 2) \cdot (\sigma^2(y) - y). \end{aligned}$$

Eventually we will be able to understand this phenomenon from a purely algebraic perspective: The perfect difference ideal generated by p inside the difference polynomial ring in the difference variable y has two prime components. Perfect difference ideals play a similar role in difference algebra as radical ideals in commutative algebra and algebraic geometry. They are precisely the ideals which can be “recovered” from their solutions.

So far we have only encountered univariate difference polynomials, an example of a difference polynomial in two difference variables y_1, y_2 would be something like

$$y_1\sigma^3(y_1)^2\sigma^4(y_2)^5 + 3\sigma^2(y_1)y_2 + y_2^2.$$

The main focus of this course is to study the set of solutions of a general system of algebraic difference polynomials in finitely many difference variables over an arbitrary difference field. These solution sets are called difference varieties.

We will address questions like: Can we replace an infinite system of algebraic difference equations by a finite system without changing the solutions? Can we decompose a system of algebraic equations into finitely many “irreducible” systems? How can we measure the size of a difference variety, e.g., is there something like a difference dimension?

The standard references for difference algebra are

- [Cohn] R. Cohn, *Difference Algebra*, Interscience Publishers, 1965.
- [Levin] A. Levin, *Difference Algebra*, Springer, 2008.

A somewhat more advanced perspective is in

- [Hrushovski] E. Hrushovski, *The elementary theory of the Frobenius automorphisms*, 2004, available from <http://www.ma.huji.ac.il/~ehud/>.

Exercise 0.0.1. *Using the notation from above: Solve the difference equation $(f(z+1) - f(z))^2 - 2(f(z+1) + f(z)) + 1 = 0$. I.e., starting from this equation, find the solutions f_1 and f_2 .*

Please report mistakes, typos or comments to

`michael.wibmer@matha.rwth-aachen.de`

Conventions

All rings are assumed to be commutative and unital. The set of natural numbers contains zero, i.e., $\mathbb{N} = \{0, 1, 2, \dots\}$. We set $\mathbb{N}_{\geq 1} = \{1, 2, \dots\}$. If $\sigma: M \rightarrow M$ is a map of sets than σ^0 is, by definition, the identity map on M . A subset N of M is called σ -stable if $\sigma(N) \subset N$. We denote by

$$\sigma^{-1}(N) = \{m \in M \mid \sigma(m) \in N\}$$

the inverse image of N under σ .

If R is a ring and \mathfrak{p} a prime ideal of R , we denote the residue field at \mathfrak{p} by

$$k(\mathfrak{p}) = R_{\mathfrak{p}}/\mathfrak{p}_{\mathfrak{p}} = \text{Quot}(R/\mathfrak{p}).$$

If F and G are subsets of a ring R we denote by

$$FG = \{fg \mid f \in F, g \in G\}$$

the set of all products of elements from F and G . However, if $L|K$ and $M|K$ are field extensions, contained in some field extension of K , we also use LM to denote the field compositum of L and M . (This should not be a source of confusion.)

Chapter 1

Basic notions of difference algebra

In this chapter we introduce the very basic definitions and constructions of difference algebra and we establish some first theorems on difference ideals. We are not yet concerned with difference varieties. Giving a course in algebraic geometry has the advantage that one can assume some knowledge about basic algebra. Of course we do not assume any knowledge about difference algebra. So all our geometric considerations will have to be accompanied by a fair bit of algebra.

1.1 Difference rings

In this section we introduce basic objects, like difference rings, difference ideals, difference fields, difference polynomials,.... Several of the standard constructions from commutative algebra (e.g., quotients, localization, tensor products,...) are shown to have meaningful difference analogs.

Definition 1.1.1. A difference ring, or σ -ring for short, is a pair (R, σ) where R is a ring and $\sigma: R \rightarrow R$ a morphism of rings.

Recall that “ring” always means “commutative ring”. If the ring R is a field, we say that (R, σ) is a σ -field. For ease of notation we will usually write R instead of (R, σ) . For brevity we will often use the prefix “ σ ” in lieu of “difference”. Note that σ is automatically injective if R is a field. Some authors also require σ to be injective in the definition of difference ring. We do not make this requirement.

Definition 1.1.2. A morphism between σ -rings R and S is a morphism of rings $\phi: R \rightarrow S$ such that $\phi\sigma = \sigma\phi$. In this situation we also say that S is an R - σ -algebra.

It usually does not cause confusion to use the same letter σ for different endomorphisms. For example, in the above definition the σ in $\phi\sigma$ is the endomorphism of R and the σ in $\sigma\phi$ is the endomorphism of S .

If S and R are σ -rings such that R is a subring of S . Then we say that R is a σ -subring of S (or $R \subset S$ is an inclusion of σ -rings) if the inclusion map $R \rightarrow S$ is a morphism of σ -rings. If, in this situation, R and S are fields, we say that S is a σ -field extension of R . The notion of R - σ -subalgebra is defined analogously. A morphism of R - σ -algebras is a morphism of R -algebras which is a morphism of σ -rings.

Example 1.1.3. (i) Any ring R can be considered as a difference ring by taking $\sigma: R \rightarrow R$ as the identity map. Such a σ -ring is called a *constant* σ -ring.

(ii) The only way to consider \mathbb{Z} and \mathbb{Q} as difference rings is to take σ as the identity map. So \mathbb{Z} and \mathbb{Q} are constant.

- (iii) If (R, σ) is a difference ring then (R, σ^d) is also a difference ring for every $d \in \mathbb{N}_{\geq 1}$.
- (iv) Let $K = \mathbb{C}(z)$ denote the field of rational functions in the variable z . We can consider K as σ -field by setting $\sigma(f(z)) = f(z+1)$. Difference equations over this base σ -field are sometimes called finite difference equations. If we set $\sigma(f(z)) = f(qz)$ where q is a non-zero complex number, then the resulting difference equations are called q -difference equations. If we set $\sigma(f(z)) = f(z^d)$, ($d \geq 2$ an integer) the corresponding difference equations are called Mahler difference equations.
Let \mathcal{M} denote the field of meromorphic functions on \mathbb{C} . Then \mathcal{M} is a field extension of K and if we equip \mathcal{M} with an action of σ by using the same formula as on K , then \mathcal{M} becomes a σ -field extension of K .
- (v) Let K be a field of characteristic $p > 0$ and $e \in \mathbb{N}_{\geq 1}$. Setting $\sigma(f) = f^{p^e}$ for $f \in K$ defines a difference field structure on K . Such a difference field is called a Frobenius difference field.
- (vi) Let k be a field, X an irreducible variety over k and $f: X \dashrightarrow X$ a dominant rational map¹. Then f induces a morphism $f^*: k(X) \rightarrow k(X)$ on the function field $k(X)$ of X . The pair $(k(X), f^*)$ is a difference field.
- (vii) There are two possible ways to turn $\mathbb{Q}(\sqrt{2})$ into a difference field. Either $\sigma(\sqrt{2}) = \sqrt{2}$ or $\sigma(\sqrt{2}) = -\sqrt{2}$. Obviously these two difference fields are not isomorphic. In fact, there does not exist a difference field which contains these two difference fields as difference subfields.
- (viii) Let S be a ring and $\text{Seq}_S := S^{\mathbb{N}}$ the ring of sequences in S . The ring structure is given by componentwise addition and multiplication. We can define an endomorphism of Seq_S by shifting a sequences one step to the left, i.e.,

$$\sigma((a_i)_{i \in \mathbb{N}}) := (a_{i+1})_{i \in \mathbb{N}}.$$

If we consider S as a constant σ -ring then the map $\phi: S \rightarrow \text{Seq}_S$ which sends an element $b \in S$ to the constant sequence $\phi(b) := (b, b, \dots)$ is a morphism of σ -rings.

- (ix) Let $\mathbb{C}[z]$ denote the polynomial ring in the variable z and consider $\mathbb{C}[z]$ as σ -ring via $\sigma(f(z)) = f(z+1)$. Then the map

$$\phi: \mathbb{C}[z] \rightarrow \text{Seq}_{\mathbb{C}}, f \mapsto (f(n))_{n \in \mathbb{N}}$$

is a morphism of difference rings.

- (x) Let R be a σ -ring. Then

$$\iota: R \rightarrow \text{Seq}_R, f \mapsto (\sigma^n(f))_{n \in \mathbb{N}}$$

is a morphism of difference rings. It is called the universal Euler morphism.

Exercise 1.1.4. Find and prove the universal property of the universal Euler morphism.

¹Never mind if you do not know precisely what these things are.

1.1.1 Quotients

Definition 1.1.5. Let R be a σ -ring and $\mathfrak{a} \subset R$ an ideal. Then \mathfrak{a} is called a σ -ideal if $\sigma(\mathfrak{a}) \subset \mathfrak{a}$

Exercise 1.1.6. Let $R = \mathbb{C}[z]$ with $\sigma(f(z)) = f(z+1)$. Determine all difference ideals of R .

It is easy to see that the kernel of a morphism of difference rings is a σ -ideal.

Lemma 1.1.7. Let \mathfrak{a} be a σ -ideal of a σ -ring R . Then there exists a unique difference structure on R/\mathfrak{a} such that the canonical map

$$R \rightarrow R/\mathfrak{a}, f \mapsto \bar{f}$$

is a morphism of difference rings.

Proof. Because \mathfrak{a} is a σ -ideal the map

$$\sigma: R/\mathfrak{a} \rightarrow R/\mathfrak{a}, \bar{f} \mapsto \overline{\sigma(f)}$$

is well-defined. □

In the sequel we will always consider R/\mathfrak{a} as a difference ring in the way explained above.

Definition 1.1.8. Let R be a σ -ring and F a subset of R . An intersection of σ -ideals containing F is a σ -ideal containing F . Therefore, there exists a unique smallest difference ideal containing F , which we call the difference ideal generated by F . It is denoted by $[F]$. A difference ideal \mathfrak{a} of R is called finitely σ -generated if there exists a finite subset F of \mathfrak{a} such that $\mathfrak{a} = [F]$.

Because the set of all R -linear combinations of elements of the form $\sigma^i(f)$, $f \in F$, $i \in \mathbb{N}$ is a σ -ideal, it is clear that the elements of $[F]$ are precisely the R -linear combinations of elements of the form $\sigma^i(f)$.

1.1.2 Localization

Recall that a subset S of a ring R is called multiplicatively closed if $1 \in S$ and $ss' \in S$ for $s, s' \in S$. We can then speak of the localization $S^{-1}R$ of R by S .

Lemma 1.1.9. Let R be a σ -ring and $S \subset R$ a multiplicatively closed, σ -stable subset of R . Then there exists a unique difference structure on $S^{-1}R$ such that the canonical map

$$R \rightarrow S^{-1}R, f \mapsto \frac{f}{1}$$

is a morphism of σ -rings.

Proof. Because S is σ -stable, the map

$$\sigma: S^{-1}R \rightarrow S^{-1}R, \frac{f}{s} \mapsto \frac{\sigma(f)}{\sigma(s)}$$

is well-defined. □

In the sequel we will always consider $S^{-1}R$ as a difference ring in the way explained above.

The intersection of multiplicatively closed, σ -stable subsets of R is multiplicatively closed and σ -stable. Therefore every subset F of R is contained in a unique smallest multiplicatively closed, σ -stable subsets of R , which we call the *multiplicatively closed, σ -stable subset of R generated by F* . Let S denote the multiplicatively closed, σ -stable subsets of R generated by a single element $f \in R$. Note that the elements of S are precisely the products of elements of the form $\sigma^i(f)$, $i \in \mathbb{N}$. We set

$$R\{1/f\} := S^{-1}R.$$

Exercise 1.1.10. *A σ -ring is called σ -simple if the only σ -ideals are the whole ring and the zero ideal. Show that the set of non-zero divisors of a σ -simple σ -ring is σ -stable.*

1.1.3 Tensor products

Lemma 1.1.11. *Let R be a σ -ring and let S, T be R - σ -algebras. Then there exists a unique difference structure on $S \otimes_R T$ such that the canonical maps $S \rightarrow S \otimes_R T$ and $T \rightarrow S \otimes_R T$ are morphisms of difference rings.*

Proof. We only have to make sure that the map

$$\sigma: S \otimes_R T \longrightarrow S \otimes_R T, f \otimes g \mapsto \sigma(f) \otimes \sigma(g)$$

is well-defined. To see this consider $S \otimes_R T$ as R -algebra via the map $R \xrightarrow{\sigma} R \rightarrow S \otimes_R T$. Then $S \rightarrow S \otimes_R T$, $f \mapsto \sigma(f) \otimes 1$ and $T \rightarrow S \otimes_R T$, $g \mapsto 1 \otimes \sigma(g)$ are morphisms of R -algebras. \square

In the sequel we will always consider the tensor products of difference algebras as a difference ring in the way explained above. The tensor product gives the coproduct in the category of algebras. It is clear that it also gives the coproduct in the category of difference algebras.

1.1.4 Difference polynomials

Let R be a σ -ring. The σ -polynomial ring $R\{y\}$ over R in the σ -variables $y = (y_1, \dots, y_n)$ is the polynomial ring over R in the variables² $\sigma^i(y_j)$, $i \in \mathbb{N}$, $j = 1, \dots, n$. That is,

$$R\{y\} = R\{y_1, \dots, y_n\} = R[\sigma^i(y_j) \mid i \in \mathbb{N}, j = 1, \dots, n].$$

Here – a priori – $\sigma^i(y_j)$ is simply the name of a new variable. However, we turn $R\{y\}$ into a difference ring by extending σ from R to $R\{y\}$ by

$$\sigma(\sigma^i(y_j)) := \sigma^{i+1}(y_j), \quad i \in \mathbb{N}, \quad j = 1, \dots, n.$$

This makes $R\{y\}$ into an R - σ -algebra. The σ -polynomial ring has the following universal property: If S is an R - σ -algebra and $a \in S^n$, then there exists a unique morphism $\phi: R\{y\} \rightarrow S$ of R - σ -algebras such that $\phi(y) = a$ (i.e., $\phi(y_i) = a_i$ for $i = 1, \dots, n$ where $a = (a_1, \dots, a_n)$). The map ϕ is called “evaluation at a ” and we write

$$f(a) := \phi(f) \text{ for } f \in R\{y\}.$$

If $f(a) = 0$ we say that a is solution of f , or that f vanishes on a . Note that $\sigma(f)(a) = \sigma(f(a))$.

²Recall that $\sigma^0(y_j) = y_j$ by definition.

Example 1.1.12. Let $k = \mathbb{C}(z)$ and \mathcal{M} the field of meromorphic functions on \mathbb{C} , both considered as σ -fields via $\sigma(f(z)) = f(z + 1)$. Then the gamma function $\Gamma \in \mathcal{M}$ is a solution of the σ -polynomial $\sigma(y_1) - zy_1 \in k\{y_1\}$.

The order of a σ -polynomial $f \in R\{y\} \setminus R$, denoted by

$$\text{ord}(f)$$

is the largest integer i such that some $\sigma^i(y_j)$ ($i \in \mathbb{N}, 1 \leq j \leq n$) effectively appears in f . The effective order of a σ -polynomial $f \in R\{y\} \setminus R$, denoted by

$$\text{Eord}(f)$$

is the difference $\text{ord}(f) - i$, where i denotes the smallest integer such that some $\sigma^i(y_j)$ ($i \in \mathbb{N}, 1 \leq j \leq n$) effectively appears in f .

Example 1.1.13. The order of $\sigma(y_2) + \sigma^3(y_1)^5\sigma^2(y_2)$ is 3, while the effective order is 2.

1.1.5 Adjunction of elements

Let R be a σ -ring, S an R - σ -algebra and A a subset of S . Obviously the intersection of R - σ -subalgebras of S is an R - σ -subalgebra. Therefore there exists a unique smallest R - σ -subalgebra of S which contains A . We denote it by

$$R\{A\} \subset S$$

and call it the R - σ -subalgebra of S generated by A . The structure of $R\{A\}$ is easy to describe: It is simply the R -subalgebra of S generated by $A, \sigma(A), \dots$. If A is finite, say $A = \{a_1, \dots, a_n\}$, then $R\{A\}$ is the image of the σ -polynomial ring $R\{y_1, \dots, y_n\}$ under the R - σ -morphism $y_i \mapsto a_i$ ($i = 1, \dots, n$). If there exists a finite subset A of S such that $S = R\{A\}$ we say that S is a finitely generated³ R - σ -algebra, or that S is finitely σ -generated over R . In this situation we also that A is a finite σ -generating set of S over R . It follows that S is finitely σ -generated over R if and only if we can write S as

$$S \simeq R\{y\}/\mathfrak{a}$$

for some σ -ideal \mathfrak{a} in some σ -polynomial ring $R\{y\}$ over R .

Let $L|K$ be an extension of σ -fields and A a subset of L . Clearly the intersection of σ -subfields of L is a σ -subfield of L . Therefore there exists a smallest σ -subfield of L which contains K and A . We denote it by

$$K\langle A \rangle \subset L$$

and call it the σ -subfield of L σ -generated by A over K . If $L = K\langle A \rangle$ for some finite subset A of L we say that L is finitely σ -generated over K (as σ -field extension of K).

1.1.6 Inversive closure

Definition 1.1.14. A σ -ring R is called *inversive* if $\sigma: R \rightarrow R$ is an automorphism (i.e., injective and surjective).

³I prefer “finitely σ -generated” to “finitely generated” to avoid confusion with “ S being finitely generated as R -algebra.”

Example 1.1.15. The base σ -field $\mathbb{C}(z)$ for finite difference equations ($\sigma(f(z)) = f(z+1)$) and for q -difference equations ($\sigma(f(z)) = f(qz)$) is inversive.

Non-example 1.1.16. The base σ -field $\mathbb{C}(z)$ for Mahler difference equations ($\sigma(f(z)) = f(z^d)$) is not inversive.

Example 1.1.17. A Frobenius difference field is inversive if and only if it is perfect.

It is sometimes easier to work with inversive σ -rings, or over an inversive σ -field. To pass to the inversive closure (defined below) is an important proof-technique.

Proposition 1.1.18. *Let R be a σ -ring. There exists an inversive σ -ring R^* and a morphism $\psi: R \rightarrow R^*$ of σ -rings satisfying the following universal property: For every morphism $\psi': R \rightarrow S$ with S inversive there exists a unique morphism $\phi: R^* \rightarrow S$ making*

$$\begin{array}{ccc} R & \xrightarrow{\psi} & R^* \\ & \searrow \psi' & \swarrow \phi \\ & & S \end{array} \quad (1.1)$$

commutative. The σ -ring R^ is unique up to unique isomorphisms.*

Proof. Consider the set M of all pairs (i, f) , where $i \in \mathbb{N}$ and $f \in R$. We say that two pairs (i, f) and (j, g) from M are equivalent if there exist $m, n \in \mathbb{N}$ such that

$$(i + m, \sigma^m(f)) = (j + n, \sigma^n(g)).$$

Obviously this defines an equivalence relation on M . The equivalence classes form naturally a σ -ring⁴ which we denote by R^* : To add the equivalence classes of (i, f) and (j, g) , choose $m, n \in \mathbb{N}$ such that $i + m = j + n$ and consider the equivalence class of $(i + m, \sigma^m(f) + \sigma^n(g))$. Multiplication is defined similarly. The action of σ on the equivalence class of (i, f) is given by considering the equivalence class of $(i, \sigma(f))$. It is clear that R^* is inversive and we have a natural morphism $\psi: R \rightarrow R^*$ which sends $f \in R$ to the equivalence class of $(0, f)$.

Now let $\psi': R \rightarrow S$ be a morphism from R to an inversive σ -ring S . Since S is inversive we can define a map $\phi: R^* \rightarrow S$ by sending the equivalence class of (i, f) to $\sigma^{-i}(\psi'(f)) \in S$. To see that this is well-defined we only have to observe that

$$\sigma^{-i}(\psi'(f)) = \sigma^{-(i+m)}(\psi'(\sigma^m(f))) = \sigma^{-(j+n)}(\psi'(\sigma^n(g))) = \sigma^{-j}(\psi'(g))$$

if $(i + m, \sigma^m(f)) = (j + n, \sigma^n(g))$. It is straight forward to check that ϕ is the unique morphism of σ -rings making (1.1) commutative. \square

Corollary 1.1.19. *The kernel of $\psi: R \rightarrow R^*$ is $\{f \in R \mid \exists n \in \mathbb{N} : \sigma^n(f) = 0\}$. In particular, if $\sigma: R \rightarrow R$ is injective, then ψ is injective and so we can consider R as a σ -subring of R^* .*

Proof. This is clear from the construction of ψ in the proof of Proposition 1.1.18. \square

Definition 1.1.20. *Let R be a σ -ring. The σ -ring R^* defined in Proposition 1.1.18 is called the inversive closure of R .*

⁴For those who know what is a direct limit: R^* is simply the direct limit of the direct system given by the maps $\sigma^{j-i}: R_i \rightarrow R_j$ for $i < j$ where R_i is an isomorphic copy of R .

Example 1.1.21. The inversive closure of a Frobenius difference field is the perfect closure.

If $R \rightarrow S$ is a morphism of σ -rings, then so is the composition $R \rightarrow S \rightarrow S^*$ and we obtain a morphism $R^* \rightarrow S^*$ of σ -rings. We conceive that “inversive closure” is a functor from the category of difference rings to the category of inversive difference rings.

Lemma 1.1.22. *If a σ -ring R is an integral domain, then also R^* is an integral domain. If K is a σ -field, then K^* is a σ -field.*

Proof. This is clear from the construction of the inversive closure in the proof of Proposition 1.1.18. \square

The following lemma gives another characterization of the inversive closure in case $\sigma: R \rightarrow R$ is injective.

Lemma 1.1.23. *Let R be a σ -ring. If R is a σ -subring of an inversive σ -ring S such⁵ that for every element $f \in S$ there exists $n \in \mathbb{N}$ with $\sigma^n(f) \in R$, then S is the inversive closure of R . Moreover, every inversive σ -ring containing R contains an inversive closure of R .*

Proof. We have to verify that S satisfies the universal property. Let $\psi: R \rightarrow S'$ be a morphism of σ -rings with S' inversive. We need to define $\phi: S \rightarrow S'$ making

$$\begin{array}{ccc} R^{\subset} & \xrightarrow{\quad} & S \\ & \searrow \psi & \swarrow \phi \\ & & S' \end{array} \tag{1.2}$$

commutative. For $f \in S$ there exists $n \in \mathbb{N}$ with $\sigma^n(f) \in R$. We can then set

$$\phi(f) = \sigma^{-n}(\psi(\sigma^n(f))).$$

Because ψ is a σ -morphism the definition does not depend of the choice of n . Clearly ϕ is the unique morphism of σ -rings making (1.2) commutative.

To verify the last claim of the lemma, we only have to observe that if R is a σ -subring of an inversive σ -ring S' , then $S := \{f \in S' \mid \exists n \in \mathbb{N} : \sigma^n(f) \in R\} \subset S'$ is an inversive σ -ring. By the first part of the lemma S is the inversive closure of R . \square

Exercise 1.1.24. *Let $k = \mathbb{C}(z)$ with $\sigma(f(z)) = f(z^d)$ for $d \geq 2$. Show that k^* is algebraic over k and that k^* is not finitely σ -generated as σ -field extension of k . Give an example of a σ -field k such that k^* is not algebraic over k .*

Exercise 1.1.25. *If $L|K$ is a (separably) algebraic extension of σ -fields then $L^*|K^*$ is a (separably) algebraic extension of σ -fields. Similarly, if $L|K$ is normal, then $L^*|K^*$ is normal.*

1.1.7 From algebras to difference algebras

Let k be a σ -field and R a k -algebra. We would like to enlarge R to a k - σ -algebra. The following proposition asserts that this can be done in a unique generic way.

⁵This implies that σ is injective on R .

Proposition 1.1.26. *Let k be a σ -field and let R be a k -algebra. There exists a k - σ -algebra $[\sigma]_k R$ together with a morphism $\psi: R \rightarrow [\sigma]_k R$ of k -algebras satisfying the following universal property: For every k - σ -algebra S and every morphism $\psi': R \rightarrow S$ of k -algebras there exists a unique morphism $\phi: [\sigma]_k R \rightarrow S$ of k - σ -algebras making*

$$\begin{array}{ccc} R & \xrightarrow{\psi} & [\sigma]_k R \\ & \searrow \psi' & \swarrow \phi \\ & S & \end{array} \quad (1.3)$$

commutative. The k - σ -algebra $[\sigma]_k R$ is unique up to unique isomorphisms.

Proof. For $i \in \mathbb{N}$ we set

$$\sigma^i R = R \otimes_k k,$$

where the tensor product is formed by using $\sigma^i: k \rightarrow k$ on the right hand side. We consider $\sigma^i R$ as k -algebra via the right factor. We have morphisms of rings⁶

$$\psi_i: \sigma^i R \longrightarrow \sigma^{i+1} R, f \otimes \lambda \mapsto f \otimes \sigma(\lambda).$$

If we set

$$R_i = R \otimes_k \sigma R \otimes_k \cdots \otimes_k \sigma^i R$$

we have natural inclusions

$$R_i \hookrightarrow R_{i+1}, f_0 \otimes \cdots \otimes f_i \mapsto f_0 \otimes \cdots \otimes f_i \otimes 1$$

of k -algebras and ring morphisms⁷

$$\sigma_i: R_i \rightarrow R_{i+1}, f_0 \otimes \cdots \otimes f_i \mapsto 1 \otimes \psi_0(f_0) \otimes \cdots \otimes \psi_i(f_i).$$

The σ_i 's are not morphisms of k -algebras but make the diagram

$$\begin{array}{ccc} R_i & \xrightarrow{\sigma_i} & R_{i+1} \\ \uparrow & & \uparrow \\ k & \xrightarrow{\sigma} & k \end{array}$$

commutative. Now we can define $[\sigma]_k R$ as the limit (i.e., the union) of the R_i 's ($i \geq 0$). Taking the limit of the σ_i 's yields a morphism $\sigma: [\sigma]_k R \rightarrow [\sigma]_k R$, turning $[\sigma]_k R$ into a k - σ -algebra. Let ψ denote the natural map $R = R_0 \hookrightarrow [\sigma]_k R$ and let $\psi': R \rightarrow S$ be a morphism of k -algebras from R into a k - σ -algebra S . For every $i \in \mathbb{N}$ we have a morphism of k -algebras⁸

$$\phi_i: \sigma^i R \rightarrow S, f \otimes \lambda \mapsto \sigma^i(\psi'(f))\lambda$$

and we obtain a morphism of k -algebras

$$\phi: [\sigma]_k R \rightarrow S, f_0 \otimes \cdots \otimes f_i \mapsto \phi_0(f_0) \cdots \phi_i(f_i).$$

For $f \otimes \lambda \in \sigma^i R$ we have

$$\phi_{i+1}(\psi_i(f \otimes \lambda)) = \phi_{i+1}(f \otimes \sigma(\lambda)) = \sigma^{i+1}(\psi'(f))\sigma(\lambda) = \sigma(\sigma^i(\psi'(f))\lambda) = \sigma(\phi_i(f \otimes \lambda)).$$

⁶To see that ψ_i is well-defined consider $\sigma^{i+1} R$ as k -algebra via the map $k \rightarrow \sigma^{i+1} R, \lambda \mapsto \lambda \otimes 1 = 1 \otimes \sigma^{i+1}(\lambda)$.

⁷To see that σ_i is well-defined consider R_{i+1} as k -algebra via the map $k \xrightarrow{\sigma} k \rightarrow R_{i+1}$.

⁸To see that ϕ_i is well-defined consider S as k -algebra via the map $k \xrightarrow{\sigma^i} k \rightarrow S$.

This implies

$$\begin{aligned}\phi(\sigma(f_0 \otimes \cdots \otimes f_i)) &= \phi(1 \otimes \psi_0(f_0) \otimes \cdots \otimes \psi_i(f_i)) = \phi_1(\psi_0(f_0)) \cdots \phi_{i+1}(\psi_i(f_i)) = \\ &= \sigma(\phi_0(f_0)) \cdots \sigma(\phi_i(f_i)) = \sigma(\phi_0(f_0) \cdots \phi_i(f_i)) = \sigma(\phi(f_0 \otimes \cdots \otimes f_i)).\end{aligned}$$

So ϕ is a morphism of k - σ -algebras. It is clear that ϕ makes diagram (1.3) commutative. To see that ϕ is unique we only have to observe that any other k - σ -morphism $\phi': [\sigma]_k R \rightarrow S$ with $\phi'\psi = \psi'$ has to agree with ϕ_i on $\sigma^i R$, since $\sigma^i(\psi(f)) = f \otimes 1 \in \sigma^i R \subset [\sigma]_k R$ and $\phi'(\sigma^i(\psi(f))) = \sigma^i(\phi'(\psi(f))) = \sigma^i(\psi'(f))$ for $f \in R$.

The uniqueness of $[\sigma]_k R$ can be seen as usual: Assume that $\psi_1: R \rightarrow S_1$ also satisfies the universal property. From the universal property of $[\sigma]_k R$ we obtain a morphism of k - σ -algebras $\phi: [\sigma]_k R \rightarrow S_1$ with $\phi\psi = \psi_1$. The universal property of S_1 yields a morphism of k - σ -algebras $\phi_1: S_1 \rightarrow [\sigma]_k R$ such that $\phi_1\psi_1 = \psi$. Consequently $\phi_1\phi\psi = \psi$. But also $\text{id}\psi = \psi$, therefore $\phi_1\phi = \text{id}$. Similarly $\phi\phi_1 = \text{id}$. \square

It follows from the above construction of $[\sigma]_k R$ that $\psi: R \rightarrow [\sigma]_k R$ is injective. In essence Proposition 1.1.26 gives a coordinate free way of thinking of algebraic equations as difference equations:

Example 1.1.27. Let k be a σ -field and $R = k[y_1, \dots, y_n]$ the polynomial ring over k in the variables y_1, \dots, y_n . Then $[\sigma]_k R = k\{y_1, \dots, y_n\}$ is the σ -polynomial ring over k in the σ -variables y_1, \dots, y_n .

Proof. You can either see this by looking at the construction of $[\sigma]_k R$ in the proof of Proposition 1.1.26, or you can use the universal property of the σ -polynomial ring. \square

Example 1.1.28. Let k be a σ -field and R a finitely generated k -algebra. We can write R as a quotient of a polynomial ring, say $R = k[y_1, \dots, y_n]/\mathfrak{a}$. Then $[\sigma]_k R = k\{y_1, \dots, y_n\}/[\mathfrak{a}]$. In particular, $[\sigma]_k R$ is finitely σ -generated over k .

Proof. Again you can see this by looking at the construction of $[\sigma]_k R$ in the proof of Proposition 1.1.26. (Note that $\sigma^i R \simeq k[y_1, \dots, y_n]/\mathfrak{a}_i$ where \mathfrak{a}_i denotes the ideal of $k[y_1, \dots, y_n]$ generated by all polynomials which are obtained from polynomials in \mathfrak{a} by applying σ^i to the coefficients.) Alternatively, you can verify the universal property directly. \square

Exercise 1.1.29. Make sure you understand all details of examples 1.1.27 and 1.1.28.

If $\varphi: R \rightarrow R'$ is a morphism of k -algebras, then so is $R \rightarrow R' \rightarrow [\sigma]_k R'$, and from the universal property we obtain a morphism $[\sigma]_k(\varphi): [\sigma]_k R \rightarrow [\sigma]_k R'$. We thus see that $[\sigma]_k$ is a functor from the category of k -algebras to the category of k - σ -algebras.

If S is a k - σ -algebra we denote with S^\sharp the underlying k -algebra of S . I.e., $(-)^{\sharp}$ is the forgetful functor from k - σ -algebras to k -algebras that forgets σ . For every k -algebra R and every k - σ -algebra S we have a natural bijection between the k - σ -algebra morphisms from $[\sigma]_k R$ to S and the k -algebra morphisms from R to S^\sharp . In other words, $[\sigma]_k$ is left adjoint to $(-)^{\sharp}$.

Exercise 1.1.30. Let k be a σ -field and R a k -algebra which is an integral domain. When is $[\sigma]_k R$ an integral domain?

Exercise 1.1.31. Recall the very basic notions of category theory in class. This should include the notion of equivalence of categories and adjoint functors.

1.2 Difference ideals

In this section we introduce and study basic properties of difference ideals. We also introduce a natural topology which is analogous to the Zariski topology on the set of prime ideals of a ring.

1.2.1 Properties of difference ideals

Definition 1.2.1. Let R be a σ -ring. A σ -ideal \mathfrak{a} of R is called

- reflexive if $\sigma^{-1}(\mathfrak{a}) = \mathfrak{a}$, i.e., $\sigma(f) \in \mathfrak{a}$ implies $f \in \mathfrak{a}$ for $f \in R$,
- well-mixed if $fg \in \mathfrak{a}$ implies $f\sigma(g) \in \mathfrak{a}$ for $f, g \in R$,
- perfect if $\sigma^{i_1}(f) \cdots \sigma^{i_n}(f) \in \mathfrak{a}$ implies $f \in \mathfrak{a}$ for $f \in R$ and $i_1, \dots, i_n \in \mathbb{N}$,
- σ -prime if \mathfrak{a} is reflexive and a prime ideal,
- σ -maximal if \mathfrak{a} is a maximal element in the set of all σ -ideals of R , which are not equal to R . (Here “maximal” refers to the partial order given by inclusion of σ -ideals.)

Exercise 1.2.2. Show that a σ -maximal σ -ideal is radical and reflexive. Give an example of a σ -maximal σ -ideal which is not σ -prime.

Definition 1.2.3. A σ -ring R is called a σ -domain if R is an integral domain and $\sigma: R \rightarrow R$ is injective.

So a σ -ring R is a σ -domain if and only if the zero ideal of R is σ -prime.

Lemma 1.2.4.

- (i) A perfect σ -ideal is reflexive, well-mixed and radical.
- (ii) A σ -prime ideal is perfect.

Proof. (i): Assume that \mathfrak{a} is a perfect σ -ideal in a σ -ring R . If $f \in R$ with $\sigma(f) \in \mathfrak{a}$, then $f \in \mathfrak{a}$. (This is the case $n = 1$ and $i_1 = 1$ in the definition of perfect.) Therefore \mathfrak{a} is reflexive. If $fg \in \mathfrak{a}$, then $\sigma(f)\sigma(g) \in \mathfrak{a}$. Thus, $f\sigma(g)\sigma(f\sigma(g)) \in \mathfrak{a}$. Because \mathfrak{a} is perfect this implies $f\sigma(g) \in \mathfrak{a}$. If $f^n \in \mathfrak{a}$ then $f \in \mathfrak{a}$. (This is the case $i_1 = \dots = i_n = 0$ in the definition of perfect.)

(ii): Let \mathfrak{p} be a σ -prime ideal of a σ -ring R . If $f \in R$ with $\sigma^{i_1}(f) \cdots \sigma^{i_n}(f) \in \mathfrak{p}$, then $\sigma^{i_j}(f) \in \mathfrak{p}$ for some $j \in \{1, \dots, n\}$ because \mathfrak{p} is prime. Since \mathfrak{p} is reflexive this implies $f \in \mathfrak{p}$. \square

Lemma 1.2.5. A σ -ideal \mathfrak{a} of a σ -ring R is perfect if and only if $f\sigma(f) \in \mathfrak{a}$ implies $f \in \mathfrak{a}$ for $f \in R$.

Proof. Let \mathfrak{a} be a σ -ideal such that $f\sigma(f) \in \mathfrak{a}$ implies $f \in \mathfrak{a}$ for $f \in R$.

We first show that \mathfrak{a} is radical: Let $f \in R$ with $f \notin \mathfrak{a}$. We have to show that $f^n \notin \mathfrak{a}$ for $n \geq 1$. It suffices to show that $f^{2^n} \notin \mathfrak{a}$ for $n \geq 0$. We will prove this by induction on n : So we assume that $n \geq 1$ and $f^{2^{n-1}} \notin \mathfrak{a}$. By assumption, also $f^{2^{n-1}}\sigma(f^{2^{n-1}}) \notin \mathfrak{a}$. Applying the same argument again gives

$$f^{2^{n-1}}\sigma(f^{2^{n-1}})\sigma(f^{2^{n-1}}\sigma(f^{2^{n-1}})) \notin \mathfrak{a}.$$

But then also $\sigma(f^{2^n}) = \sigma(f^{2^{n-1}})\sigma(f^{2^{n-1}}) \notin \mathfrak{a}$. Because \mathfrak{a} is a σ -ideal this yields $f^{2^n} \notin \mathfrak{a}$.

Now assume that $\sigma^{i_1}(f) \cdots \sigma^{i_n}(f) \in \mathfrak{a}$ for some $f \in R$. We have to show that $f \in \mathfrak{a}$. Let m be the maximum of $\{i_1, \dots, i_n\}$. After multiplying $\sigma^{i_1}(f) \cdots \sigma^{i_n}(f) \in \mathfrak{a}$ with appropriate factors we see that $(f\sigma(f) \cdots \sigma^m(f))^l \in \mathfrak{a}$ for some $l \geq 1$. Because \mathfrak{a} is radical, $f\sigma(f) \cdots \sigma^m(f) \in \mathfrak{a}$. If m equals 0 or 1 we are done. Otherwise we set $g := f\sigma(f) \cdots \sigma^{m-1}(f)$. Then $g\sigma(g) \in \mathfrak{a}$ because it is a multiple of $f\sigma(f) \cdots \sigma^m(f)$. By assumption, $g = f\sigma(f) \cdots \sigma^{m-1}(f) \in \mathfrak{a}$. Continuing in this way we arrive at $f\sigma(f) \in \mathfrak{a}$. So $f \in \mathfrak{a}$, as desired. \square

Lemma 1.2.6. *Let $\phi: R \rightarrow S$ be a morphism of σ -rings. If $\mathfrak{a} \subset S$ is a σ -ideal then also $\phi^{-1}(\mathfrak{a})$ is a σ -ideal. Moreover, if \mathfrak{a} is reflexive/well-mixed/perfect/ σ -prime, then so is $\phi^{-1}(\mathfrak{a})$.*

Exercise 1.2.7. *Proof Lemma 1.2.6 above.*

Proposition 1.2.8. *Let R be a σ -ring, \mathfrak{a} a σ -ideal of R and $\phi: R \rightarrow R/\mathfrak{a}$ the canonical map. The map $\mathfrak{b} \mapsto \phi^{-1}(\mathfrak{b})$ is a bijection between the set of σ -ideals of R/\mathfrak{a} and the set of σ -ideals of R which contain \mathfrak{a} . This map stays bijective if we restrict to reflexive/well-mixed/perfect/ σ -prime ideals on both sides.*

Proof. If \mathfrak{c} is a σ -ideal of R then $\phi(\mathfrak{c})$ is a σ -ideal of R/\mathfrak{a} . If $\mathfrak{a} \subset \mathfrak{c}$ then $\phi^{-1}(\phi(\mathfrak{c})) = \mathfrak{c} + \mathfrak{a} = \mathfrak{c}$. Clearly $\phi(\phi^{-1}(\mathfrak{b})) = \mathfrak{b}$ for every σ -ideal \mathfrak{b} of R/\mathfrak{a} . This shows that $\mathfrak{b} \mapsto \phi^{-1}(\mathfrak{b})$ is bijective with inverse $\mathfrak{c} \mapsto \phi(\mathfrak{c})$.

Because of Lemma 1.2.6 it now suffices to show that $\phi(\mathfrak{c})$ is reflexive/well-mixed/perfect/ σ -prime if $\mathfrak{c} \subset R$, $\mathfrak{a} \subset \mathfrak{c}$ has the corresponding property. But this is easy. For example, assume that \mathfrak{c} is perfect. Let $f \in R$ such that $\bar{f}\sigma(\bar{f}) \in \phi(\mathfrak{c})$. Then $f\sigma(f) \in \mathfrak{c}$ and so $f \in \mathfrak{c}$ and $\bar{f} \in \phi(\mathfrak{c})$. \square

Let R be ring, $S \subset R$ a multiplicatively closed subset and $\mathfrak{a} \subset R$ an ideal. Recall that

$$\mathfrak{a} : S = \{f \in R \mid \exists s \in S \text{ with } sf \in \mathfrak{a}\}$$

is called the saturation of \mathfrak{a} with respect to S . If $\mathfrak{a} = \mathfrak{a} : S$, we say that \mathfrak{a} is saturated with respect to S . Note that if R is a σ -ring, \mathfrak{a} a σ -ideal and S σ -stable, then $\mathfrak{a} : S$ is a σ -ideal. Let

$$\phi: R \rightarrow S^{-1}R, f \mapsto f/1$$

denote the canonical map. Recall that $\phi^{-1}(S^{-1}\mathfrak{a}) = \mathfrak{a} : S$ and that $\mathfrak{b} = S^{-1}\phi^{-1}(\mathfrak{b})$ for every ideal \mathfrak{b} of $S^{-1}R$. In particular, $\mathfrak{b} \mapsto \phi^{-1}(\mathfrak{b})$ is a bijection between the ideals of $S^{-1}R$ and the ideals of R which are saturated with respect to S . The inverse is $\mathfrak{a} \mapsto S^{-1}\mathfrak{a}$.

Proposition 1.2.9. *Let R be a σ -ring, $S \subset R$ a multiplicatively closed, σ -stable subset and $\phi: R \rightarrow S^{-1}R$ the canonical map. The map $\mathfrak{b} \mapsto \phi^{-1}(\mathfrak{b})$ is a bijection between the σ -ideals of $S^{-1}R$ and the σ -ideals of R which are saturated with respect to S . This map stays bijective if we restrict to reflexive/well-mixed/perfect/ σ -prime ideals on both sides.*

Proof. We only have to verify that $S^{-1}\mathfrak{a}$ is reflexive/well-mixed/perfect/ σ -prime if \mathfrak{a} has the corresponding property. We do this exemplarily for the case that \mathfrak{a} is perfect. Let $f \in R$ and $s \in S$ such that $f/s \cdot \sigma(f/s) \in S^{-1}\mathfrak{a}$. But then $f\sigma(f)/1 \in S^{-1}\mathfrak{a}$, i.e.,

$$f\sigma(f) \in \phi^{-1}(S^{-1}\mathfrak{a}) = \mathfrak{a} : S = \mathfrak{a}.$$

So $f \in \mathfrak{a}$ and $f/s \in S^{-1}\mathfrak{a}$ as desired. \square

Exercise 1.2.10. *Do all the cases in the proof of Proposition 1.2.8 and 1.2.9*

Lemma 1.2.11. *Let R be a σ -ring, $S \subset R$ a multiplicatively closed, σ -stable subset and \mathfrak{a} a perfect σ -ideal of R . Then $\mathfrak{a} : S$ is a perfect σ -ideal of R .*

Proof. Let $f \in R$ such that $f\sigma(f) \in \mathfrak{a} : S$. Then $sf\sigma(f) \in \mathfrak{a}$ for some $s \in S$. It follows that $sf\sigma(sf) \in \mathfrak{a}$. Because \mathfrak{a} is perfect this yields $sf \in \mathfrak{a}$, and so $f \in \mathfrak{a} : S$. \square

If R is a σ -ring and \mathfrak{p} a σ -prime ideal of R , then $S := R \setminus \mathfrak{p}$ is multiplicatively closed and σ -stable. Therefore

$$R_{\mathfrak{p}} := S^{-1}R$$

carries naturally the structure of a σ -ring. This ring has a unique maximal ideal, namely the ideal $S^{-1}\mathfrak{p}$. Note that this ideal is a σ -ideal. The quotient of $R_{\mathfrak{p}}$ by $S^{-1}\mathfrak{p}$ is denoted by

$$k(\mathfrak{p}).$$

This is naturally a σ -field and called the residue σ -field at \mathfrak{p} . Another way to construct $k(\mathfrak{p})$ is to take the quotient field of R/\mathfrak{p} .

Exercise 1.2.12. *Let R be a σ -ring, $S \subset R$ a multiplicatively closed, σ -stable subset and \mathfrak{a} a σ -ideal of R . If \mathfrak{a} is reflexive/well-mixed/ σ -prime, does $\mathfrak{a} : S$ have the same property?*

Definition 1.2.13. *Let R be a σ -ring and \mathfrak{a} a σ -ideal of R . The intersection of reflexive σ -ideals is a reflexive σ -ideal. Therefore there exists a unique smallest reflexive σ -ideal of R containing \mathfrak{a} . It is called the reflexive closure of \mathfrak{a} and denoted by \mathfrak{a}^* .*

Example 1.2.14. Let k be a σ -field and $k\{y_1\}$ the univariate σ -polynomial ring over k . The reflexive closure of $[\sigma(y_1)]$ is $[y_1]$.

Lemma 1.2.15. *Let R be a σ -ring and \mathfrak{a} a σ -ideal of R . Then*

$$\mathfrak{a}^* = \{f \in R \mid \exists n \in \mathbb{N} \text{ with } \sigma^n(f) \in \mathfrak{a}\}.$$

Proof. Let $\mathfrak{b} = \{f \in R \mid \exists n \in \mathbb{N} \text{ with } \sigma^n(f) \in \mathfrak{a}\}$. Then \mathfrak{b} is a reflexive σ -ideal of R . So $\mathfrak{a}^* \subset \mathfrak{b}$. On the other hand, if $\sigma^n(f) \in \mathfrak{a} \subset \mathfrak{a}^*$, then $f \in \mathfrak{a}^*$. So, $\mathfrak{b} \subset \mathfrak{a}^*$. \square

Lemma 1.2.16. *Let R be a σ -ring and \mathfrak{p} a prime σ -ideal of R , then \mathfrak{p}^* is σ -prime.*

Proof. First of all, note that $1 \notin \mathfrak{p}^*$. If $fg \in \mathfrak{p}^*$ then $\sigma^n(fg) = \sigma^n(f)\sigma^n(g) \in \mathfrak{p}$ for some $n \in \mathbb{N}$. Therefore $f \in \mathfrak{p}^*$ or $g \in \mathfrak{p}^*$. \square

1.2.2 Perfect difference ideals

Definition 1.2.17. *Let R be a σ -ring. The intersection of perfect difference ideals is a perfect difference ideal. Therefore, every subset F of R is contained in a smallest perfect difference ideal, which we call the perfect closure of F or the perfect difference ideal generated by F . It is denoted by $\{F\}$.*

🐞 1.2.18. *The notation $\{f_1, \dots, f_n\}$ can mean either “the perfect σ -ideal generated by f_1, \dots, f_n ” or “the set with elements f_1, \dots, f_n ”. Which one it is, should be clear from the context. I agree that this is not perfect, but this is standard notation.*

As it was already indicated in the “Motivation” the perfect difference ideals play a distinguished role in difference algebra. (See also Lemma 2.1.7 below.) One of the very basic challenges in difference algebra is that the perfect closure of a σ -ideal has no simple description. For comparison, in commutative algebra, the radical $\sqrt{\mathfrak{a}}$ of an ideal \mathfrak{a} of a ring R can be defined as the smallest radical ideal of R containing \mathfrak{a} . The radical has the

simple description $\sqrt{\mathfrak{a}} = \{f \in R \mid \exists n \in \mathbb{N}_{\geq 1} \text{ such that } f^n \in \mathfrak{a}\}$. The perfect closure of a σ -ideal can only be described by a recursive procedure which we will now explain. This procedure is sometimes called *shuffling*. For a σ -ideal \mathfrak{a} of a σ -ring R we set

$$\mathfrak{a}' := \{f \in R \mid \exists i_1, \dots, i_n \in \mathbb{N} \text{ such that } \sigma^{i_1}(f) \cdots \sigma^{i_n}(f) \in \mathfrak{a}\}.$$

Note that \mathfrak{a}' need not be an ideal. Let F be a subset of R . We define⁹

$$F^{\{1\}} := [F]'$$

and recursively

$$F^{\{i\}} := (F^{\{i-1\}})^{\{1\}} = [F^{\{i-1\}}]'$$
 for $i = 2, 3, \dots$

We claim that $\{F\} = \bigcup_{i \geq 1} F^{\{i\}}$. Because $[F^{\{i-1\}}] \subset F^{\{i\}}$ it is clear that the union is a σ -ideal. It is also clear from the construction that the union $\bigcup_{i \geq 1} F^{\{i\}}$ is a perfect σ -ideal. Therefore $\{F\} \subset \bigcup_{i \geq 1} F^{\{i\}}$. The reverse inclusion follows by induction.

Lemma 1.2.19. *Let F and G be subsets of a σ -ring R .*

- (i) *If F and G are σ -stable then $F^{\{1\}}G^{\{1\}} \subset (FG)^{\{1\}}$.*
- (ii) *$F^{\{i\}}G^{\{i\}} \subset (FG)^{\{i+1\}}$ for $i = 1, 2, \dots$*
- (iii) *$(FG)^{\{i\}} \subset F^{\{i\}} \cap G^{\{i\}}$ for $i = 1, 2, \dots$*
- (iv) *$F^{\{i\}} \cap G^{\{i\}} \subset (FG)^{\{i+1\}}$ for $i = 1, 2, \dots$*

Proof. (i): Let $f \in F^{\{1\}}$ and $g \in G^{\{1\}}$. Then there exist $i_1, \dots, i_n \in \mathbb{N}$ such that $f' = \sigma^{i_1}(f) \cdots \sigma^{i_n}(f) \in [F]$. Similarly, there exist $j_1, \dots, j_m \in \mathbb{N}$ such that $g' = \sigma^{j_1}(g) \cdots \sigma^{j_m}(g) \in [G]$. Because F is σ -stable the elements of $[F]$ are R -linear combinations of elements of F ; similarly for G . This implies that $f'g' \in [FG]$. Since $\sigma^{i_1}(fg) \cdots \sigma^{i_n}(fg)\sigma^{j_1}(fg) \cdots \sigma^{j_m}(fg)$ is a multiple of $f'g'$, it belongs to $[FG]$. Hence, $fg \in (FG)^{\{1\}}$ and $F^{\{1\}}G^{\{1\}} \subset (FG)^{\{1\}}$.

(ii): Set $\tilde{F} = \{\sigma^i(f) \mid f \in F, i \in \mathbb{N}\}$ and $\tilde{G} = \{\sigma^i(g) \mid g \in G, i \in \mathbb{N}\}$. We claim that $\tilde{F}\tilde{G} \subset (FG)^{\{1\}}$: Let $f \in F, g \in G$ and $i, j \in \mathbb{N}$. Since $\sigma^i(\sigma^i(f)\sigma^j(g))\sigma^j(\sigma^i(f)\sigma^j(g))$ is a multiple of $\sigma^{i+j}(fg) \in [FG]$ it follows that $\sigma^i(f)\sigma^j(g) \in (FG)^{\{1\}}$. So $\tilde{F}\tilde{G} \subset (FG)^{\{1\}}$ as claimed.

Because \tilde{F}, \tilde{G} are σ -stable, we can use (i) to obtain

$$F^{\{1\}}G^{\{1\}} \subset \tilde{F}^{\{1\}}\tilde{G}^{\{1\}} \subset (\tilde{F}\tilde{G})^{\{1\}} \subset ((FG)^{\{1\}})^{\{1\}} = (FG)^{\{2\}}.$$

This proves the case $i = 1$. For $i \geq 2$ we use induction: Assume $F^{\{i-1\}}G^{\{i-1\}} \subset (FG)^{\{i\}}$. Because $F^{\{i-1\}}$ and $G^{\{i-1\}}$ are σ -stable

$$F^{\{i\}}G^{\{i\}} = (F^{\{i-1\}})^{\{1\}}(G^{\{i-1\}})^{\{1\}} \subset (F^{\{i-1\}}G^{\{i-1\}})^{\{1\}} \subset ((FG)^{\{i\}})^{\{1\}} = (FG)^{\{i+1\}}.$$

(iii): We have $[FG] \subset [F]$ and therefore $(FG)^{\{i\}} \subset F^{\{i\}}$. Likewise, $(FG)^{\{i\}} \subset G^{\{i\}}$.

(iv): Let $f \in F^{\{i\}} \cap G^{\{i\}}$. Then $f^2 \in F^{\{i\}}G^{\{i\}} \subset (FG)^{\{i+1\}}$ by (ii). Hence, $f \in (FG)^{\{i+1\}}$. \square

Proposition 1.2.20. *Let F and G be subsets of a σ -ring R . Then*

$$\{F\} \cap \{G\} = \{FG\}.$$

⁹In [Cohn] and [Levin] the notation F_i is used for $F^{\{i\}}$.

Proof. This is clear from (iii) and (iv) of Lemma 1.2.19. \square

Lemma 1.2.21. *Let R be a σ -ring and \mathfrak{a} a maximal element in the set of all perfect σ -ideals of R which are not equal to R . Then \mathfrak{a} is a σ -prime ideal.*

Proof. We only have to show that \mathfrak{a} is a prime ideal. Let $f, g \in R$ with $fg \in \mathfrak{a}$. Assume that $f \notin \mathfrak{a}$. Then $\{\mathfrak{a}, f\} = R$ by the maximality of \mathfrak{a} . We have to show that $g \in \mathfrak{a}$. By Proposition 1.2.20

$$\{\mathfrak{a}, f\} \cap \{\mathfrak{a}, g\} = \{\mathfrak{a}\mathfrak{a}, f\mathfrak{a}, g\mathfrak{a}, fg\} \subset \mathfrak{a}.$$

Therefore

$$\mathfrak{a} = \{\mathfrak{a}, f\} \cap \{\mathfrak{a}, g\} = R \cap \{\mathfrak{a}, g\} = \{\mathfrak{a}, g\}.$$

So $g \in \mathfrak{a}$. \square

Proposition 1.2.22. *Let R be a σ -ring and $F \subset R$. Then $\{F\}$ is the intersection of all σ -prime ideals of R which contain F . In particular, every perfect σ -ideal is the intersection of σ -prime ideals¹⁰.*

Proof. Because σ -prime ideals are perfect it is clear that $\{F\}$ is contained in every σ -prime ideal containing F . So it suffices to show that every perfect σ -ideal \mathfrak{a} of R is the intersection of σ -prime ideals. Replacing R by R/\mathfrak{a} we may assume that \mathfrak{a} is the zero ideal (Proposition 1.2.8). So we assume that R is a σ -ring with perfect zero ideal.

Let $f \in R$ such that f is contained in every σ -prime ideal of R . We have to show that $f = 0$. Suppose the contrary. An expression of the form $\sigma^{i_1}(f) \cdots \sigma^{i_n}(f)$ cannot be zero, since otherwise $f = 0$ as the zero ideal is perfect. Therefore $R\{1/f\}$ is not the zero ring. (Unless R is the zero ring, in which case the statement is trivial.) Let \mathfrak{b} denote the saturation of the zero ideal with respect to the multiplicatively closed, σ -stable subset of R generated by f . We know from Lemma 1.2.11 that \mathfrak{b} is a perfect σ -ideal, moreover $1 \notin \mathfrak{b}$. It is clear from Proposition 1.2.9 that the ideal generated by \mathfrak{b} in $R\{1/f\}$ is perfect and does not contain 1.

Clearly the union of a chain of perfect σ -ideals not containing 1 is a perfect σ -ideal not containing 1. It therefore follows from Zorn's lemma that there exists a maximal element \mathfrak{p} in the set of perfect σ -ideals of $R\{1/f\}$ not containing 1. We know from Lemma 1.2.21 that \mathfrak{p} is σ -prime and by Proposition 1.2.9 the inverse image \mathfrak{q} of \mathfrak{p} under $R \rightarrow R\{1/f\}$ is a σ -prime ideal of R with $f \notin \mathfrak{q}$; a contradiction. \square

Exercise 1.2.23. *Let R be a σ -ring. An element $f \in R$ is called a σ -unit if $1 \in \{f\}$. Show that the set of σ -units is multiplicatively closed and σ -stable. Find an example of a σ -unit which is not a unit.*

Exercise 1.2.24. ** Let k be a σ -field. Is every σ -unit in the σ -polynomial ring $k\{y_1, \dots, y_n\}$ a unit, i.e. an element of k ?*

1.2.3 Well-mixed difference ideals

A nice feature of well-mixed difference ideals is that their perfect closure can easily be described:

Lemma 1.2.25. *Let R be a σ -ring and $\mathfrak{a} \subset R$ a well-mixed σ -ideal. Then*

$$\{\mathfrak{a}\} = \{f \in R \mid \exists n \in \mathbb{N}, m \in \mathbb{N}_{\geq 1} : \sigma^n(f)^m \in \mathfrak{a}\} (= \sqrt{\mathfrak{a}^*} = \sqrt{\mathfrak{a}}^*).$$

¹⁰If $\{F\} = R$ the intersection is understood to be empty.

Proof. Set $\mathfrak{b} = \{f \in R \mid \exists n \in \mathbb{N}, m \in \mathbb{N}_{\geq 1} : \sigma^n(f)^m \in \mathfrak{a}\}$. Note that \mathfrak{b} is a σ -ideal and $\mathfrak{a} \subset \mathfrak{b}$. Clearly $\mathfrak{b} \subset \{\mathfrak{a}\}$. So it suffices to show that \mathfrak{b} is perfect. Let $f \in R$ with $f\sigma(f) \in \mathfrak{b}$. Then $\sigma^n(f\sigma(f))^m \in \mathfrak{a}$ for some $n \in \mathbb{N}, m \in \mathbb{N}_{\geq 1}$. Because \mathfrak{a} is well-mixed $\sigma^n(f)^m \cdot \sigma^{n+1}(f)^m = \sigma^n(f\sigma(f))^m \in \mathfrak{a}$ implies $\sigma(\sigma^n(f)^m) \cdot \sigma^{n+1}(f)^m = \sigma^{n+1}(f)^{2m} \in \mathfrak{a}$. So $f \in \mathfrak{b}$. \square

The above lemma shows that well-mixed σ -ideals are rather close to being perfect.

Let R be a σ -ring. Since the intersection of radical, well-mixed σ -ideals is a radical, well-mixed σ -ideal, every subset F of R is contained in a unique smallest radical, well-mixed σ -ideal. We denote it by

$$\{F\}_{wm}.$$

Exercise 1.2.26. Let R be a σ -ring and $F \subset R$. Show that $(\{F\}_{wm})^* = \{F\}$.

In the next proof will use the following three basic facts from commutative algebra:

- In a Noetherian ring every radical ideal is the finite intersection of prime ideals.
- In a Noetherian ring R every minimal prime ideal of R is the annihilator of an element of R .
- If $R \subset S$ is an inclusion of rings and \mathfrak{q} a minimal prime ideal of R , then there exists a minimal prime ideal \mathfrak{q}' of S with $\mathfrak{q}' \cap R = \mathfrak{q}$.

Exercise 1.2.27. If you are not familiar with the above three facts we can discuss this during the exercise session.

Proposition 1.2.28. Let R be a σ -ring and $F \subset R$. Then $\{F\}_{wm}$ is the intersection of all prime σ -ideals which contain F . In particular, every radical, well-mixed σ -ideal is the intersection of prime σ -ideals.

Proof. Clearly a prime σ -ideal is well-mixed. Therefore $\{F\}_{wm}$ is contained in the intersection of all prime σ -ideals which contain F . It remains to show that $\{F\}_{wm}$ is the intersection of prime σ -ideals. Using Proposition 1.2.8 we can replace R with $R/\{F\}_{wm}$. We then have to prove the following

Claim: If the zero ideal of R is radical and well-mixed, then the zero ideal is the intersection of prime σ -ideals.

We first assume that R is finitely σ -generated over \mathbb{Z} . Let $f \in R$ be contained in every prime σ -ideal of R . We have to show that $f = 0$. Suppose $f \neq 0$. Choose a finite subset A of R which σ -generates R over \mathbb{Z} . Then $f \in \mathbb{Z}[A, \sigma(A), \dots, \sigma^n(A)]$ for some $n \in \mathbb{N}$. Because $\mathbb{Z}[A, \dots, \sigma^n(A)]$ is Noetherian and reduced the zero ideal is the finite intersection of prime ideals. So there exists a minimal prime ideal \mathfrak{p}_0 of $\mathbb{Z}[A, \dots, \sigma^n(A)]$ with $f \notin \mathfrak{p}_0$. Considering the inclusion $\mathbb{Z}[A, \dots, \sigma^n(A)] \subset \mathbb{Z}[A, \dots, \sigma^{n+1}(A)]$ we find that there exists a minimal prime ideal \mathfrak{p}_1 of $\mathbb{Z}[A, \dots, \sigma^{n+1}(A)]$ with $\mathfrak{p}_1 \cap \mathbb{Z}[A, \dots, \sigma^n(A)] = \mathfrak{p}_0$. Going on like this we obtain for every $i \in \mathbb{N}$ a minimal prime ideal \mathfrak{p}_i of $\mathbb{Z}[A, \dots, \sigma^{n+i}(A)]$ such that $\mathfrak{p}_i \cap \mathbb{Z}[A, \dots, \sigma^{n+i-1}(A)] = \mathfrak{p}_{i-1}$. Set

$$\mathfrak{p} = \bigcup_{i \in \mathbb{N}} \mathfrak{p}_i.$$

Clearly \mathfrak{p} is a prime ideal of R with $f \notin \mathfrak{p}$. We claim that \mathfrak{p} is a σ -ideal: Let $g \in \mathfrak{p}$. Then $g \in \mathfrak{p}_i$ for some $i \in \mathbb{N}$. Increasing i if necessary we can assume that $\sigma(g) \in \mathbb{Z}[A, \dots, \sigma^{n+i}(A)]$. There exists an element $h \in \mathbb{Z}[A, \dots, \sigma^{n+i}(A)]$ such that

$$\mathfrak{p}_i = \text{Ann}(h) := \{h' \in R \mid h'h = 0\}.$$

Since $g \in \mathfrak{p}_i$ we have $gh = 0$. Because R is well-mixed this implies $\sigma(g)h = 0$, i.e., $\sigma(g) \in \mathfrak{p}_i$. This shows that \mathfrak{p} is a σ -ideal. But $f \notin \mathfrak{p}$ contradicts the assumption that f is contained in every σ -prime ideal of R . Thus the claim is proved in the case that R is finitely σ -generated over \mathbb{Z} .

To deduce the general case from the case just proved we will use the concept of filters. Recall that a filter \mathcal{F} on a set M is a set of subsets of M such that

- $\emptyset \notin \mathcal{F}$, $M \in \mathcal{F}$.
- If $V \subset M$, $U \in \mathcal{F}$ and $U \subset V$ then $V \in \mathcal{F}$.
- A finite intersection of elements in \mathcal{F} is in \mathcal{F} .

A filter is called an ultrafilter if for $V \subset M$ either $V \in \mathcal{F}$ or $M \setminus V \in \mathcal{F}$. An ultrafilter is the same thing as maximal filter and using Zorn's lemma one can show that every filter is contained in an ultrafilter.

Now let R be a σ -ring such that the zero ideal of R is reduced and well-mixed. Let $f \in R$, $f \neq 0$. We have to find a prime σ -ideal of R which does not contain f . Let M denote the set of all σ -subrings of R which are finitely σ -generated over \mathbb{Z} . Let \mathcal{F} denote the set of all subsets of M which contain a set of the form $\{S \in M \mid F \subset S\}$ for some finite set $F \subset R$. Then \mathcal{F} is a filter on M . Let \mathcal{G} denote an ultrafilter which contains \mathcal{F} .

The product

$$P := \prod_{S \in M} S$$

is naturally a difference ring with componentwise addition, multiplication and action of σ . We say that two elements $(g_S)_{S \in M}, (h_S)_{S \in M} \in P$ are equivalent (modulo \mathcal{G}) if $\{S \in M \mid g_S = h_S\} \in \mathcal{G}$. It is easy to see that this defines an equivalence relation and that the set of equivalence classes $T := P/\mathcal{G}$ is naturally a difference ring. Moreover, we have a map $\phi: R \rightarrow T$ which sends $g \in R$ to the equivalence class of $(g_S)_{S \in M}$ where $g_S = g$ if $g \in S$. (Since we are working with equivalence classes the value of g_S for $g \notin S$ is irrelevant.) Note that $\phi: R \rightarrow T$ is a morphism of difference rings.

For $S \in M$ let \mathfrak{p}_S be a prime σ -ideal of S . By the case of the claim already established we can choose \mathfrak{p}_S such that $f \notin \mathfrak{p}_S$ for every $S \in M$. The set \mathfrak{p} of all equivalence classes of elements $(g_S)_{S \in M}$ of P such that $\{S \in M \mid g_S \in \mathfrak{p}_S\} \in \mathcal{G}$ is called the ultraproduct of the \mathfrak{p}_S . Clearly, \mathfrak{p} is a σ -ideal of T . Because \mathcal{G} is an ultrafilter \mathfrak{p} is also prime: Let $(g_S)_{S \in M}, (h_S)_{S \in M} \in P$ such that $\{S \in M \mid g_S h_S \in \mathfrak{p}_S\} \in \mathcal{G}$. Assume that $V := \{S \in M \mid g_S \in \mathfrak{p}_S\} \notin \mathcal{G}$, i.e., the equivalence class of $(g_S)_{S \in M}$ does not lie in \mathfrak{p} . Then $M \setminus V \in \mathcal{G}$. But $(M \setminus V) \cap \{S \in M \mid g_S h_S \in \mathfrak{p}_S\} \subset \{S \in M \mid h_S \in \mathfrak{p}_S\}$. Therefore the equivalence class of $(h_S)_{S \in M}$ lies in \mathfrak{p} . Consequently \mathfrak{p} is prime. By construction $\phi(f) \notin \mathfrak{p}$. So $\phi^{-1}(\mathfrak{p})$ is a prime σ -ideal of R with $f \notin \mathfrak{p}$. \square

Exercise 1.2.29. *Is there an analog of the shuffling process for radical, well-mixed σ -ideals?*

1.2.4 The Cohn topology

In this subsection we encounter a first, somewhat geometric perspective in difference algebra. How this is related to the idea of difference varieties will be explained in the next chapter.

Definition 1.2.30. *Let R be a σ -ring. The set of all σ -prime ideals of R is denoted by*

$$\text{Spec}^\sigma(R),$$

and called the difference spectrum of R .

Note that $\text{Spec}^\sigma(R)$ is the set of fixed points of the map $\mathfrak{p} \mapsto \sigma^{-1}(\mathfrak{p})$ on $\text{Spec}(R)$.

1.2.31. $\text{Spec}^\sigma(R)$ can be the empty set, even if R is not the zero ring. For example, let S be a σ -ring and consider $R := S \oplus S$ as ring with componentwise addition and multiplication. Define a difference structure on R by $\sigma(s_1 \oplus s_2) = \sigma(s_2) \oplus \sigma(s_1)$. A prime ideal of R must either contain $S \oplus 0$ or $0 \oplus S$. Therefore, a σ -prime ideal of S would necessarily contain all of R . Consequently, $\text{Spec}^\sigma(R) = \emptyset$.

Exercise 1.2.32. Let $R \subset \text{Seq}_{\mathbb{C}}$ denote the σ -subring of all periodic sequences. Show that $\text{Spec}^\sigma(R)$ is empty.

Exercise 1.2.33. Show that for a Noetherian σ -ring R , there exists an integer $l \geq 1$ such that $\text{Spec}^{\sigma^l}(R)$ is non-empty.

Let R be a σ -ring and $F \subset R$. We set

$$\mathcal{V}(F) := \{\mathfrak{p} \in \text{Spec}^\sigma(R) \mid F \subset \mathfrak{p}\} \subset \text{Spec}^\sigma(R).$$

Obviously, $\mathcal{V}(F) = \mathcal{V}([F])$.

Lemma 1.2.34. Let R be a σ -ring.

- (i) $\mathcal{V}(0) = \text{Spec}^\sigma(R)$ and $\mathcal{V}(R) = \emptyset$.
- (ii) If $\mathfrak{a}, \mathfrak{b}$ are σ -ideals of R then $\mathcal{V}(\mathfrak{a}) \cup \mathcal{V}(\mathfrak{b}) = \mathcal{V}(\mathfrak{a} \cap \mathfrak{b})$.
- (iii) If $(\mathfrak{a}_i)_{i \in I}$ is a family of σ -ideals of R then $\bigcap_i \mathcal{V}(\mathfrak{a}_i) = \mathcal{V}(\sum_i \mathfrak{a}_i)$.

Proof. (i) and (iii) are obvious. Also the inclusion $\mathcal{V}(\mathfrak{a}) \cup \mathcal{V}(\mathfrak{b}) \subset \mathcal{V}(\mathfrak{a} \cap \mathfrak{b})$ is clear. Let $\mathfrak{p} \in \mathcal{V}(\mathfrak{a} \cap \mathfrak{b})$. If $\mathfrak{a} \subset \mathfrak{p}$ we are done. Thus, we may assume that there exists an $f \in \mathfrak{a}$ with $f \notin \mathfrak{p}$. Let $g \in \mathfrak{b}$. Then $fg \in \mathfrak{a} \cap \mathfrak{b}$ and therefore $fg \in \mathfrak{p}$. As $f \notin \mathfrak{p}$ this implies $g \in \mathfrak{p}$. Hence, $\mathfrak{b} \subset \mathfrak{p}$, i.e., $\mathfrak{p} \in \mathcal{V}(\mathfrak{b})$. \square

The above lemma shows that the sets of the form $\mathcal{V}(\mathfrak{a})$ are the closed sets of a topology on $\text{Spec}^\sigma(R)$. This topology is sometimes called the *Cohn topology*. In the sequel we will always consider $\text{Spec}^\sigma(R)$ as a topological space endowed with this topology. Note that the Cohn topology is the restriction of the usual (Zariski) topology on $\text{Spec}(R)$.

For $f \in R$ we set

$$D^\sigma(f) := \text{Spec}^\sigma(R) \setminus \mathcal{V}(f) = \{\mathfrak{p} \in \text{Spec}^\sigma(R) \mid f \notin \mathfrak{p}\}.$$

This is sometimes called a basic open subset of $\text{Spec}^\sigma(R)$.

Recall that a topological space X is called *irreducible* if $X = X_1 \cup X_2$ with $X_1, X_2 \subset X$ closed implies $X = X_1$ or $X = X_2$. A subset Y of X is called irreducible if it is irreducible with respect to the induced topology. This is equivalent to saying that $Y \subset X_1 \cup X_2$ with $X_1, X_2 \subset X$ closed implies $Y \subset X_1$ or $Y \subset X_2$.

A *generic point* of a closed subset Y of X is an element $y \in Y$ such that the closure of y equals Y .

Proposition 1.2.35. Let R be a σ -ring.

- (i) We have $\mathcal{V}(F) = \mathcal{V}(\{F\})$ for $F \subset R$.
- (ii) The map $\mathfrak{a} \mapsto \mathcal{V}(\mathfrak{a})$ defines an inclusion-reversing bijection between the set of all perfect σ -ideals of R and the set of all closed subsets of $\text{Spec}^\sigma(R)$.

- (iii) If \mathfrak{a} is a perfect σ -ideal of R , then $\mathcal{V}(\mathfrak{a})$ is irreducible if and only if \mathfrak{a} is a σ -prime ideal.
- (iv) $\text{Spec}^\sigma(R)$ is quasi-compact.
- (v) The open sets of the form $D^\sigma(f)$, $f \in R$ are a basis of the Cohn topology.
- (vi) Every irreducible closed subset of $\text{Spec}^\sigma(R)$ has a unique generic point.

Proof. (i): Clear because σ -prime ideals are perfect.

(ii): If $\mathfrak{a} \subset \mathfrak{b}$ then $\mathcal{V}(\mathfrak{b}) \subset \mathcal{V}(\mathfrak{a})$. If \mathfrak{a} is a perfect σ -ideal of R then $\mathfrak{a} = \bigcap_{\mathfrak{p} \in \mathcal{V}(\mathfrak{a})} \mathfrak{p}$ by Proposition 1.2.22. Thus the assignment $\mathfrak{a} \mapsto \mathcal{V}(\mathfrak{a})$ is injective. Because $\mathcal{V}(F) = \mathcal{V}(\{F\})$ for $F \subset R$ it is also surjective.

(iii): Let \mathfrak{a} be a σ -prime ideal. Assume $\mathcal{V}(\mathfrak{a}) \subset \mathcal{V}(F) \cup \mathcal{V}(G)$ for subsets F, G of R . Assume $\mathcal{V}(\mathfrak{a}) \not\subset \mathcal{V}(F)$. We have to show that $\mathcal{V}(\mathfrak{a}) \subset \mathcal{V}(G)$. As $\mathcal{V}(\mathfrak{a}) \not\subset \mathcal{V}(F)$ there must exist $f \in F$ with $f \notin \mathfrak{a}$. Let $g \in G$. Then every σ -prime ideal in $\mathcal{V}(\mathfrak{a})$ contains fg . Because $\mathfrak{a} = \bigcap_{\mathfrak{p} \in \mathcal{V}(\mathfrak{a})} \mathfrak{p}$ this shows that $fg \in \mathfrak{a}$. But as \mathfrak{a} is prime and $f \notin \mathfrak{a}$ we have $g \in \mathfrak{a}$. Therefore $G \subset \mathfrak{a}$, i.e., $\mathcal{V}(\mathfrak{a}) \subset \mathcal{V}(G)$ as desired.

Now assume that $\mathcal{V}(\mathfrak{a})$ is irreducible. Let $f, g \in R$ with $fg \in \mathfrak{a}$. Then $\mathcal{V}(\mathfrak{a}) \subset \mathcal{V}(f) \cup \mathcal{V}(g)$. So $\mathcal{V}(\mathfrak{a}) \subset \mathcal{V}(f)$ or $\mathcal{V}(\mathfrak{a}) \subset \mathcal{V}(g)$. But then $f \in \mathfrak{a}$ or $g \in \mathfrak{a}$. So \mathfrak{a} is prime.

(iv): Let $\mathcal{V}(\mathfrak{a}_i)_{i \in I}$ be a family of closed subsets of $\text{Spec}^\sigma(R)$ such that $\bigcap_{i \in J} \mathcal{V}(\mathfrak{a}_i) \neq \emptyset$ for every finite subset J of I . We have to show that $\bigcap_{i \in I} \mathcal{V}(\mathfrak{a}_i) \neq \emptyset$.

Note that $\bigcap_{i \in I} \mathcal{V}(\mathfrak{a}_i) = \mathcal{V}(\sum_{i \in I} \mathfrak{a}_i)$ is empty if and only if 1 lies in the perfect closure of $\sum_{i \in I} \mathfrak{a}_i$. The description of the perfect closure via the shuffling process shows that, in this case, 1 already lies in the perfect closure of a finite sum $\sum_{i \in J} \mathfrak{a}_i$.

(v): Let U be an open subset of $\text{Spec}^\sigma(R)$. Then U is of the form $U = \text{Spec}^\sigma(R) \setminus \mathcal{V}(F)$, $F \subset R$ and $U = \bigcup_{f \in F} D^\sigma(f)$.

(vi): Note that for an arbitrary subset Y of $\text{Spec}^\sigma(R)$ the closure \overline{Y} of Y is given by

$$\overline{Y} = \mathcal{V}\left(\bigcap_{\mathfrak{p} \in Y} \mathfrak{p}\right).$$

So if Y is an irreducible closed subset of $\text{Spec}^\sigma(R)$, then Y is of the form $Y = \mathcal{V}(\mathfrak{p})$ for some σ -prime ideal \mathfrak{p} of R . It follows that the closure of $y := \mathfrak{p}$ is equal to Y . The uniqueness is clear. \square

Exercise 1.2.36. Let R be a σ -ring. Define a topology on the set of all prime σ -ideals of R and work out an analog of Proposition 1.2.35.

Lemma 1.2.37. Let $\phi: R \rightarrow S$ be a morphism of σ -rings. Then

$$\phi^*: \text{Spec}^\sigma(S) \rightarrow \text{Spec}^\sigma(R), \mathfrak{p} \mapsto \phi^{-1}(\mathfrak{p})$$

is a continuous map. Indeed,

$$\phi^{*-1}(\mathcal{V}(F)) = \mathcal{V}(\phi(F))$$

for $F \subset R$.

Proof. We have

$$\begin{aligned} \phi^{*-1}(\mathcal{V}(F)) &= \{\mathfrak{p} \in \text{Spec}^\sigma(S) \mid F \subset \phi^{-1}(\mathfrak{p})\} = \\ &= \{\mathfrak{p} \in \text{Spec}^\sigma(S) \mid \phi(F) \subset \mathfrak{p}\} = \mathcal{V}(\phi(F)). \end{aligned}$$

\square

We conceive that Spec^σ is a contravariant functor from the category of difference rings to the category of topological spaces. Considering the analogy with algebraic geometry, one might feel tempted to add some extra structure, i.e., a structure sheaf, to the topological space, to make this functor fully faithful. Because of [§ 1.2.31](#) there is no hope for this to work. Nevertheless considering a certain naturally defined structure sheaf on $\text{Spec}^\sigma(R)$ is useful and natural. See [Hrushovski]. Since we are only interested in the “affine” situation, we are not inclined to consider sheaves.

Exercise 1.2.38. *Let R be a σ -ring. Show that the map $\text{Spec}^\sigma(R^*) \rightarrow \text{Spec}^\sigma(R)$ induced by $R \rightarrow R^*$ is a homeomorphism.*

Exercise 1.2.39. *Let R be a σ -ring and $S \subset R$ a multiplicatively closed σ -stable subset consisting of σ -units. (See [Exercise 1.2.23](#).) Show that $\text{Spec}^\sigma(R)$ and $\text{Spec}^\sigma(S^{-1}R)$ are homeomorphic.*

Lemma 1.2.40. *Let $\phi: R \rightarrow S$ be a morphism of σ -rings and \mathfrak{a} a perfect σ -ideal of S . Then the closure of $\phi^*(\mathcal{V}(\mathfrak{a}))$ in $\text{Spec}^\sigma(R)$ equals $\mathcal{V}(\phi^{-1}(\mathfrak{a}))$.*

Proof. We have

$$\overline{\phi^*(\mathcal{V}(\mathfrak{a}))} = \mathcal{V} \left(\bigcap_{\mathfrak{p} \in \phi^*(\mathcal{V}(\mathfrak{a}))} \mathfrak{p} \right) = \mathcal{V} \left(\bigcap_{\mathfrak{q} \in \mathcal{V}(\mathfrak{a})} \phi^{-1}(\mathfrak{q}) \right) = \mathcal{V} \left(\phi^{-1} \left(\bigcap_{\mathfrak{q} \in \mathcal{V}(\mathfrak{a})} \mathfrak{q} \right) \right) = \mathcal{V}(\phi^{-1}(\mathfrak{a})).$$

□

Exercise 1.2.41. *Let $\phi: R \rightarrow S$ be a surjective morphism of σ -rings. Show that the image of ϕ^* is closed. Which ideal defines the image?*

Chapter 2

Difference varieties

In this chapter we introduce the main protagonists of this course: Difference varieties. Roughly speaking, a difference variety is the set of solutions of a system of algebraic difference equations over some σ -field k . It is however quite crucial to specify where we are looking for the solutions. Here we will only be looking for solutions in σ -field extensions of k . If we allow ourselves to look for solutions in arbitrary k - σ -algebras this leads to the notion of a difference scheme.

2.1 Difference varieties and their coordinate rings

Let k be a σ -field and $k\{y\} = k\{y_1, \dots, y_n\}$ the σ -polynomial ring in the σ -variables y_1, \dots, y_n over k . By a system of algebraic difference equations over k we mean any subset F of $k\{y\}$. The main objective in difference algebra is to study the solutions of such systems. If K is a σ -field extension of k , by definition, the solutions of F in K are

$$\mathbb{V}_K(F) := \{a \in K^n \mid f(a) = 0 \text{ for all } f \in F\}.$$

If $\phi: K \rightarrow L$ is a morphism of σ -field extension of k , then we have an induced map $\mathbb{V}_K(F) \rightarrow \mathbb{V}_L(F)$, $a \mapsto \phi(a)$. This (and a couple of other obvious things) means that $\mathbb{V}(F)$ is a functor from the category of σ -field extension of k to the category of sets. The following definition defines the central objects of study of difference algebraic geometry.

Definition 2.1.1. *Let k be a σ -field. An (affine) σ -variety over k is a functor X from the category of σ -field extension of k to the category of sets which is of the form $X = \mathbb{V}(F)$ for some subset F of $k\{y_1, \dots, y_n\}$. In this situation we say that X is the σ -variety defined by F .*

Since, in this course, we will only consider affine σ -varieties we usually drop the adjective “affine”. We will also say that X is a k - σ -variety to express that X is a σ -variety over k . To facilitate easier notation we denote the category of σ -field extensions of k by

$$\sigma\text{-field}_k$$

and write $K \in \sigma\text{-field}_k$ to express that K is a σ -field extension of k .

Exercise 2.1.2. *Determine the σ -variety defined by $y_1 + \sigma(y_1) - 1, y_1\sigma(y_1)$.*

Example 2.1.3. Every system of algebraic equations $F \subset k[y_1, \dots, y_n] \subset k\{y_1, \dots, y_n\}$ over a σ -field k defines a k - σ -variety.

Definition 2.1.4. The functor \mathbb{A}_k^n given by $\mathbb{A}_k^n(K) = K^n$ for $K \in \sigma\text{-field}_k$ is called difference affine (n -)space over k . Obviously, $\mathbb{A}_k^n = \mathbb{V}(0)$ is a σ -variety over k .

Example 2.1.5. The following equations all define the same σ -variety, namely the point at the origin: $y_1, y_1^3, \sigma^2(y_1), y_1\sigma(y_1)$.

Non-example 2.1.6. Let X be the subfunctor of \mathbb{A}_k^1 given by $X(K) = K \setminus \{0\}$ for $K \in \sigma\text{-field}_k$. Then X is not a σ -variety.

Proof. If X was a σ -variety then $X = \mathbb{V}(F)$ for some $F \subset k\{y_1\}$. But if $f \in k\{y_1\}$ satisfies $f(a) = 0$ for all $a \in K \setminus \{0\}$, $K \in \sigma\text{-field}_k$ then $f = 0$. (Take $K = k\langle y_1 \rangle$ and $a = y_1$.) Therefore $X = \mathbb{V}(0) = \mathbb{A}_k^1$; a contradiction. \square

If X and Y are σ -varieties over k , then we write $X \subset Y$ to indicate that X is a subfunctor of Y . This simply means that $X(K) \subset Y(K)$ for every $K \in \sigma\text{-field}_k$. In this situation we also say that X is a σ -subvariety of Y .

Let X be σ -subvariety of \mathbb{A}_k^n . Then

$$\mathbb{I}(X) := \{f \in k\{y_1, \dots, y_n\} \mid f(a) = 0 \text{ for all } a \in X(K) \text{ and all } K \in \sigma\text{-field}_k\}$$

is called the *vanishing ideal* of X .

The following lemma explains why perfect difference ideals are ubiquitous in difference algebra.

Lemma 2.1.7. Let X be a σ -subvariety of \mathbb{A}_k^n . Then $\mathbb{I}(X)$ is a perfect σ -ideal of $k\{y_1, \dots, y_n\}$.

Proof. Obviously $\mathbb{I}(X)$ is an ideal. Let $f \in \mathbb{I}(X)$. Then $f(a) = 0$ for every $a \in X(K)$ and every σ -field extension K of k . As $\sigma(f)(a) = \sigma(f(a)) = \sigma(0) = 0$ we see that $\mathbb{I}(X)$ is a σ -ideal.

It remains to show that $\mathbb{I}(X)$ is perfect. Assume that $f \in k\{y\}$ with $f\sigma(f) \in \mathbb{I}(X)$. Then $(f\sigma(f))(a) = 0$ for every $a \in X(K)$ and every σ -field extension K of k . But $0 = (f\sigma(f))(a) = f(a)\sigma(f(a)) \in K$. Because K is a field one of the factors $f(a)$ or $\sigma(f(a))$ has to be zero. If $\sigma(f(a)) = 0$ we can also conclude that $f(a) = 0$ because σ is injective on K (since K is a field). Therefore $f \in \mathbb{I}(X)$ and $\mathbb{I}(X)$ is perfect. \square

Definition 2.1.8. Let X be a σ -subvariety of \mathbb{A}_k^n . Then the k - σ -algebra

$$k\{X\} := k\{y_1, \dots, y_n\}/\mathbb{I}(X)$$

is called the σ -coordinate ring of X .

Let K be a σ -field extension of k and $a \in X(K)$. If $f = \bar{g}$, $g \in k\{y_1, \dots, y_n\}$ is an element from $k\{X\}$, then $f(a) := g(a)$ does not depend on the choice of g : If $\bar{g}' = \bar{g}$ then $g' - g \in \mathbb{I}(X)$ and therefore $g'(a) = g(a)$.

We can thus think of $k\{X\}$ as a ring of functions on X . If $f \in k\{X\}$ vanishes on X , that is, if $f(a) = 0$ for all $a \in X(K)$ and all $K \in \sigma\text{-field}_k$, then $f = 0$.

Example 2.1.9. Let $a \in k^n$ and let $X \subset \mathbb{A}_k^n$ be the σ -variety defined by $y_1 - a_1, \dots, y_n - a_n$. Then $X(K) = \{a\}$ for $K \in \sigma\text{-field}_k$ and $k\{X\} = k$. Conversely, every k - σ -variety X with $k\{X\} = k$ is of this form.

Remark 2.1.10. Let X be a k - σ -variety. For $K \in \sigma\text{-field}_k$ there is a natural bijection between $X(K)$ and the set of k - σ -morphisms from $k\{X\}$ to K . Indeed,

$$X \simeq \text{Hom}(k\{X\}, -)$$

as functors.

Proof. Let \bar{y} denote the coordinate functions on X . If $a \in X(K)$ we can define a k - σ -morphism from $k\{X\}$ to K by sending \bar{y} to a . Conversely, given a k - σ -morphism $\phi: k\{X\} \rightarrow K$, then $a := \phi(\bar{y})$ belongs to $X(K)$. \square

Proposition 2.1.11. *Let $F \subset k\{y\} = k\{y_1, \dots, y_n\}$ be a system of algebraic difference equations over k . Then*

$$\mathbb{I}(\mathbb{V}(F)) = \{F\}.$$

Proof. Clearly $F \subset \mathbb{I}(\mathbb{V}(F))$. Since $\mathbb{I}(\mathbb{V}(F))$ is perfect by Lemma 2.1.7 we have $\{F\} \subset \mathbb{I}(\mathbb{V}(F))$.

Now let $f \in \mathbb{I}(\mathbb{V}(F))$. We have to show that $f \in \{F\}$. Because $\{F\}$ is the intersection of all σ -prime ideals of $k\{y\}$ which contain F (Proposition 1.2.22), it suffices to show that f lies in every σ -prime ideal \mathfrak{p} of $k\{y\}$ with $F \subset \mathfrak{p}$. Set $K := k(\mathfrak{p})$ and let a denote the image of y in K . Since $F \subset \mathfrak{p}$ we have $a \in \mathbb{V}_K(F)$. As $f \in \mathbb{I}(\mathbb{V}(F))$ this implies that $f(a) = 0$. Therefore $f \in \mathfrak{p}$. \square

Theorem 2.1.12. *The maps $X \mapsto \mathbb{I}(X)$ and $\mathfrak{a} \mapsto \mathbb{V}(\mathfrak{a})$ define inclusion reversing bijections between the set of all σ -subvarieties of \mathbb{A}_k^n and the set of all perfect σ -ideals in $k\{y_1, \dots, y_n\}$.*

Proof. Let $X = \mathbb{V}(F)$ be a σ -subvariety of \mathbb{A}_k^n . Because $F \subset \mathbb{I}(\mathbb{V}(F))$ we have

$$X = \mathbb{V}(F) \supset \mathbb{V}(\mathbb{I}(\mathbb{V}(F))) = \mathbb{V}(\mathbb{I}(X)).$$

The reverse inclusion, $X \subset \mathbb{V}(\mathbb{I}(X))$ is obvious. Therefore $\mathbb{V}(\mathbb{I}(X)) = X$.

Let \mathfrak{a} be a perfect σ -ideal of $k\{y\}$. From Proposition 2.1.11 we obtain $\mathbb{I}(\mathbb{V}(\mathfrak{a})) = \{\mathfrak{a}\} = \mathfrak{a}$. \square

The above theorem can be generalized to arbitrary σ -varieties in place of \mathbb{A}_k^n :

Corollary 2.1.13. *Let Y be a σ -variety. The map*

$$X \mapsto \{f \in k\{Y\} \mid f(a) = 0 \forall a \in X(K) \forall K \in \sigma\text{-field}_k\}$$

is an inclusion reversing bijection between the set of σ -subvarieties of Y and the set of perfect σ -ideals in $k\{Y\}$.

Proof. Assume $Y \subset \mathbb{A}_k^n$. The image of X under the above map is $\phi(\mathbb{I}(X))$ where

$$\phi: k\{y\} \rightarrow k\{y\}/\mathbb{I}(Y) = k\{Y\}$$

is the canonical map and $k\{y\} = k\{y_1, \dots, y_n\}$. The claim now follows from Theorem 2.1.12 and Proposition 1.2.8. \square

Definition 2.1.14. *Let X and Y be k - σ -varieties. Say $X \subset \mathbb{A}_k^n$ and $Y \subset \mathbb{A}_k^m$. A morphism of functors $f: X \rightarrow Y$ is called a σ -polynomial map (or morphism of k - σ -varieties) if there exist σ -polynomials $f_1, \dots, f_m \in k\{y_1, \dots, y_n\}$ such that $f(a) = (f_1(a), \dots, f_m(a))$ for every $a \in X(K)$ and all $K \in \sigma\text{-field}_k$.*

It is in principle possible to consider more general maps between σ -varieties than σ -polynomial maps: One could consider mappings which are locally given by fractions of σ -polynomials. However, the resulting category does not seem to have a nice algebraic description. It seems interesting to note that the standard textbooks ([Cohn], [Levin]) survive very well without ever defining the notion of a morphism of σ -varieties.

Note that the composition of two σ -polynomial maps is again a σ -polynomial map. The composition corresponds to composition of σ -polynomials.

Example 2.1.15. Let X be σ -subvariety of \mathbb{A}_k^n and $1 \leq m \leq n$. The projection

$$X \rightarrow \mathbb{A}_k^m, (a_1, \dots, a_n) \mapsto (a_1, \dots, a_m)$$

onto the first m coordinates is a morphism of σ -varieties.

Example 2.1.16. The map $f: \mathbb{A}_k^1 \rightarrow \mathbb{A}_k^1, a \mapsto a\sigma^3(a)^2 + \sigma(a)$ is an endomorphism of \mathbb{A}_k^1 .

Note the abuse of notation in the above examples: First of all, f is not a map but a functor, and we do not really specify where a belongs to. A pedant would need to say that f is given by $f(a) = a\sigma^3(a)^2 + \sigma(a)$ for all $a \in \mathbb{A}_k^1(K)$ and all $K \in \sigma\text{-field}_k$.

Example 2.1.17. Let Y be a σ -variety and X and σ -subvariety of Y . Then the inclusion map $X \rightarrow Y$ is a morphism of σ -varieties.

Example 2.1.18. Let X be k - σ -variety. The morphisms from X to \mathbb{A}_k^1 are in one-to-one correspondence with the elements in $k\{X\}$.

Non-example 2.1.19. Let $X \subset \mathbb{A}_k^1$ denote the σ -subvariety defined by $\sigma(y_1) = y_1 + 1$. For $a \in X(K)$ we have $a \neq 0$ (otherwise $0 = 1$). Therefore we can define a functor $f: X \rightarrow \mathbb{A}_k^1$ by setting $f(a) = \frac{1}{a}$ for $a \in X(K)$, $K \in \sigma\text{-field}_k$. Then f is not a morphism of σ -varieties.

Proof. Suppose f is given by a σ -polynomial $g \in k\{y_1\}$. Then $\bar{g} \cdot \bar{y}_1 = 1$ in $k\{X\}$. Note that $k\{X\}$ can be identified with the univariate polynomial ring $k[y_1]$ with action of σ given by $\sigma(y_1) = y_1 + 1$. Since y_1 is not invertible in $k[y_1]$ we obtain a contradiction. \square

The above non-example is in stark contrast to what happens in usual algebraic geometry: An everywhere defined rational function on an affine variety is necessarily polynomial. Note however that f does not extend to a functor on k - σ -algebras. I.e., you can not define $f(a)$ for $a \in X(R) := \{a \in R \mid \sigma(a) = a + 1\}$ where R is an arbitrary k - σ -algebra. In this sense f is not everywhere defined.

With the notation of Definition 2.1.14, note that the images $\overline{f_1}, \dots, \overline{f_m}$ of f_1, \dots, f_m in $k\{X\} = k\{y\}/\mathbb{I}(X)$ are uniquely determined by f : If $f_i(a) = \overline{g_i(a)}$ for all $a \in X(K)$ and all $K \in \sigma\text{-field}_k$ for some $g_i \in k\{y\}$, then $f_i - \overline{g_i} \in \mathbb{I}(X)$, i.e., $\overline{f_i} = \overline{g_i}$. Define

$$\phi: k\{z_1, \dots, z_m\} \rightarrow k\{X\}, z_i \mapsto \overline{f_i}.$$

Then for $a \in X(K)$ and $g \in k\{z_1, \dots, z_m\}$

$$\phi(g)(a) = g(f_1(a), \dots, f_m(a)) = g(f(a)).$$

Because $f(a) \in Y(K)$ we have $\phi(g)(a) = 0$ if $g \in \mathbb{I}(Y)$. So, $\phi(g) = 0$ for $g \in \mathbb{I}(Y)$. This shows that ϕ induces a morphism of k - σ -algebras

$$f^*: k\{Y\} = k\{z_1, \dots, z_m\}/\mathbb{I}(Y) \longrightarrow k\{X\}, \overline{z_i} \mapsto \overline{f_i}$$

We call f^* the morphism *dual* to f . The relation between f and f^* is given by

$$f^*(g)(a) = g(f(a)), g \in k\{Y\}, a \in X(K)$$

for every σ -field extension K of k .

If $X \xrightarrow{f} Y \xrightarrow{g} Z$ are morphisms of k - σ -varieties, it is easy to see that $(gf)^* = f^*g^*$. This means that $X \mapsto k\{X\}$ and $f \mapsto f^*$ defines a (contravariant) functor from the category of k - σ -varieties to the category of k - σ -algebras.

Definition 2.1.20. A σ -ring R is called *perfectly reduced* if the zero ideal of R is perfect.

Theorem 2.1.21. Let k be a σ -field. The category of k - σ -varieties is (anti-)equivalent to the category of perfectly reduced k - σ -algebras which are finitely σ -generated over k .

Proof. By Lemma 2.1.7 the σ -coordinate ring $k\{X\}$ of a σ -variety X is perfectly reduced. It is also clear that $k\{X\}$ is finitely σ -generated over k . Let X and Y be k - σ -varieties. We will show that

$$\mathrm{Hom}(X, Y) \rightarrow \mathrm{Hom}(k\{Y\}, k\{X\}), f \mapsto f^*$$

is bijective. We first show the injectivity. So let $f, g \in \mathrm{Hom}(X, Y)$ with $f^* = g^*$. Then

$$h(f(a)) = f^*(h)(a) = g^*(h)(a) = h(g(a))$$

for every $h \in k\{Y\}, a \in X(K)$ and $K \in \sigma\text{-field}_k$. For example, h could be one of the coordinate functions. This shows that $f(a) = g(a)$, so $f = g$.

Now for surjectivity: Let $\phi: k\{Y\} \rightarrow k\{X\}$ be a morphism of k - σ -algebras. Assume $X \subset \mathbb{A}_k^n$ and $Y \subset \mathbb{A}_k^m$. So $k\{Y\} = k\{z_1, \dots, z_m\}/\mathbb{I}(Y)$. Let $f_1, \dots, f_m \in k\{y_1, \dots, y_n\}$ such that

$$\phi(\overline{z_i}) = \overline{f_i} \in k\{y_1, \dots, y_n\}/\mathbb{I}(X) = k\{X\} \text{ for } i = 1, \dots, m.$$

Define a morphism $f: X \rightarrow \mathbb{A}_k^m$ by

$$f(a) = (f_1(a), \dots, f_m(a)) \in \mathbb{A}_k^m(K)$$

for $a \in X(K)$ and $K \in \sigma\text{-field}_k$. We will show that f is actually mapping into Y . It suffices to show that $h(f(a)) = 0$ for $h \in \mathbb{I}(Y)$. By definition $h(\overline{z_1}, \dots, \overline{z_m}) = 0$. Because ϕ is a morphism of k - σ -algebras this implies $h(\overline{f_1}, \dots, \overline{f_m}) = 0$. But then also $h(f(a)) = 0$ for all $a \in X(K), K \in \sigma\text{-field}_k$.

We thus obtain a morphism $f: X \rightarrow Y$ and it is clear from the construction of f that $f^* = \phi$.

It remains to see that $X \mapsto k\{X\}$ is essentially surjective: Let R be a perfectly reduced finitely σ -generated k - σ -algebra. Since R is finitely σ -generated over k we can write R as

$$R \simeq k\{y\}/\mathfrak{a}$$

for some σ -polynomial ring $k\{y\} = k\{y_1, \dots, y_n\}$ and σ -ideal $\mathfrak{a} \subset k\{y\}$. As R is perfectly reduced \mathfrak{a} is perfect. Let $X := \mathbb{V}(\mathfrak{a})$. Then $\mathbb{I}(X) = \mathfrak{a}$ by Proposition 2.1.11. Therefore $R \simeq k\{X\}$. \square

Corollary 2.1.22. Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. Then f is an isomorphism if and only if $f^*: k\{Y\} \rightarrow k\{X\}$ is an isomorphism.

Proof. I guess you know that. \square

2.2 The topological space of a difference variety

In this section we explain how we can view the solutions of a system of algebraic difference equations as a topological space.

Let k be a σ -field and $F \subset k\{y_1, \dots, y_n\}$ a system of algebraic difference equations over k . Let K and L be σ -field extensions of k . Two solutions $a \in \mathbb{V}_K(F)$ and $b \in \mathbb{V}_L(F)$ are called *equivalent* if there exists a k - σ -isomorphism between $k\langle a \rangle$ and $k\langle b \rangle$ which maps a to b . Obviously this defines an equivalence relation. There is really not much point in distinguishing between two equivalent solutions. The following theorem gives a natural way to choose a single element from each equivalence class.

Theorem 2.2.1. *Let $X = \mathbb{V}(F)$ be a k - σ -variety. There is a natural bijection between the set of equivalence classes of solutions of F and $\text{Spec}^\sigma(k\{X\})$.*

Proof. If $a \in K^n$ is a solution of F then F lies in the kernel of $k\{y\} \rightarrow K$, $y \mapsto a$. As the kernel is a σ -prime ideal (K is a σ -field!) it contains $\{F\} = \mathbb{I}(X)$. Therefore we obtain an induced k - σ -morphism $k\{X\} \rightarrow K$, $\bar{y} \mapsto a$. The kernel of this map, call it \mathfrak{p}_a , is a σ -prime ideal of $k\{X\}$. It is clear that \mathfrak{p}_a only depends on the equivalence class of a .

Conversely, if \mathfrak{p} is a σ -prime ideal of $k\{X\}$, set $K = k(\mathfrak{p})$ and let $a_{\mathfrak{p}} \in K^n$ denote the image of y . Then $a_{\mathfrak{p}}$ is a solution of F .

It is immediate that these two constructions are inverse to each other. \square

Definition 2.2.2. *Let k be a σ -field and X a k - σ -variety. The topological space of X is $\text{Spec}^\sigma(k\{X\})$ equipped with the Cohn topology.*

We shall not make a big fuss about distinguishing between a σ -variety and its topological space. For example if we speak of the closed subsets of a σ -variety X , we of course mean the closed subsets of $\text{Spec}^\sigma(k\{X\})$. We write $x \in X$ to mean that x is a point of the topological space of X . From a formal point of view this is simply saying that x is a σ -prime ideal of $k\{X\}$, but this is not what you should have in mind. The notation $x \in X$ suggests that you should think of x as a point on some topological space. If we want to speak about σ -prime ideals of $k\{X\}$ we will rather say something like “Let \mathfrak{p} be a σ -prime ideal of $k\{X\}$.”

A σ -variety is called irreducible, if its topological space is irreducible. Note that the closed subsets of X are in one-to-one correspondence with the σ -subvarieties of X . (Clear from Proposition 1.2.35 (ii) and Corollary 2.1.13.)

Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. If $a \in X(K)$ and $b \in X(L)$ ($K, L \in \sigma\text{-field}_k$) are equivalent, then also $f(a)$ and $f(b)$ are equivalent. It follows that $a \mapsto f(a)$ induces a map from the equivalence classes of solutions of $\mathbb{I}(X)$ to the equivalence classes of solutions of $\mathbb{I}(Y)$, which by Theorem 2.2.1 corresponds to a map $\text{Spec}^\sigma(k\{X\}) \rightarrow \text{Spec}^\sigma(k\{Y\})$. On the other hand, by Theorem 2.1.21 the morphism $f: X \rightarrow Y$ corresponds to a morphism $f^*: k\{Y\} \rightarrow k\{X\}$ which, by Lemma 1.2.37, also corresponds to a map $\text{Spec}^\sigma(k\{X\}) \rightarrow \text{Spec}^\sigma(k\{Y\})$. These two maps from $\text{Spec}^\sigma(k\{X\})$ to $\text{Spec}^\sigma(k\{Y\})$ are the same.

Exercise 2.2.3. *Check that last sentence.*

In particular, every morphism $f: X \rightarrow Y$ of k - σ -varieties induces a continuous map on the corresponding topological spaces. Indeed, we have a functor from k - σ -varieties to topological spaces.

Example 2.2.4. Let $k = \mathbb{C}$ considered as a constant σ -field. Let $X \subset \mathbb{A}_k^1$ be the σ -subvariety defined by $\sigma(y_1) = y_1 + 1$. Then the topological space of X is a single point.

Proof. The σ -coordinate ring of X is $\mathbb{C}[y_1]$ (cf. Example 2.1.19) which has only one σ -prime ideal, namely the zero ideal, by Exercise 1.1.6. \square

Note that the topological space of the σ -variety in Example 2.1.9 also is just a point. So, from a formal point of view, the σ -varieties in Example 2.1.9 and in Example 2.2.4 have the same topological space. It however makes kind of sense to think that the point of the σ -variety in Example 2.2.4 is a somewhat fat point. This is in analogy to usual algebraic geometry where you can think of $\text{Spec}(k')$ with k' a finite field extension of k or $\text{Spec}(k[y_1]/(y_1^2))$ as a point which is fatter than $\text{Spec}(k)$. However, in difference algebraic geometry already points can be quite complicated...

Exercise 2.2.5. Let $k = \mathbb{C}$ considered as a constant σ -field. Let $X \subset \mathbb{A}_k^1$ be the σ -subvariety defined by $\sigma(y_1) = y_1$. How does the topological space of X look like?

2.3 Morphisms of difference varieties

In this section we show how some topological properties of morphisms of σ -varieties are reflected by the dual morphism on the σ -coordinate rings.

Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. The image of the corresponding map on topological spaces is denoted by

$$\mathrm{Im}(f) \subset \mathrm{Spec}^\sigma(k\{Y\}).$$

Definition 2.3.1. A morphism $f: X \rightarrow Y$ of σ -varieties is called *dominant* if $\mathrm{Im}(f)$ is dense in Y .

Lemma 2.3.2. Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. The closure of $\mathrm{Im}(f)$ is $\mathcal{V}(\mathfrak{a})$ where \mathfrak{a} is the kernel of $f^*: k\{Y\} \rightarrow k\{X\}$.

Proof. Since $X = \mathcal{V}(0)$ the closure of $\mathrm{Im}(f)$ equals $\mathcal{V}(f^{*-1}(0))$ by Lemma 1.2.40. \square

Proposition 2.3.3. Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. Then f is dominant if and only if the dual map $f^*: k\{Y\} \rightarrow k\{X\}$ is injective.

Proof. Let \mathfrak{a} denote the kernel of $f^*: k\{Y\} \rightarrow k\{X\}$. If f is dominant, then $\mathcal{V}(\mathfrak{a}) = Y = \mathcal{V}(0)$ by Lemma 2.3.2. Since $k\{Y\}$ is perfectly reduced we must have $\mathfrak{a} = 0$ by Proposition 1.2.35 (ii). So f^* is injective.

Conversely, if f^* is injective then $\overline{\mathrm{Im}(f)} = \mathcal{V}(0) = Y$ by Lemma 2.3.2. \square

Definition 2.3.4. Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. Setting

$$f(X)(K) := f(X(K)) \subset Y(K) \text{ for } K \in \sigma\text{-field}_k$$

defines a subfunctor $f(X)$ of Y . We say that f is a *closed embedding* if $f(X)$ is a σ -subvariety of Y (i.e., $f(X)$ is a σ -variety) and the induced map $f: X \rightarrow f(X)$ is an isomorphism of σ -varieties.

Example 2.3.5. Let X be a σ -subvariety of Y . Then the canonical map $X \rightarrow Y$ is a closed embedding. Every closed embedding is isomorphic to a closed embedding of this form.

Non-example 2.3.6. Let $X \subset \mathbb{A}_k^2$ be the σ -variety defined by $y_1 y_2 = 1$ and

$$f: X \rightarrow \mathbb{A}_k^1, (a_1, a_2) \mapsto a_1.$$

Then $f(X) \subset \mathbb{A}_k^1$ is not a σ -variety.

Proof. Indeed $f(X)$ is the functor defined in Non-example 2.1.6. \square

Note that the above non-example is rather of an algebraic nature, than of a difference algebraic nature.

Exercise 2.3.7. Let $f: X \rightarrow Y$ be a morphism of σ -varieties. Show that there is a bijection between the equivalence classes of elements from $f(X)$ and $\mathrm{Im}(f)$.

Proposition 2.3.8. *Let $f: X \rightarrow Y$ be a morphism of k - σ -varieties. Then f is a closed embedding if and only if $f^*: k\{Y\} \rightarrow k\{X\}$ is surjective.*

Proof. Let \mathfrak{a} denote the kernel of $f^*: k\{Y\} \rightarrow k\{X\}$ and let Z denote the σ -subvariety of Y corresponding to \mathfrak{a} (cf. Corollary 2.1.13). Then $k\{Z\}$ can be identified with $k\{Y\}/\mathfrak{a}$.

First assume that f is a closed embedding. By Lemma 2.3.2 the closure of $\text{Im}(f)$ equals $\mathcal{V}(\mathfrak{a})$. By assumption $\text{Im}(f)$ is closed, so $\text{Im}(f) = \mathcal{V}(\mathfrak{a})$. The dual of the induced map $f: X \rightarrow f(X) = Z$ is simply given by

$$\phi: k\{Z\} = k\{Y\}/\mathfrak{a} \longrightarrow k\{X\}, \quad \bar{g} \mapsto f^*(g)$$

for $g \in k\{Y\}$. By assumption ϕ is an isomorphism. Therefore f^* is surjective.

Now assume that $f^*: k\{Y\} \rightarrow k\{X\}$ is surjective. Using Proposition 1.2.8 we find that

$$\text{Im}(f) = \{f^{*-1}(\mathfrak{p}) \mid \mathfrak{p} \in \text{Spec}^\sigma(k\{X\})\} = \mathcal{V}(\mathfrak{a}).$$

In particular $\text{Im}(f)$ is closed and $f(X) = Z$. The dual map of $f: X \rightarrow f(X) = Z$ is again ϕ as defined above. Since f^* is surjective, ϕ is an isomorphism. Therefore $f: X \rightarrow f(X) = Z$ is an isomorphism. \square

Exercise 2.3.9. *Show that every morphism of σ -varieties can be factored as a dominant morphism followed by a closed embedding.*

Exercise 2.3.10. *Let $\phi: R \rightarrow S$ be a morphism of σ -rings. Is it true that ϕ induces an homeomorphism between $\text{Spec}^\sigma(S)$ and a closed subset of $\text{Spec}^\sigma(R)$ if and only if ϕ is surjective?*

Chapter 3

Around the basis theorem

The main result of this chapter is the Ritt-Raudenbusch basis theorem, which can be seen as a σ -analog of Hilbert's basis theorem. Hilbert's basis theorem, in its most elementary form, states that the polynomial ring $k[y_1, \dots, y_n]$ over a field k is Noetherian. I.e., every ideal of $k[y_1, \dots, y_n]$ is finitely generated. Of course the σ -polynomial ring $k\{y_1, \dots, y_n\}$ over a σ -field k is not Noetherian, for example $(y_1) \subset (y_1, \sigma(y_1)) \subset \dots$ is an infinite ascending chain of ideals. However, one might hope that $k\{y_1, \dots, y_n\}$ is " σ -Noetherian", in the sense that every σ -ideal of $k\{y_1, \dots, y_n\}$ is finitely σ -generated. The following example shatters that hope.

Example 3.0.11. Let k be a σ -field and \mathfrak{a} the σ -ideal of $k\{y_1\}$ σ -generated by

$$y_1\sigma(y_1), y_1\sigma^2(y_1), y_1\sigma^3(y_1), \dots$$

Then \mathfrak{a} is not finitely σ -generated.

Proof. Suppose that \mathfrak{a} is finitely σ -generated. Then \mathfrak{a} is σ -generated by $y_1\sigma(y_1), \dots, y_1\sigma^n(y_1)$ for some $n \in \mathbb{N}_{\geq 1}$. We will show that $y_1\sigma^{n+1}(y_1) \notin [y_1\sigma(y_1), \dots, y_1\sigma^n(y_1)] = \mathfrak{a}$. Suppose the contrary. Observe that when we consider $y_1, \sigma(y_1), \dots$, as usual variables over k then $\sigma^j(y_1\sigma^i(y_1))$ is homogeneous of degree two. This implies that $y_1\sigma^{n+1}(y_1)$ is a k -linear combination of polynomials of the form $\sigma^j(y_1\sigma^i(y_1))$ ($1 \leq i \leq n, j \in \mathbb{N}$). But then $y_1\sigma^{n+1}(y_1)$ must already be equal to some $\sigma^j(y_1\sigma^i(y_1))$ with $1 \leq i \leq n, j \in \mathbb{N}$. Clearly this is not possible. \square

Obviously the ideal \mathfrak{a} in the above example is not perfect. So, we can still hope that $k\{y_1, \dots, y_n\}$ satisfies the ascending chain condition on perfect σ -ideals. That this is true is the content of the Ritt-Raudenbusch basis theorem. The geometric meaning of the basis theorem is that the topological space of a σ -variety X is a Noetherian topological space. This immediately yields a decomposition of X into finitely many irreducible components.

Hilbert's basis theorem asserts that every ascending chain of ideals in $k[y_1, \dots, y_n]$ is finite and not only that every ascending chain of radical ideals in $k[y_1, \dots, y_n]$ is finite (which is all you need to conclude that usual affine varieties are Noetherian). This strong property allows you to obtain a primary decomposition for arbitrary ideals in $k[y_1, \dots, y_n]$. The problem of " σ -primary decomposition" is largely open. For example, the question whether every ascending chain of radical, well-mixed σ -ideals in $k\{y_1, \dots, y_n\}$ is finite has been open for several years now. See [Hrushovski, Section 4.6]. It was recently shown by A. Levin that $k\{y_1, \dots, y_n\}$ does not satisfy the ascending chain condition for well-mixed σ -ideals.

3.1 Characteristic sets

Characteristic sets are the main computational tool in difference algebra. However, occasionally, they can also provide some theoretical insight. For example, the use of characteristic sets makes the proof of the Ritt-Raudenbusch basis theorem (Theorem 3.3.8) more transparent. (See Theorem 3.2.6 for another theoretical application of characteristic sets.) The idea behind characteristic sets is a little bit similar to the notion of a Gröbner basis. However, the true analog of characteristic sets in algebraic geometry are the (Ritt) characteristic sets (of algebraic ideals).

Exercise 3.1.1. Search the literature¹ for (Ritt) characteristic sets of ideals in a (usual) polynomial ring and present your main findings in class.

Let R be a σ -ring. The σ -polynomial ring $R\{y\} = R\{y_1, \dots, y_n\}$ over R in the σ -variables y_1, \dots, y_n can be viewed as a polynomial ring over R in the variables

$$\Theta y := \{\sigma^i(y_j) \mid i \geq 0, j = 1, \dots, n\}.$$

A total order \leq on Θy is called a *ranking* if the following properties are satisfied:

- (i) We have $u \leq \sigma(u)$ for every $u \in \Theta y$.
- (ii) If $u, v \in \Theta y$ with $u \leq v$, then $\sigma(u) \leq \sigma(v)$.

A ranking is called *orderly* if $\text{ord}(u) < \text{ord}(v)$ implies $u < v$ for $u, v \in \Theta y$. (The notation $u < v$ is short hand for $u \leq v$ and $u \neq v$.)

Example 3.1.2. The *standard ranking* is defined by:

$$\sigma^i(y_j) \leq \sigma^{i'}(y_{j'}) \text{ if and only if } i < i' \text{ or } i = i' \text{ and } j \leq j'.$$

This is an example of an orderly ranking.

Example 3.1.3. The *variable ranking* is defined by:

$$\sigma^i(y_j) \leq \sigma^{i'}(y_{j'}) \text{ if and only if } j < j' \text{ or } j = j' \text{ and } i \leq i'.$$

This ranking is not orderly if $n > 1$.

Lemma 3.1.4. Every ranking is a well-ordering. I.e., every non-empty subset of Θy has a least element.

Proof. Let U be a non-empty subset of Θy . For $j \in \{1, \dots, n\}$ such that there exists $i \in \mathbb{N}$ with $\sigma^i(y_j) \in U$, let i be minimal with this property and set $u_j = \sigma^i(y_j)$. Then the minimum of the u_j 's is a least element of U . \square

Exercise 3.1.5. Find some more examples of rankings.

Until the end of Section 3.1 we assume that a ranking has been fixed. The *leader* $u_f \in \Theta y$ of a σ -polynomial $f \in R\{y\} \setminus R$ (with respect to the given ranking) is the largest element of Θy which appears effectively in f . We can write f (uniquely) as a polynomial in u_f :

$$f = \sum_{i=0}^d g_i u_f^i \text{ with } g_i \in R\{y\} \text{ not involving } u_f \text{ and } g_d \neq 0.$$

¹It is not a bad idea to start with Wikipedia.

The σ -polynomial $I_f := g_d$ is called the *initial* of f (with respect to the given ranking). The pair

$$\text{rk}(f) := (u_f, d) = (u_f, \deg_{u_f}(f))$$

is called the *rank* of f . We can use the given ranking to define a total order on the set of all ranks, $\Theta y \times \mathbb{N}_{\geq 1}$:

$$(u, d) \leq (u', d') \text{ if and only if } u < u' \text{ or } u = u' \text{ and } d \leq d'.$$

Clearly, this defines a well-ordering on the set of all ranks. In particular, every non-empty subset of $R\{y\} \setminus R$ contains an element of lowest rank.

For $f \in R\{y\} \setminus R$ with $I_f \notin R$ we have $\text{rk}(I_f) < \text{rk}(f)$. Note that condition (iii) in the definition of a ranking implies that $u_{\sigma(f)} \leq \sigma(u_f)$ for $f \in R\{y\} \setminus R$ with $\sigma(f) \notin R$. If R is a σ -domain we have $u_{\sigma(f)} = \sigma(u_f)$ and $I_{\sigma(f)} = \sigma(I_f)$.

A σ -polynomial $g \in R\{y\}$ is called *reduced* with respect to a σ -polynomial $f \in R\{y\} \setminus R$ if g does not contain a power of $\sigma^i(u_f)$ whose exponent is greater than or equal to $\deg_{u_f}(f)$. It follows from condition (ii) in the definition of a ranking that g is reduced with respect to f if $\text{rk}(g) < \text{rk}(f)$. Moreover, g is called reduced with respect to a subset F of $R\{y\} \setminus R$ if g is reduced with respect to every element of F . Finally, a subset F of $R\{y\} \setminus R$ is called *autoreduced* if every element of F is reduced with respect to all the other elements of F . (By definition, the empty set is autoreduced.)

Lemma 3.1.6. *Every autoreduced subset of $R\{y\}$ is finite.*

Proof. Suppose that $F \subset R\{y\}$ is an infinite autoreduced subset. Clearly, the leaders of distinct elements of F are distinct. Since the set of leaders of elements of F is infinite there must exist an index $j \in \{1, \dots, n\}$ and an infinite subset $I \subset \mathbb{N}$ such that $\sigma^i(x_j)$ is a leader of an element of F for every $i \in I$. Say, $\sigma^i(x_j) = u_{f_i}$ ($f_i \in F$). If $k < l$ are elements from I then $\deg_{\sigma^k(x_j)}(f_k) > \deg_{\sigma^l(x_j)}(f_l)$ because f_l is reduced with respect to f_k . But then $(\deg_{\sigma^i(x_j)}(f_i))_{i \in I}$ would be an infinite strictly decreasing sequence of natural numbers. \square

In the sequel we will always write the elements of an autoreduced set in the order of increasing rank, i.e., if $F = \{f_1, \dots, f_p\}$ is an autoreduced set then $\text{rk}(f_1) < \dots < \text{rk}(f_p)$. If $F = \{f_1, \dots, f_p\}$ and $G = \{g_1, \dots, g_q\}$ are autoreduced sets we say that F has lower rank than G and write $\text{rk}(F) < \text{rk}(G)$ if and only if one of the following conditions is satisfied:

- There exists an $m \in \mathbb{N}$ with $1 \leq m \leq \min\{p, q\}$, such that $\text{rk}(f_1) = \text{rk}(g_1), \dots, \text{rk}(f_{m-1}) = \text{rk}(g_{m-1})$ and $\text{rk}(f_m) < \text{rk}(g_m)$.
- $p > q$ and $\text{rk}(f_1) = \text{rk}(g_1), \dots, \text{rk}(f_q) = \text{rk}(g_q)$.

If neither $\text{rk}(F) < \text{rk}(G)$ nor $\text{rk}(G) < \text{rk}(F)$ we say that F and G are of the same rank and write $\text{rk}(F) = \text{rk}(G)$. This is the case if and only if $p = q$ and $\text{rk}(f_1) = \text{rk}(g_1), \dots, \text{rk}(f_p) = \text{rk}(g_p)$.

Proposition 3.1.7. *In every non-empty set of autoreduced subsets of $R\{y\}$ there exists an autoreduced subset of lowest rank.*

Proof. It suffices to show that there is no infinite sequence F_1, F_2, \dots of autoreduced sets with $\text{rk}(F_1) > \text{rk}(F_2) > \dots$. Write $F_i = \{f_1^i, \dots, f_{p_i}^i\}$ with $\text{rk}(f_1^i) < \dots < \text{rk}(f_{p_i}^i)$. Then $\text{rk}(f_1^1) \geq \text{rk}(f_1^2) \geq \dots$. This sequence must eventually stabilize, say $\text{rk}(f_1^{i_1}) = \text{rk}(f_1^{i_1+1}) = \dots$. Similarly, the sequence $\text{rk}(f_2^{i_1}) \geq \text{rk}(f_2^{i_1+1}) \geq \dots$ must eventually stabilize, i.e., $\text{rk}(f_2^{i_2}) = \text{rk}(f_2^{i_2+1}) = \dots$. Continuing in this way we obtain an infinite decreasing sequence $\text{rk}(f_1^{i_1}) > \text{rk}(f_2^{i_2}) > \dots$; a contradiction. \square

Let \mathfrak{a} be a σ -ideal of $R\{y\}$. Then \mathfrak{a} contains an autoreduced set, for example the empty set. If $f \in \mathfrak{a} \setminus R$ then also $\{f\}$ is an autoreduced subset of \mathfrak{a} . By Proposition 3.1.7 the following definition makes sense:

Definition 3.1.8. *Let \mathfrak{a} be a σ -ideal of $R\{y\}$. An autoreduced subset of \mathfrak{a} of lowest rank is called a characteristic set of \mathfrak{a} (with respect to the given ranking).*

A characteristic set of \mathfrak{a} (with respect to a given ranking) need not be unique, for example we can multiply the elements of a characteristic set with a unit from R . However, if F and G are characteristic sets of \mathfrak{a} then $\text{rk}(F) = \text{rk}(G)$.

A characteristic set of a σ -ideal \mathfrak{a} of $R\{y\}$ can be obtained by the following (non-constructive) procedure: Choose $f_1 \in M_1 := \mathfrak{a} \setminus R$ of minimal rank. Let M_2 denote the set of all σ -polynomials in $\mathfrak{a} \setminus R$ which are reduced with respect to f_1 . Choose f_2 of minimal rank in M_2 . Then f_1, f_2 is autoreduced. (f_1 is reduced with respect to f_2 since $\text{rk}(f_1) < \text{rk}(f_2)$.) Let M_3 denote the set of all σ -polynomials in $\mathfrak{a} \setminus R$ which are reduced with respect to f_1, f_2 . Choose $f_3 \in M_3$ of minimal rank. Then f_1, f_2, f_3 is autoreduced. Continue like this. The process must terminate because an autoreduced set is finite. In the end we have obtained an autoreduced subset $F = \{f_1, \dots, f_p\}$ of \mathfrak{a} such that no σ -polynomial of $\mathfrak{a} \setminus R$ is reduced with respect² to F . It is clear from the construction of F that F is a characteristic set of \mathfrak{a} .

Definition 3.1.9. *Let $F = \{f_1, \dots, f_p\}$ be an autoreduced set. The multiplicatively closed σ -stable subset of $R\{y\}$ generated by the initials I_{f_1}, \dots, I_{f_p} is denoted by I_F . (See Subsection 1.1.2.)*

⚠ 3.1.10. *As we do not assume that R is a σ -domain it can happen that $0 \in S_F$. In this case Prop. 3.1.11 below is trivial.*

The following proposition is a version of division with remainder for difference polynomials.

Proposition 3.1.11. *Let $F = \{f_1, \dots, f_p\}$ be an autoreduced set and $g \in R\{y\}$. Then there exist σ -polynomials $g_0 \in R\{y\}$ and $h \in I_F$ such that g_0 is reduced with respect to F and $hg - g_0 \in [F]$.*

Proof. For a difference polynomial g which is not reduced with respect to F we define the F -rank of g as the largest rank (u, d) such that u^d appears in g and $u = \sigma^i(u_f)$, $\deg_{u_f}(f) \leq d$ for some $f \in F$.

Suppose, for a contradiction, that the statement of the proposition is false and let M denote the set of all $g \in R\{y\}$ such that the statement is false for g . An element $g \in M$ can not be reduced with respect to F because in this case we could take $h = 1$ and $g_0 = g$ to satisfy the proposition. Choose a $g \in M$ with minimal F -rank (u, d) . Then we can write $g = g_1 u^d + g_2$, where g_1 is a σ -polynomial not involving u and g_2 is a σ -polynomial with $\deg_u(g_2) < d$. Choose $f \in F$ with $u = \sigma^i(u_f)$ and $\deg_{u_f}(f) \leq d$. Consider the σ -polynomial

$$g' := \sigma^i(I_f)g - g_1 u^{d - \deg_{u_f}(f)} \sigma^i(f).$$

We claim that $g' \notin M$: If g' is reduced with respect to F then $g' \notin M$. So we can assume that g' is not reduced with respect to F . It suffices to show that the F -rank of g' is strictly smaller than the F -rank of g : If $\sigma^i(I_f) \notin R$ then $u_{\sigma^i(I_f)} \leq \sigma^i(u_{I_f}) < \sigma^i(u_f) = u$, consequently the F -rank of $\sigma^i(I_f)g$ is at most the F -rank of g . If $\sigma^i(f) \notin R$ then $u_{\sigma^i(f)} \leq$

²Theorem 3.1.14 (i) gives another proof of this property of characteristic sets.

$\sigma^i(u_f) = u$. Also, the F -rank of g_1 is lesser than the F -rank of g . Consequently, the F -rank of $g_1 u^{d - \deg_{u_f}(f)} \sigma^i(f)$ is at most the F -rank of g . This shows that the F -rank of g' is at most the F -rank of g . By construction $\deg_u(g') < \deg_u(g)$. Therefore, the F -rank of g' is less than the F -rank of g and so $g' \notin M$.

Thus, there exist σ -polynomials $h' \in I_F$ and $g'_0 \in R\{y\}$ such that g'_0 is reduced with respect to F and $h'g' \equiv g'_0 \pmod{[F]}$. As $g' \equiv \sigma^i(I_f)g \pmod{[F]}$ we find that $h'\sigma^i(I_f)g \equiv g'_0 \pmod{[F]}$. Now the proposition is satisfied with $h := h'\sigma^i(I_f)$ and $g_0 := g'_0$; a contradiction to $g \in M$. \square

The element g_0 in the above proposition is sometimes called a *remainder* of g with respect to F .

Exercise 3.1.12. * Give an algorithm to compute a remainder with respect to an autoreduced set. Implement this in your favorite computer algebra system.

Ex 3.1.13. The set of ranks $\{\text{rk}(f) \mid f \in \mathfrak{a}\}$ is not a reasonable invariant of a σ -ideal \mathfrak{a} . For example, consider the difference ideals $[y_1]$ and $[y_1, y_2]$ in the σ -polynomial ring $k\{y\} = k\{y_1, y_2\}$ over some σ -field k . Then, with respect to the standard ranking,

$$\{\text{rk}(f) \mid f \in [y_1], f \neq 0\} = \Theta y \times \mathbb{N}_{\geq 1} = \{\text{rk}(f) \mid f \in [y_1, y_2], f \neq 0\}.$$

(Note that $y_1 y_2 \in [y_1]$ and $\text{rk}(y_1 y_2) = (y_2, 1)$.) So, $[y_1]$ and $[y_1, y_2]$ have the same ranks. However, $\{y_1\}$ is a characteristic set of $[y_1]$ and $\{y_1, y_2\}$ is a characteristic set of $[y_1, y_2]$ (Check this!). In particular, the cardinality of a characteristic set of \mathfrak{a} can not be read from the set of ranks of \mathfrak{a} .

Theorem 3.1.14. Let F be a characteristic set of a σ -ideal \mathfrak{a} of $R\{y\}$.

- (i) If $g \in \mathfrak{a}$ is reduced with respect to F then $g \in R$.
- (ii) In particular, if $\mathfrak{a} \cap R = \{0\}$, then 0 is the only element of \mathfrak{a} which is reduced with respect to F and $I_f \notin \mathfrak{a}$ for $f \in F$. Moreover,

$$[F] \subset \mathfrak{a} \subset [F] : I_F.$$

If $\mathfrak{a} : I_F = \mathfrak{a}$ then $\mathfrak{a} = [F] : I_F$.

Proof. If $g \in \mathfrak{a}$ is reduced with respect to F but $g \notin R$, then $G := \{f \in F \mid u_f < u_g\} \cup \{g\}$ is an autoreduced subset of \mathfrak{a} with $\text{rk}(G) < \text{rk}(F)$; In contradiction to the assumption that F is a characteristic set of \mathfrak{a} . This proves (i).

Now for (ii): Since I_f is reduced with respect to F we have $I_f \notin \mathfrak{a}$ by (i). Let $g \in \mathfrak{a}$, $g \neq 0$. By Prop. 3.1.11 there exists $h \in I_F$ and $g_0 \in R\{y\}$ such that g_0 is reduced with respect to F and $hg - g_0 \in [F] \subset \mathfrak{a}$. Then $g_0 \in \mathfrak{a}$ is reduced with respect to $[F]$. From (i) we obtain $g_0 = 0$. Consequently, $hg \in [F]$. I.e., $g \in [F] : I_F$.

If $\mathfrak{a} : I_F = \mathfrak{a}$ then $[F] : I_F \subset \mathfrak{a} : I_F = \mathfrak{a}$. \square

Corollary 3.1.15. If F is a characteristic set of a σ -prime ideal $\mathfrak{p} \subset R\{y\}$ with³ $\mathfrak{p} \cap R = \{0\}$, then

$$\mathfrak{p} = [F] : I_F.$$

Proof. This is clear from (ii) of Theorem 3.1.14: As $I_f \notin \mathfrak{p}$ for $f \in F$ and \mathfrak{p} is σ -prime it follows that $\mathfrak{p} : I_F = \mathfrak{p}$. \square

Exercise 3.1.16. * Is f a characteristic set of $[f]$?

³This implies that R is a σ -domain.

3.2 Difference algebras of finite presentation

In this section we present a coordinate free version of Corollary 3.1.15.

Definition 3.2.1. Let R be a σ -ring and S an R - σ -algebra. We say that S is finitely σ -presented over R , if S is isomorphic to a quotient of a σ -polynomial ring over R modulo a finitely σ -generated σ -ideal. That is, $S \simeq R\{y_1, \dots, y_n\}/[f_1, \dots, f_m]$.

Example 3.2.2. Let k be a σ -field and R a finitely generated k -algebra. Then $[\sigma]_k R$ is finitely σ -presented over k .

Exercise 3.2.3. Let $R \rightarrow S \rightarrow T$ be morphisms of σ -rings. If S is finitely σ -presented over R and T is finitely σ -presented over S , then T is finitely σ -presented over R .

Our first concern is to show that a finitely σ -presented difference algebra is “always” finitely σ -presented:

Lemma 3.2.4. Let R be a σ -ring, S an R - σ -algebra and $\phi: R\{x_1, \dots, x_l\} \rightarrow S$ a surjective morphism of R - σ -algebras from a σ -polynomial ring over R onto S . If S is finitely σ -presented over R , then the kernel of ϕ is a finitely σ -generated σ -ideal.

Proof. We identify S with $R\{y_1, \dots, y_n\}/[f_1, \dots, f_m]$. Choose $g_i \in R\{y\}$ such that $\phi(x_i) = \overline{g_i}$ for $i = 1, \dots, l$. Then $S \simeq R\{x, y\}/[x_i - g_i, f_j]$, $\overline{y_j} \mapsto \overline{y_j}$; this is clear because the inverse map is given by $\overline{x_i} \mapsto \overline{g_i}$, $\overline{y_j} \mapsto \overline{y_j}$.

Choose $h_j \in R\{x\}$ such that $\phi(h_j) = \overline{y_j}$ for $j = 1, \dots, n$ and consider the map $\psi: R\{x, y\} \rightarrow R\{x\}$, $x_i \mapsto x_i$, $y_j \mapsto h_j$. Then $\phi \circ \psi$ is $R\{x, y\} \rightarrow S$, $x_i \mapsto \overline{g_i}$, $y_j \mapsto \overline{y_j}$. As the kernel of $\phi \circ \psi$ is $[x_i - g_i, f_j]$, it follows that the kernel of ϕ is σ -generated by $\psi(x_i - g_i), \psi(f_j)$. \square

By Hilbert’s basis theorem every finitely generated algebra over a Noetherian ring is of finite presentation. In difference algebra the situation is somewhat more delicate:

Non-example 3.2.5. Let k be a σ -field and $\mathfrak{a} := [y_1\sigma(y_1), y_1\sigma^2(y_1), \dots]$ the σ -ideal from Example 3.0.11. It follows from Lemma 3.2.4 that $S := k\{y_1\}/\mathfrak{a}$ is not finitely σ -presented over k .

The above non-example shows that even over a σ -field a finitely σ -generated σ -algebra need to be finitely σ -presented. However, the theory of characteristic sets yields the following:

Theorem 3.2.6. Let $R \subset S$ be an inclusion of σ -domains such that S is finitely σ -generated over R . Then there exists a non-zero element $s \in S$ such that $S\{\frac{1}{s}\}$ is finitely σ -presented over R .

Proof. We may write $S = R\{y\}/\mathfrak{p}$ for some σ -prime ideal \mathfrak{p} of $R\{y\} = R\{y_1, \dots, y_n\}$ with $\mathfrak{p} \cap R = 0$. Let F be a characteristic set of \mathfrak{p} and let $g \in R\{y\}$ denote the product of all initials of elements from F . It follows from Corollary 3.1.15 that the σ -ideal generated by \mathfrak{p} in $R\{y\}\{\frac{1}{g}\}$ is σ -generated by F . Now it is easy to see that the kernel of

$$R\{y_1, \dots, y_n, x\} \rightarrow (R\{y\}/\mathfrak{p})\{\frac{1}{g}\}, y_i \mapsto \overline{y_i}, x \mapsto \frac{1}{g}$$

is $[F, gx - 1]$. \square

Exercise 3.2.7. Define what it means for an algebra to be finitely presented. Show that if $R \subset S$ is an inclusion of integral domains such that S is finitely generated over R , then there exists a non-zero element $f \in S$ such that S_f is finitely presented over R . Can you also do this with $f \in R$, only assuming that R is an integral domain?

Exercise 3.2.8. Let k be a σ -field and R a finitely σ -generated k - σ -algebra. Let \mathfrak{p} be a σ -prime ideal of R . Show that the maximal ideal of $R_{\mathfrak{p}}$ is finitely σ -generated.

3.3 Ritt difference rings and the basis theorem

The role of Ritt difference rings in difference algebra is similar to the role of Noetherian rings in commutative algebra. However, there are some subtle differences...

Let R be a σ -ring and F a subset of R . We say that F has a *basis* if there exists a finite subset B of F such that the perfect σ -ideal generated by B equals the perfect σ -ideal generated by F , i.e., $\{B\} = \{F\}$. The set B is then called a basis of F . If $B^{(m)} = \{F\}$, we say that B is an *m-basis*.

☠ 3.3.1. *If B is a basis of a σ -ideal $\mathfrak{a} \subset R$ then B does not necessarily generate \mathfrak{a} as a σ -ideal. To avoid this misunderstanding, a finite subset B of \mathfrak{a} such that $\mathfrak{a} = [B]$ is sometimes called a strong basis of \mathfrak{a} .*

Because, even in a univariate σ -polynomial ring over a σ -field, not every σ -ideal has a strong basis (Example 3.0.11) the notion of a strong basis is not so useful.

In the following (and similar) statements the integer m is not fixed but allowed to depend on the set.

Lemma 3.3.2. *Let R be a σ -ring and assume that there exists a subset of R without basis/m-basis. Then there exists a σ -prime ideal \mathfrak{p} of R which is maximal among the subsets of R without basis/m-basis. That is, \mathfrak{p} has no basis/m-basis and every subset of R which properly contains \mathfrak{p} has a basis/m-basis.*

Proof. We only treat the case of a basis. The case of an m -basis is similar. The union of a chain of subsets of R without basis is a subset of R without basis. Therefore, by Zorn's lemma, there exists a maximal element \mathfrak{p} in the set of all subsets of R without basis. We claim that \mathfrak{p} is a σ -prime ideal. We start by showing that \mathfrak{p} is a σ -ideal: Suppose $\mathfrak{p} \subsetneq [\mathfrak{p}]$. Then $[\mathfrak{p}]$ has a basis, i.e., there exists a finite set $B \subset [\mathfrak{p}]$ such that $\{B\} = \{[\mathfrak{p}]\} = \{\mathfrak{p}\}$. There exists a finite subset F of \mathfrak{p} such that every element of B is an R -linear combination of elements of the form $\sigma^i(f)$ with $f \in F$ and $i \in \mathbb{N}$. Then $\{\mathfrak{p}\} = \{B\} \subset \{F\} \subset \{\mathfrak{p}\}$. So $\{\mathfrak{p}\} = \{F\}$ and F is a basis of \mathfrak{p} ; a contradiction. Therefore $\mathfrak{p} = [\mathfrak{p}]$ is a σ -ideal.

Next we will show that \mathfrak{p} is reflexive. Suppose $\mathfrak{p} \subsetneq \mathfrak{p}^*$. Then \mathfrak{p}^* has a basis. So there exists a finite set $B \subset \mathfrak{p}^*$ with $\{B\} = \{\mathfrak{p}^*\} = \{\mathfrak{p}\}$. By Lemma 1.2.15 there exists an $i \in \mathbb{N}$ such that $\sigma^i(B) \subset \mathfrak{p}$. Note that $B \subset \{\sigma^i(B)\}$. Then $\{\mathfrak{p}\} = \{B\} \subset \{\sigma^i(B)\} \subset \{\mathfrak{p}\}$. So $\{\mathfrak{p}\} = \{\sigma^i(B)\}$ and $\sigma^i(B)$ is a basis of \mathfrak{p} ; a contradiction.

Finally we show that \mathfrak{p} is prime. Suppose $fg \in \mathfrak{p}$ with $f, g \notin \mathfrak{p}$. Then both sets \mathfrak{p}, f and \mathfrak{p}, g have bases. So there exist finite sets $F, G \subset \mathfrak{p}$ such that $\{\mathfrak{p}, f\} = \{F, f\}$ and $\{\mathfrak{p}, g\} = \{G, g\}$. Then

$$\{\mathfrak{p}\} \subset \{\mathfrak{p}, f\} \cap \{\mathfrak{p}, g\} = \{FG, Fg, fG, fg\} \subset \{\mathfrak{p}\}$$

by Proposition 1.2.20 and FG, Fg, fG, fg is a basis for \mathfrak{p} ; a contradiction. □

Exercise 3.3.3. *Do the m -basis case of Lemma 3.3.2.*

Recall that a topological space is called Noetherian if every descending chain of closed subsets is finite.

Proposition 3.3.4. *Let R be a σ -ring. The following statements are equivalent:*

- (i) *Every subset of R has a basis.*
- (ii) *Every perfect σ -ideal of R has a basis.*

- (iii) Every σ -prime ideal of R has a basis.
- (iv) Every ascending chain of perfect σ -ideals of R is finite.
- (v) $\text{Spec}^\sigma(R)$ is a Noetherian topological space.

Proof. The implications (i) \Rightarrow (ii) \Rightarrow (iii) are trivial and (iii) \Rightarrow (i) follows from Lemma 3.3.2. We show (ii) \Rightarrow (iv): Let $\mathfrak{a}_1 \subset \mathfrak{a}_2 \subset \dots$ be an ascending chain of perfect σ -ideals. The union $\mathfrak{a} := \cup_{i \geq 1} \mathfrak{a}_i$ is a perfect σ -ideal of R . So, by assumption, there exists a finite set $B \subset \mathfrak{a}$ such that $\{B\} = \mathfrak{a}$. But $B \subset \mathfrak{a}_{i_0}$ for some $i_0 \geq 1$ and so, $\mathfrak{a} = \mathfrak{a}_{i_0} = \mathfrak{a}_{i_0+1} = \dots$.

Next we will show (iv) \Rightarrow (i): Let F be a subset of R . Choose $f_1 \in F$. If $\{f_1\} = \{F\}$ we are done. Otherwise, there exists $f_2 \in F \setminus \{f_1\}$. Then $\{f_1\} \subsetneq \{f_1, f_2\}$. Continuing in this way we obtain a strictly ascending chain of perfect σ -ideals. By assumption, this process must terminate, i.e., $\{F\} = \{f_1, \dots, f_n\}$ for some n .

Finally, the equivalence of (iv) and (v) is clear from Proposition 1.2.35 (ii). \square

Exercise 3.3.5. *Is there an analog of Proposition 3.3.4 for radical well-mixed σ -ideals in place of perfect σ -ideals?*

Definition 3.3.6. *A difference ring satisfying the equivalent definitions of Proposition 3.3.4 is called a Ritt difference ring.*

Example 3.3.7. Obviously a σ -field is Ritt. More generally, any σ -Ring which is Noetherian is Ritt.

The following theorem can be seen as a σ -analog of Hilbert's basis theorem. Needless to say that this is a very fundamental result.

Theorem 3.3.8 (Ritt-Raudenbush basis theorem). *Let R be a Ritt difference ring. Then the σ -polynomial ring $R\{y_1, \dots, y_n\}$ over R is also a Ritt difference ring.*

Proof. By induction it suffices to treat the case $y = y_1$. Suppose that $R\{y\} = R\{y_1\}$ is not Ritt. By Lemma 3.3.2 there exists a σ -prime ideal \mathfrak{p} of $R\{y\}$ with no basis such that \mathfrak{p} is maximal in the set of all subsets of $R\{y\}$ with no basis. The ideal $\mathfrak{p}' := R \cap \mathfrak{p}$ is a perfect σ -ideal of R . Because R is Ritt, there exists a finite set $F \subset R$ such that $\mathfrak{p}' = \{F\}$. Let M denote the set of all non-zero σ -polynomials $f \in \mathfrak{p}$ such that no non-zero coefficient of f lies in \mathfrak{p}' .

Suppose $M = \emptyset$. Let $f \in \mathfrak{p}$. If we subtract from f all the terms which have their coefficient in \mathfrak{p}' we obtain a polynomial $f' \in \mathfrak{p}$. As $M = \emptyset$ we must have $f' = 0$. But then \mathfrak{p} is generated as an ideal by \mathfrak{p}' and so F is a basis for \mathfrak{p} ; a contradiction. Therefore $M \neq \emptyset$.

We are going to construct an autoreduced subset G of M such that every polynomial in \mathfrak{p} which is reduced with respect to G does not lie in M , i.e., lies in $R\{y\}\mathfrak{p}'$. Note that because $y = y_1$ there is only one possible ranking, namely the one given by $y < \sigma(y) < \dots$. We proceed as outlined after Definition 3.1.8: Choose $g_1 \in M$ of minimal rank. If the set of all elements in M which are reduced with respect g_1 is non-empty, choose g_2 from this set of minimal rank. We go on like this: If the set of all elements in M which are reduced with respect to g_1, g_2 is non-empty, choose g_3 from this set of minimal rank.

Because an autoreduced set is finite this process terminates. In the end we have obtained an autoreduced subset $G = g_1, \dots, g_n$ of M such that every σ -polynomial $f \in \mathfrak{p}$ which is reduced with respect to G lies in $R\{y\}\mathfrak{p}'$.

Let $I = I_{g_1} \dots I_{g_n}$ denote the product of the initials of elements from G . It is clear that I_{g_i} ($i = 1, \dots, n$) is reduced with respect to G . Because g_i has no coefficient in \mathfrak{p}' ,

the same holds for I_{g_i} . Therefore $I_{g_i} \notin \mathfrak{p}$. Because \mathfrak{p} is prime $I \notin \mathfrak{p}$. By choice of \mathfrak{p} , the set \mathfrak{p}, I has a basis. So there exists a finite subset H of \mathfrak{p} such that $\{H, I\} = \{\mathfrak{p}, I\}$.

Let $f \in \mathfrak{p}$. By Proposition 3.1.11 there exists $h \in I_G$ such that $hf - f_0 \in [G]$ and f_0 is reduced with respect to G . Because $f_0 \in \mathfrak{p}$ is reduced with respect to G we have $f_0 \in R\{y\}\mathfrak{p}'$. Consequently $hf \in \{F, G\}$. We can multiply hf with appropriate factors of the form $\sigma^i(I_{g_j})$ to obtain $\sigma^{i_1}(I) \cdots \sigma^{i_m}(I)f \in \{F, G\}$. Then also $If\sigma^{i_1}(If) \cdots \sigma^{i_m}(If) \in \{F, G\}$ and therefore $If \in \{F, G\}$. We have shown that $I\mathfrak{p} \subset \{F, G\}$. This yields

$$\mathfrak{p} = \mathfrak{p} \cap \{\mathfrak{p}, I\} = \mathfrak{p} \cap \{H, I\} = \{\mathfrak{p}H, \mathfrak{p}I\} \subset \{H, F, G\} \subset \mathfrak{p}.$$

So H, F, G is a basis for \mathfrak{p} ; a contradiction. \square

Proposition 3.3.9. *Let R be a σ -ring such that every perfect σ -ideal of R has an m -basis, then every perfect σ -ideal of the σ -polynomial ring $R\{y_1, \dots, y_n\}$ has an m -basis.*

Proof. Analogous to Theorem 3.3.8. \square

Exercise 3.3.10. *Prove Proposition 3.3.9 above.*

Corollary 3.3.11. *Let k be a σ -field and R a finitely σ -generated k - σ -algebra. Then R is Ritt.*

Proof. Since we can write R as a quotient of a σ -polynomial ring over k , the claim follows from Theorem 3.3.8 and Proposition 1.2.8. \square

Exercise 3.3.12. * *Let k be a σ -field. Is every ascending chain of radical well-mixed σ -ideals in $k\{y_1, \dots, y_n\}$ finite?*

3.4 Application to difference varieties

The geometric meaning of the basis theorem is the following:

Theorem 3.4.1. *Let X be a σ -variety. Then the topological space of X is Noetherian.*

Proof. By Corollary 3.3.11 the σ -coordinate ring $k\{X\}$ is Ritt, so $\text{Spec}^\sigma(k\{X\})$ is Noetherian by Proposition 3.3.4. \square

Lemma 3.4.2. *In a Noetherian topological space X every non-empty closed subset Y of X can be expressed as a finite union*

$$Y = Y_1 \cup \cdots \cup Y_n$$

of irreducible closed subsets $Y_i \subset X$. If we require that $Y_i \not\subseteq Y_j$ for $i \neq j$, then the Y_i are uniquely determined. They are called the irreducible components of Y .

Proof. Let M denote the set of all non-empty closed subsets of X which can not be written as a finite union of irreducible closed sets. Suppose M is non-empty. Since X is Noetherian there exists a minimal (with respect to inclusion) element Y of M . Since Y is not irreducible, $Y = Y' \cup Y''$ for some closed subsets Y', Y'' of X with $Y \supsetneq Y', Y''$. It follows from the minimality of Y that $Y', Y'' \notin M$. But then also $Y \notin M$; a contradiction.

Consequently every closed subset Y of X is of the form $Y = Y_1 \cup \cdots \cup Y_n$ with $Y_i \subset X$ closed and irreducible. By deleting some of the Y_i if necessary, we can assume that $Y_i \not\subseteq Y_j$ for $i \neq j$.

It remains to see that the Y_i are unique: Assume that $Y = Y'_1 \cup \cdots \cup Y'_{n'}$ is another irredundant (i.e. $Y'_i \not\subseteq Y'_j$ for $i \neq j$) representation of Y . For $1 \leq i \leq n'$ we have

$Y'_i \subset Y = Y_1 \cup \dots \cup Y_n$. Since Y'_i is irreducible $Y'_i \subset Y_j$ for some $j = j(i) \in \{1, \dots, n\}$. Analogously we find that $Y_j \subset Y'_i$ for some $\tilde{i} = \tilde{i}(j) \in \{1, \dots, n'\}$. Therefore $Y'_i \subset Y_j \subset Y'_{\tilde{i}}$. This yields $i = \tilde{i}$ and $Y'_i = Y_j$. \square

Corollary 3.4.3. *Let R be a Ritt σ -ring. Then every perfect σ -ideal \mathfrak{a} of R can be expressed as a finite intersection*

$$\mathfrak{a} = \mathfrak{p}_1 \cap \dots \cap \mathfrak{p}_n$$

of σ -prime ideals $\mathfrak{p}_i \subset R$. If we require that $\mathfrak{p}_i \not\subseteq \mathfrak{p}_j$ for $i \neq j$, then the \mathfrak{p}_i are uniquely determined. They are called the prime components of \mathfrak{a} .

Proof. This is clear from Lemma 3.4.2 and Proposition 1.2.35. \square

Let k be a σ -field and X a k - σ -variety. By Lemma 3.4.2 the topological space of X is the finite union of its irreducible components. The σ -subvarieties of X corresponding to the irreducible components of the topological space of X are called the *irreducible components of X* .

Definition 3.4.4. *Let k be a σ -field, X a k - σ -variety and $x \in X$. We say that X is finitely σ -presented (over k) at x if there exists $f \in k\{X\}$ such that $x \in D^\sigma(f)$ and $k\{X\}\{\frac{1}{f}\}$ is finitely σ -presented over k .*

Theorem 3.4.5. *Let X be a σ -variety. The set of $x \in X$ such that X is finitely σ -presented at x is open and dense in X .*

Proof. Let Y denote the set of all $x \in X$ such that X is finitely σ -presented at x . Since $D^\sigma(f)$ is open in X for $f \in k\{X\}$ it is clear that Y is open.

It remains to show that Y is dense in X . It suffices to show that Y contains the generic points of the irreducible components of X . If we assume that X is irreducible (i.e. $k\{X\}$ is a σ -domain) the claim immediately follows from Theorem 3.2.6. In general, the zero ideal of $k\{X\}$ is the finite intersection of its prime components:

$$(0) = \mathfrak{p}_1 \cap \dots \cap \mathfrak{p}_n.$$

We confine ourselves to show that $\mathfrak{p}_1 \in Y$. There exists $f \in (\mathfrak{p}_2 \cap \dots \cap \mathfrak{p}_n) \setminus \mathfrak{p}_1$. The kernel of the canonical map $k\{X\} \rightarrow k\{X\}\{\frac{1}{f}\}$ is \mathfrak{p}_1 . So $k\{X\}\{\frac{1}{f}\}$ can be identified with $(k\{X\}/\mathfrak{p}_1)\{\frac{1}{f}\}$. This shows that $k\{X\}\{\frac{1}{f}\}$ is a σ -domain. By Theorem 3.2.6 there exists a non-zero element g of $k\{X\}\{\frac{1}{f}\}$ such that $k\{X\}\{\frac{1}{f}\}\{\frac{1}{g}\}$ is finitely σ -presented over k . We can assume that $g \in k\{X\}$. Then $g \notin \mathfrak{p}_1$, $\mathfrak{p}_1 \in D^\sigma(fg)$ and $k\{X\}\{\frac{1}{fg}\} \simeq k\{X\}\{\frac{1}{f}\}\{\frac{1}{g}\}$ is finitely σ -presented over k . \square

Exercise 3.4.6. *Let X be the σ -subvariety of \mathbb{A}_k^2 defined by $y_1 y_2$. Determine all $x \in X$ such that X is finitely σ -presented at x .*

Exercise 3.4.7. *Let k be a σ -field and X a k - σ -variety. Assume that X is finitely σ -presented at every $x \in X$. Is $k\{X\}$ finitely σ -presented over k ?*

Chapter 4

Extensions of difference fields

In this chapter we begin our systematic study of σ -field extensions. We first study difference algebraic extensions and difference transcendence bases. Here the analogy with the situation for classical field extensions is very strong. Then we introduce the limit degree, an important numerical invariant of a σ -field extension. The more important results are the fundamental theorem on finitely σ -generated σ -field extensions (stating that an intermediate σ -field extensions of a finitely σ -generated σ -field extension is finitely σ -generated), Babbitt's decomposition and the compatibility theorem.

4.1 Difference algebraic extensions

Definition 4.1.1. Let $L|K$ be an extension of σ -fields and $a \in L$. We say that a is σ -algebraic over K if there exists a non-zero σ -polynomial which annihilating a . Otherwise we say that a is σ -transcendental over K . We say that L is σ -algebraic over K , if every element from L is σ -algebraic over K .

In other words, if $K\{y\}$ is the univariate σ -polynomial ring over K , then a is σ -algebraic over K if and only if the kernel of

$$K\{y\} \rightarrow L, y \mapsto a$$

is non-zero.

Definition 4.1.2. Let $L|K$ be an extension of σ -fields. The transcendence degree of $L|K$ is also called the order of L over K . So

$$\text{ord}(L|K) = \text{trdeg}(L|K).$$

This definition is justified by the following lemma. This lemma also shows that the role of the order in difference algebra is similar to the role of the degree in usual algebra.

Lemma 4.1.3. Let K be a σ -field and a an element in a σ -field extension of K . If a is σ -algebraic over K then $\text{ord}(K\langle a \rangle|K)$ is the minimal integer n such that there exists a non-zero univariate σ -polynomial of order n with coefficients in K which annihilates a .

Proof. It follows from the minimality of n that $a, \sigma(a), \dots, \sigma^{n-1}(a)$ are algebraically independent over K . Thus $\text{ord}(K\langle a \rangle|K) \geq n$. On the other hand, let $f \in K\{y\} \setminus \{0\}$ be of order n such that $f(a) = 0$. Because f has order n , the element $\sigma^n(a)$ is algebraic over $K(a, \sigma(a), \dots, \sigma^{n-1}(a))$. Similarly, since $\sigma^i(f)(a) = 0$ we see that $\sigma^{n+i}(a)$ is algebraic over $K(a, \sigma(a), \dots, \sigma^{n+i-1}(a))$ for $i = 1, 2, \dots$. This shows that $K\langle a \rangle$ is an algebraic extension of $K(a, \sigma(a), \dots, \sigma^{n-1}(a))$. Therefore $n = \text{ord}(K\langle a \rangle|K)$. \square

Lemma 4.1.4. *Let $L|K$ be an extension of σ -fields and $a \in L$. The following statements are equivalent.*

- (i) *The element a is σ -algebraic over K .*
- (ii) *The σ -field extension $K\langle a \rangle$ has finite order over K .*
- (iii) *The element a is contained in an intermediate σ -field $K \subset M \subset L$ which has finite order over K .*

Proof. The implication (i) \Rightarrow (ii) is clear from Lemma 4.1.3 and the implication (ii) \Rightarrow (iii) is trivial. We show that (iii) implies (i): Since M has finite order over K , the elements $a, \sigma(a), \dots$ must be algebraically dependent over K . This dependence relation is given by a non-zero σ -polynomial. So a is σ -algebraic over K . \square

Note that every σ -field extension of finite order is σ -algebraic, but the converse does not hold.

Lemma 4.1.5. *Let $L|K$ be an extension of σ -fields and a, b elements from L which are σ -algebraic over K . Then $a + b, ab, \sigma(a)$ and a^{-1} (if $a \neq 0$) are σ -algebraic over K . In particular, a σ -field extension generated by σ -algebraic elements is σ -algebraic and the set of all elements in L which are σ -algebraic over K is a σ -algebraic σ -field extension of K . Moreover, a σ -field extension which is generated by finitely many σ -algebraic elements has finite order.*

Proof. Because $K\langle a \rangle$ and $K\langle b \rangle$ have finite transcendence degree over K also the composite field $K\langle a \rangle K\langle b \rangle \subset L$ has finite transcendence degree over K . But this is a σ -field containing $a + b, ab, \sigma(a)$ and a^{-1} . So the claim follows from Lemma 4.1.3. \square

Lemma 4.1.6. *Let $K \subset L \subset M$ be σ -fields. Then M is σ -algebraic over K if and only if M is σ -algebraic over L and L is σ -algebraic over K .*

Proof. One direction is trivial: If M is σ -algebraic over K then clearly M is σ -algebraic over L and L is σ -algebraic over K .

Assume that M is σ -algebraic over L and L is σ -algebraic over K . Choose $a \in M$. We have to show that a is σ -algebraic over K . By assumption there exists a non-zero σ -polynomial with coefficients in L which annihilates a . Because L is σ -algebraic over K , the σ -field L' generated by all the coefficients of f over K has finite order over K . But clearly $L'\langle a \rangle$ has finite order over L' . So $L'\langle a \rangle$ has finite order over K (by elementary properties of the transcendence degree). It follows from Lemma 4.1.3 that a is σ -algebraic over K . \square

4.1.1 Application to difference varieties

The following theorem can be seen as a σ -analog of Hilbert's (weak) Nullstellensatz. The proof is surprisingly simple, especially if you take into account that usually the σ -analogs are much harder to prove than their original commutative algebra counterparts.

Theorem 4.1.7. *Let k be a σ -field and $F \subset k\{y_1, \dots, y_n\}$ a system of algebraic difference equations over k . If F has a solution in a σ -field extension of k , then F has a solution in a σ -algebraic σ -field extension of k .*

Proof. By Corollary 3.3.11 there exists a finite subset B of F such that $\{B\} = \{F\}$. This means that F and B have the same solutions in every σ -field extension of k . Therefore, we can assume without loss of generality that F is finite.

We first reduce to the case that all the elements from F have order lesser than or equal to one: Let $m \in \mathbb{N}$ be an integer such that all elements in F have order not greater than m . Let

$$z = (z_j^i)_{0 \leq i \leq m, 1 \leq j \leq n}$$

be a new set of σ -variables over k . For $f \in F$ let \tilde{f} denote the polynomial in z which is obtained from f by replacing $\sigma^i(y_j)$ with z_j^i . Let G denote the system of algebraic difference equations in z over k , which consists of all \tilde{f} ($f \in F$) and all equations

$$\sigma(z_j^i) = z_j^{i+1} \text{ for } i = 0, \dots, m-1 \text{ and } j = 1, \dots, n.$$

Then G has a solution in a (σ -algebraic) σ -field extension of k if and only if F has a solution in a (σ -algebraic) σ -field extension of k . Note that the elements of G have order lesser than or equal to one.

So we can assume without loss of generality that all the elements from F have order lesser than or equal to one. Because F has a solution in a σ -field extension of k , there exists a σ -prime ideal \mathfrak{p} of $k\{y_1, \dots, y_n\}$ which contains F . (Cf. Theorem 2.2.1.) Set

$$\mathfrak{p}[1] := \mathfrak{p} \cap k[y_1, \dots, y_n, \sigma(y_1), \dots, \sigma(y_n)] \text{ and } \mathfrak{p}[0] := \mathfrak{p} \cap k[y_1, \dots, y_n].$$

Let $L = k(\mathfrak{p}[1])$ denote the residue field of $\mathfrak{p}[1]$. Denote the image of $y, \sigma(y)$ in L by a, b , respectively. So $L = k(a, b)$. Because $\sigma^{-1}(\mathfrak{p}[1]) = \mathfrak{p}[0]$ we have a map $\sigma: k(a) \rightarrow L$ which extends $\sigma: k \rightarrow k$ and satisfies $\sigma(a) = b$. Let \bar{L} denote an algebraic closure of L . We would like to extend $\sigma: k(a) \rightarrow \bar{L}$ to an endomorphism of L . To be able to do this we only need to know that $\text{trdeg}(L|\sigma(k)(b)) \geq \text{trdeg}(L|k(a))$. But

$$\begin{aligned} \text{trdeg}(L|\sigma(k)(b)) &= \text{trdeg}(L|\sigma(k)) - \text{trdeg}(\sigma(k)(b)|\sigma(k)) \geq \\ &\geq \text{trdeg}(L|k) - \text{trdeg}(k(a)|k) = \text{trdeg}(L|k(a)). \end{aligned}$$

Therefore we can extend σ to \bar{L} . So \bar{L} is a σ -field extension of k . Note that \bar{L} is σ -algebraic over k because \bar{L} has finite transcendence degree over k . Because $F \subset \mathfrak{p}[1]$ it is clear that a is a solution of F in \bar{L} . \square

Corollary 4.1.8. *Let k be a σ -field and R a finitely σ -generated k - σ -algebra. Let \mathfrak{p} be a maximal element in the set of all σ -prime ideals of R ordered by inclusion (i.e., a closed point of $\text{Spec}^\sigma(R)$). Then $k(\mathfrak{p})$ is σ -algebraic over k .*

Proof. We may write $R \simeq k\{y\}/\mathfrak{a}$ for some σ -ideal \mathfrak{a} in the σ -polynomial ring $k\{y\} = k\{y_1, \dots, y_n\}$. Then \mathfrak{p} corresponds to a σ -prime ideal \mathfrak{p}' of $k\{y\}$ above \mathfrak{a} . By Theorem 4.1.7 there exists a solution $a \in K^n$ of \mathfrak{p}' in some σ -algebraic σ -field extension K of k . By the maximality of \mathfrak{p} , the corresponding map $k\{y\}/\mathfrak{p}' \rightarrow K, \bar{y} \mapsto a$ is injective. Thus $k(\mathfrak{p}')$ embeds into K . Consequently $k(\mathfrak{p}) \simeq k(\mathfrak{p}')$ is σ -algebraic over k . \square

The following corollary states that difference varieties are determined by their points in σ -algebraic extensions of the base σ -field. This is analogous to the fact that varieties are determined by their points in algebraic field extensions. (In a scheme of finite type over a field the closed points are dense.)

Corollary 4.1.9. *Let k be a σ -field and X, Y difference varieties over k . If $X(K) = Y(K)$ for every σ -algebraic σ -field extension K of k , then $X = Y$.*

Proof. Suppose $X \neq Y$. Say, $X \not\subseteq Y$. Assume $X = \mathbb{V}(F)$ and $Y = \mathbb{V}(G)$ with $F, G \subset k\{y_1, \dots, y_n\}$. (Note that X and Y must lie in the same \mathbb{A}_k^n by Theorem 4.1.7.) Because $X \not\subseteq Y$ there exists a σ -field extension K of k such that $X(K) \not\subseteq Y(K)$. This means that for some $a \in X(K)$ and some $g \in G$ we have $g(a) \neq 0$.

Let \tilde{F} denote the system of algebraic difference equations over k in the σ -variables y_1, \dots, y_n, y_{n+1} given by F and $y_{n+1}g - 1$. Then $(a, 1/g(a))$ is a solution of \tilde{F} in K . By Theorem 4.1.7 there exists a σ -algebraic σ -field extension L of k and a solution $b = (b_1, \dots, b_{n+1}) \in L^{n+1}$ of \tilde{F} . But then $(b_1, \dots, b_n) \in X(L)$ and $(b_1, \dots, b_n) \notin Y(L)$ because $g(b_1, \dots, b_n) \neq 0$. This contradicts $X(L) = Y(L)$. \square

Exercise 4.1.10. Let k be a σ -field and R a finitely σ -generated k - σ -algebra. Is the set of closed points dense in $\text{Spec}^\sigma(R)$? (Hint: What are the fixed points of the map $x \mapsto x^2$?) Is the set of locally closed points dense in $\text{Spec}^\sigma(R)$? (Recall that a subset of a topological space is called locally closed, if it is open in its closure.)

4.2 Difference transcendence bases

Definition 4.2.1. Let R be a σ -ring. Elements f_1, \dots, f_n in an R - σ -algebra S are called σ -algebraically independent over R if they satisfy no σ -algebraic relation over R . I.e.,

$$R\{y_1, \dots, y_n\} \rightarrow S, \quad y_i \mapsto f_i, \quad i = 1, \dots, n$$

is injective. Otherwise, we say that f_1, \dots, f_n are σ -algebraically dependent over R . A subset of S is called σ -algebraically independent over R if all its finite subsets are σ -algebraically independent over R .

Definition 4.2.2. Let $L|K$ be an extension of σ -fields and $A \subset L$. An element $b \in L$ is σ -algebraically dependent on A (over K) if b is σ -algebraic over $K\langle A \rangle$. A subset B of L is called σ -algebraically dependent on A (over K) if every element of B is σ -algebraically dependent on A .

If it is clear that we are working over the base σ -field K , then we usually omit the “over K ”.

Lemma 4.2.3. Let $L|K$ be an extension of σ -fields, $A \subset L$ and $b \in L$. Then b is σ -algebraically dependent on A if and only if there exists a σ -polynomial $f = f(y_1, \dots, y_n, z)$ with coefficients in K and $a_1, \dots, a_n \in A$ such that $f(a_1, \dots, a_n, z) \neq 0$ but $f(a_1, \dots, a_n, b) = 0$.

Proof. Assume that b is σ -algebraic over $K\langle A \rangle$. This means that there exists a non-zero univariate σ -polynomial $g \in K\langle A \rangle\{z\}$ such that $g(b) = 0$. After multiplying g with an appropriate element from $K\{A\}$ we can assume that g has coefficients in $K\{A\}$. Now the claim follows from the explicit description of $K\{A\}$ given in Subsection 1.1.5. The converse is obvious. \square

Lemma 4.2.4. Let $L|K$ be an extension of σ -fields and A a subset of L which is σ -algebraically independent over K . If b is an element from L such that A, b is σ -algebraically dependent over K , then b is σ -algebraically dependent on A over K .

Proof. Because A, b is σ -algebraically dependent over K there exists a non-zero σ -polynomial $f = f(y_1, \dots, y_n, z)$ with coefficients in K and elements $a_1, \dots, a_n \in A$ such that $f(a_1, \dots, a_n, b) = 0$. Because a_1, \dots, a_n are σ -algebraically independent, $f(a_1, \dots, a_n, z) \neq 0$. This shows that b satisfies a non-zero σ -polynomial over $K\langle A \rangle$. \square

Lemma 4.2.5 (Transitivity of σ -algebraic dependence). *Let $L|K$ be an extension of σ -fields and A, B, C subsets of L such that A is σ -algebraically dependent on B and B is σ -algebraically dependent on C . Then A is σ -algebraically dependent on C .*

Proof. The assumptions imply that $K\langle B, C \rangle$ is σ -algebraic over $K\langle C \rangle$ and that $K\langle A, B, C \rangle$ is σ -algebraic over $K\langle B, C \rangle$. It follows from Lemma 4.1.6 that $K\langle A, B, C \rangle$ is σ -algebraic over $K\langle C \rangle$. In particular every element from A is σ -algebraic over $K\langle C \rangle$. \square

Lemma 4.2.6 (The exchange property). *Let a_1, \dots, a_n, b be elements in a σ -field extension of a σ -field K . If b is σ -algebraically dependent on a_1, \dots, a_n but not on a_1, \dots, a_{n-1} , then a_n is σ -algebraically dependent on a_1, \dots, a_{n-1}, b .*

Proof. Because b is σ -algebraically dependent on a_1, \dots, a_n , there exists a σ -polynomial $f = f(y_1, \dots, y_n, z)$ such that $f(a_1, \dots, a_n, z) \neq 0$ but $f(a_1, \dots, a_n, b) = 0$ (Lemma 4.2.3). Let $g_1, \dots, g_m \in K\{y_1, \dots, y_{n-1}, z\}$ denote the non-zero coefficients of f when considered as a univariate σ -polynomial in y_n . Then $g_i(a_1, \dots, a_{n-1}, z) \neq 0$ for some i because otherwise $f(a_1, \dots, a_n, z)$ would be zero. Because b is not σ -algebraically dependent on a_1, \dots, a_{n-1} , it follows that $g_i(a_1, \dots, a_{n-1}, b) \neq 0$. But this shows that $f(a_1, \dots, a_{n-1}, y_n, b) \neq 0$ and so a_n is σ -algebraically dependent on a_1, \dots, a_{n-1}, b . \square

Proposition 4.2.7. *Let $L|K$ be an extension of σ -fields and $A = a_1, \dots, a_n$, $B = b_1, \dots, b_m$ subsets of L . Assume that A is σ -algebraically independent over K but σ -algebraically dependent on B . Then $n \leq m$.*

Proof. Let r be the number of elements that A and B have in common. If $r = n$ we are done. So assume $r < n$ and write $B = a_1, \dots, a_r, b_{r+1}, \dots, b_m$. Since a_{r+1} is σ -algebraically dependent on $a_1, \dots, a_r, b_{r+1}, \dots, b_m$ but not on a_1, \dots, a_r , there will be a b_j , $r + 1 \leq j \leq m$ such that a_{r+1} is σ -algebraically dependent on $a_1, \dots, a_r, b_{r+1}, \dots, b_j$, but not on $a_1, \dots, a_r, b_{r+1}, \dots, b_{j-1}$. The exchange lemma (Lemma 4.2.6) then shows that b_j is σ -algebraically dependent on

$$B_1 := (B \setminus \{b_j\}) \cup \{a_{r+1}\}.$$

Therefore B is σ -algebraically dependent on B_1 . Because A is σ -algebraically dependent on B it follows from Lemma 4.2.5 that A is σ -algebraically dependent on B_1 . Note that B_1 has the same size as B , but one more element in common with A . Continuing in this way we will eventually arrive at $r = n$, i.e., $A \subset B$. \square

Definition 4.2.8. *Let $L|K$ be an extension of σ -fields. A subset A of L is called a σ -transcendence basis of $L|K$ if A is σ -algebraically independent over K and L is σ -algebraic over $K\langle A \rangle$.*

By the size of a set, we mean its cardinality if the set is finite, and infinity otherwise. In other words, we are not interested in cardinals.

Theorem 4.2.9. *Let $L|K$ be an extensions of σ -fields. Then every σ -generating set of $L|K$ contains a σ -transcendence basis of $L|K$. In particular, there exists a σ -transcendence basis of $L|K$. Moreover, any two σ -transcendence bases of $L|K$ are of the same size.*

Proof. Let M be a σ -generating set of $L|K$. Let N be the set of all σ -algebraically independent subsets of M . Clearly the union of any chain of elements in N is again in N . Therefore, by Zorn's lemma, there exists a maximal σ -algebraically independent subset A of M . We will show that A is a σ -transcendence basis of $L|K$. By Lemma 4.2.4 every element of M is σ -algebraically dependent on A . I.e., all the elements of M are σ -algebraic

over $K\langle A \rangle$. Because L is σ -generated by M , this implies that L is σ -algebraic over $K\langle A \rangle$ by Lemma 4.1.5. So A is a σ -transcendence basis of $L|K$.

Now let A and B be transcendence bases of $L|K$. By symmetry it suffices to show that the size of A is greater than or equal to the size of B . If A is infinite this is automatically the case. So we may assume that A is finite. Let B' be a finite subset of B . Because A is a transcendence basis, every element of L is σ -algebraically dependent on A , in particular B' is σ -algebraically dependent on A . By Proposition 4.2.7 the size of A is greater than or equal to the size of B' . It follows that the size of A is greater than or equal to the size of B . \square

In the course of the above proof we have established the following corollary.

Corollary 4.2.10. *Let $L|K$ be an extensions of σ -fields with a σ -generating set M . If A is a maximal σ -algebraically independent subset of M , then A is a σ -transcendence basis of $L|K$.* \square

Thanks to Theorem 4.2.9 we can make the following definition.

Definition 4.2.11. *Let $L|K$ be an extensions of σ -fields. The size of a σ -transcendence basis of $L|K$ is called the σ -transcendence degree of $L|K$. It is denoted by*

$$\sigma\text{-trdeg}(L|K).$$

Corollary 4.2.12. *Let $L|K$ be an extension of σ -fields which can be σ -generated by n elements. Then $\sigma\text{-trdeg}(L|K) \leq n$. In particular, the σ -transcendence degree of a finitely σ -generated σ -field extension is finite.*

Proof. This is clear from Corollary 4.2.10. \square

Corollary 4.2.13. *Let $L|K$ be an extension of σ -fields. If L contains n σ -independent elements, then $n \leq \sigma\text{-trdeg}(L|K)$. In fact,*

$$\sigma\text{-trdeg}(L|K) = \sup\{n \in \mathbb{N} \mid \exists a_1, \dots, a_n \in L \text{ } \sigma\text{-algebraically independent over } K\}.$$

Proof. Let $A = a_1, \dots, a_n$ be a set of n elements from L which are σ -algebraically independent over K . We can enlarge A to a σ -generating set B of $L|K$. Then A is contained in a maximal σ -algebraically independent subset A' of B . By Corollary 4.2.10 the set A' is a σ -transcendence basis of $L|K$. Therefore $n \leq \sigma\text{-trdeg}(L|K)$. This shows that the supremum is lesser than or equal to $\sigma\text{-trdeg}(L|K)$. Because a σ -transcendence basis is σ -algebraically independent, the reverse estimate is clear. \square

In the following theorem the usual rules for calculating with the symbol ∞ have to be applied.

Theorem 4.2.14. *Let $K \subset L \subset M$ be σ -fields. Then*

$$\sigma\text{-trdeg}(M|K) = \sigma\text{-trdeg}(M|L) + \sigma\text{-trdeg}(L|K).$$

Proof. Let A be a σ -transcendence basis of $L|K$ and B a σ -transcendence basis of $M|L$. It suffices to show that A, B is a σ -transcendence basis of $M|K$. Because B is σ -algebraically independent over $K\langle A \rangle$ it is clear that A, B is σ -algebraically independent over K .

It remains to show that M is σ -algebraic over $K\langle A, B \rangle$. Every element from L is σ -algebraic over $K\langle A, B \rangle$. Therefore $K\langle L, A, B \rangle = L\langle B \rangle$ is σ -algebraic over $K\langle A, B \rangle$. Because M is σ -algebraic over $L\langle B \rangle$ it follows from Lemma 4.1.6 that M is σ -algebraic over $K\langle A, B \rangle$. \square

4.2.1 Application to difference varieties

Let k be a σ -field and X an irreducible k - σ -variety. Then $k\{X\}$ is a σ -domain and we can consider the quotient field

$$k\langle X \rangle := \text{Quot}(k\{X\}).$$

It is naturally a σ -field extension of k , called the σ -function field of X . Clearly, $k\langle X \rangle$ is finitely σ -generated over k . Indeed, the coordinate functions are a canonical σ -generating set.

Definition 4.2.15. *Let k be a σ -field and X an irreducible k - σ -variety. The σ -dimension of X is the σ -transcendence degree of the σ -function field of X over k . That is,*

$$\sigma\text{-dim}(X) = \sigma\text{-trdeg}(k\langle X \rangle|k).$$

For an arbitrary k - σ -variety X , with irreducible components X_1, \dots, X_n , we set

$$\sigma\text{-dim}(X) = \max_i \sigma\text{-dim}(X_i).$$

Example 4.2.16. The σ -dimension of \mathbb{A}_k^n is n .

Exercise 4.2.17. *Let k be a σ -field. Show that \mathbb{A}_k^n is not isomorphic to \mathbb{A}_k^m unless $n = m$.*

Exercise 4.2.18. *Let k be an algebraically closed σ -field and \mathcal{X} an affine, irreducible variety over k with coordinate ring $k[\mathcal{X}]$. Let X be the k - σ -variety with σ -coordinate ring $k\{X\} = [\sigma]_k k[\mathcal{X}]$. (Cf. Exercise 1.1.30.) Show that $\sigma\text{-dim}(X) = \dim(\mathcal{X})$.*

Theorem 4.2.19. *Let X be a k - σ -variety and $Y \subset X$ a σ -subvariety. Then $\sigma\text{-dim}(Y) \leq \sigma\text{-dim}(X)$.*

Proof. First assume that X and Y are irreducible. Assume that $X \subset \mathbb{A}_k^n$ with coordinate functions y_1, \dots, y_n . By Theorem 4.2.9 there exists a subset $I \subset \{1, \dots, n\}$ such that the images of $(y_i)_{i \in I}$ in $k\{Y\}$ are a σ -transcendence basis of $k\langle Y \rangle$ over k .

The inclusion $Y \subset X$ corresponds to a surjection $k\{X\} \rightarrow k\{Y\}$ on σ -coordinate rings. The images of $(y_i)_{i \in I}$ in $k\{X\}$ can not be σ -algebraically dependent over k , because otherwise, the images of $(y_i)_{i \in I}$ in $k\{Y\}$ were also σ -algebraically dependent over k . It follows that $\sigma\text{-trdeg}(k\langle X \rangle|k) \geq |I|$ by Corollary 4.2.13. So $\sigma\text{-dim}(X) \geq \sigma\text{-dim}(Y)$.

Now let X and Y be arbitrary. Let Y_1 be an irreducible component of Y such that $\sigma\text{-dim}(Y) = \sigma\text{-dim}(Y_1)$. Then Y_1 is contained in an irreducible component X_1 of X . By the above,

$$\sigma\text{-dim}(Y) = \sigma\text{-dim}(Y_1) \leq \sigma\text{-dim}(X_1) \leq \sigma\text{-dim}(X).$$

□

Exercise 4.2.20. *Let X be an irreducible σ -variety and $Y \subset X$ a σ -subvariety with $\sigma\text{-dim}(Y) = \sigma\text{-dim}(X)$. Is $Y = X$?*

Theorem 4.2.21. *Let $f: X \rightarrow Y$ be a dominant morphism of k - σ -varieties. Then $\sigma\text{-dim}(X) \geq \sigma\text{-dim}(Y)$.*

Proof. By Proposition 2.3.3, the dual map $f^*: k\{Y\} \rightarrow k\{X\}$ is injective. First assume that X and Y are irreducible. Then we obtain an inclusion of σ -function fields $k\langle Y \rangle \hookrightarrow k\langle X \rangle$. Therefore $\sigma\text{-dim}(Y) \leq \sigma\text{-dim}(X)$.

Now let X and Y be arbitrary. Let Y_1 be an irreducible component of Y such that $\sigma\text{-dim}(Y_1) = \sigma\text{-dim}(Y)$. For every minimal prime ideal \mathfrak{p} of $k\{Y\}$ there exists a minimal

prime ideal \mathfrak{p}' of $k\{X\}$ with $(f^*)^{-1}(\mathfrak{p}') = \mathfrak{p}$. (Cf. Exercise 1.2.27.) But the minimal prime ideals of $k\{Y\}$ correspond to the irreducible components of Y . Thus, there exists an irreducible component X_1 of X , such that we have a well-defined, dominant induced morphism $X_1 \rightarrow Y_1$. By the above,

$$\sigma\text{-dim}(Y) = \sigma\text{-dim}(Y_1) \leq \sigma\text{-dim}(X_1) \leq \sigma\text{-dim}(X).$$

□

Exercise 4.2.22. * Let k be a σ -field and $\mathfrak{p} \subset k\{y_1, \dots, y_n\}$ a σ -prime ideal. Can you compute the σ -dimension of $X := \mathbb{V}(\mathfrak{p})$ from a characteristic set of \mathfrak{p} ?

4.3 The limit degree

In this section we define and study an important invariant of σ -algebraic σ -field extensions: The limit degree.

As usual, if $L|K$ is an extension of fields, we denote the dimension of L as a vector space over K by $[L : K]$ and call this the *degree of L over K* . The degree is either a natural number or ∞ . The usual rules apply for calculations or estimates involving the symbol ∞ . We will need to use the following properties of the degree:

Lemma 4.3.1. *Let $K \subset L \subset M$ be fields.*

- (i) *We have $[M : K] = [M : L] \cdot [L : K]$.*
- (ii) *If $A \subset M$ then $[L(A) : L] \leq [K(A) : K]$.*
- (iii) *If $A \subset M$ then $[L(A) : K(A)] \leq [L : K]$. If moreover A is algebraically independent over L then $[L(A) : K(A)] = [L : K]$.*
- (iv) *If $A_1 \subset A_2$ are finite subsets of M , then there exists a finite subset B of L such that*

$$[L(A_2) : L(A_1)] = [K(B, A_2) : K(B, A_1)].$$

Exercise 4.3.2. *Prove the above lemma.*

Theorem 4.3.3 (Existence of the limit degree). *Let $L|K$ be an extension of σ -fields and A a finite σ -generating set of L over K . Then the sequence*

$$d_i := [K(A, \dots, \sigma^i(A)) : K(A, \dots, \sigma^{i-1}(A))], \quad i = 0, 1, \dots$$

is non-increasing and therefore stabilizes. The eventual value

$$d = \lim_{i \rightarrow \infty} d_i$$

does not depend on the choice of the σ -generating set A .

Proof. To simplify the notation we write K_i for $K(A, \dots, \sigma^i(A))$. Then, using Lemma 4.3.1 (ii), we find that

$$\begin{aligned} d_i &= [K_i : K_{i-1}] = [K_{i-1}(\sigma^i(A)) : K_{i-1}] = \\ &= [\sigma(K_{i-1})(\sigma^{i+1}(A)) : \sigma(K_{i-1})] \geq [K_i(\sigma^{i+1}(A)) : K_i] = d_{i+1}. \end{aligned}$$

It remains to show that $d := \lim_{i \rightarrow \infty} d_i$ is independent of the choice of A . So let A' be another finite σ -generating set of $L|K$. Set $K'_i = K(A', \dots, \sigma^i(A'))$ and $d' = \lim_{i \rightarrow \infty} [K'_i : K'_{i-1}]$.

There exists an $m \geq 0$ such that A lies in K'_m and A' lies in K_m . Then $K'_i \subset K_{m+i}$ and $K_i \subset K'_{m+i}$ for $i = 0, 1, \dots$. So for $j \geq m$ we have

$$K_i \subset K'_{m+i} \subset K'_{j+i} \subset K_{m+j+i}.$$

This implies $[K_{m+j+i} : K_i] \geq [K'_{j+i} : K'_{m+i}]$.

On the other hand, for $i \gg 0$ we $[K_{m+j+i} : K_i] = d^{m+j}$ since $[K_{i+1} : K_i] = d$ for $i \gg 0$. Similarly, $[K'_{j+i} : K'_{m+i}] = d'^{j-m}$ for $i \gg 0$. In summary we obtain

$$d^{m+j} = [K_{m+j+i} : K_i] \geq [K'_{j+i} : K'_{m+i}] = d'^{j-m}.$$

Letting j tend to infinity, this implies $d \geq d'$. By symmetry $d = d'$ and we win. \square

Definition 4.3.4. Let $L|K$ be a finitely σ -generated extension of σ -fields. The value d , whose existence is established by Theorem 4.3.3, is called the limit degree of $L|K$ and denoted by

$$\text{ld}(L|K).$$

Exercise 4.3.5. Consider $K := \mathbb{C}(z)$ as σ -field via $\sigma(f(z)) = f(z+1)$. Extend σ to the algebraic closure $\overline{\mathbb{C}(z)}$ of $\mathbb{C}(z)$. Determine $\text{ld}(K\langle\sqrt{z}\rangle|K)$.

Remark 4.3.6. Let $L|K$ be a finitely σ -generated extension of σ -fields. Then $\text{ld}(L|K)$ is finite if and only if L is σ -algebraic over K .

Proof. Let A be a finite σ -generating set of $L|K$. If $\text{ld}(L|K)$ is finite, there exists an $n \in \mathbb{N}$ such that

$$[K(A, \dots, \sigma^i(A)) : K(A, \dots, \sigma^{i-1}(A))]$$

is finite for all $i > n$. This implies that $L = K\langle A \rangle$ is algebraic over $K(A, \dots, \sigma^n(A))$. In particular L has finite transcendence degree over K and so L is σ -algebraic over K .

Conversely, assume that L is σ -algebraic over K . Then for every element $a \in A$ there exists an integer n such that $\sigma^{n+1}(a)$ is algebraic over $K(a, \dots, \sigma^n(a))$. This shows that for n large enough $K(A, \dots, \sigma^{n+1}(A))$ is algebraic over $K(A, \dots, \sigma^n(A))$. Then $\text{ld}(L|K) \leq [K(A, \dots, \sigma^{n+1}(A)) : K(A, \dots, \sigma^n(A))]$ and the latter value is finite. \square

Remark 4.3.7. Let $L|K$ be a finitely σ -generated extension of σ -fields. Then there exists a finite σ -generating set A of $L|K$ such that

$$\text{ld}(L|K) = [K(A, \sigma(A)) : K(A)].$$

Moreover, A can be chosen to contain any given finite subset of L .

Proof. Let $B \subset L$ be finite and extend B to a finite σ -generating set C of $L|K$. By Theorem 4.3.3 we have

$$\text{ld}(L|K) = [K(C, \dots, \sigma^n(C)) : K(C, \dots, \sigma^{n-1}(C))]$$

for some $n \in \mathbb{N}_{\geq 1}$. So we can take $A := C, \dots, \sigma^{n-1}(C)$. \square

Theorem 4.3.8. Let $L|K$ be a finitely σ -generated extension of σ -fields. Then $\text{ld}(L|K) = 1$ if and only if L is finitely generated as a field extension of K .

Proof. Assume that $\text{ld}(L|K) = 1$ and choose a finite σ -generating set A for $L|K$ as in Remark 4.3.7. Then $[K(A, \sigma(A)) : K(A)] = 1$, i.e., $\sigma(A) \subset K(A)$. This implies

$$L = K\langle A \rangle \subset K(A) \subset L.$$

So A generates L as a field extension of K .

Conversely, if A is a finite set which generates L as a field extension of K , then $[K(A, \sigma(A)) : K(A)] = 1$. So $\text{ld}(L|K) = 1$. \square

Similarly to the usual degree, the limit degree is multiplicative in towers:

Proposition 4.3.9. *Let $K \subset L \subset M$ be σ -fields such that L and M are finitely σ -generated over K . Then*

$$\text{ld}(M|K) = \text{ld}(M|L) \cdot \text{ld}(L|K).$$

Proof. Let $B \subset M$ be a finite set such that $M = L\langle B \rangle$ and $\text{ld}(M|L) = [L(B, \sigma(B)) : L(B)]$. (See Remark 4.3.7.) By Lemma 4.3.1 (iv) there exists a finite subset C of L such that

$$[L(B, \sigma(B)) : L(B)] = [K(C, B, \sigma(B)) : K(C, B)].$$

By Remark 4.3.7 there exists a finite subset A of L with $C \subset A$ and $L = K\langle A \rangle$ such that $\text{ld}(L|K) = [K(A, \sigma(A)) : K(A)]$. Then

$$\begin{aligned} \text{ld}(M|L) &= [L(B, \sigma(B)) : L(B)] = [K(C, B, \sigma(B)) : K(C, B)] \geq \\ &\geq [K(A, \sigma(A), B, \sigma(B)) : K(A, \sigma(A), B)] \end{aligned}$$

by Lemma 4.3.1 (iii). Note that A, B is a σ -generating set for $M|K$. Applying Lemma 4.3.1 (i) to the inclusions $K(A, B) \subset K(A, \sigma(A), B) \subset K(A, \sigma(A), B, \sigma(B))$ yields

$$\begin{aligned} \text{ld}(M|K) &\leq [K(A, \sigma(A), B, \sigma(B)) : K(A, B)] = \\ &= [K(A, \sigma(A), B, \sigma(B)) : K(A, \sigma(A), B)] \cdot [K(A, \sigma(A), B) : K(A, B)] \leq \\ &\leq \text{ld}(M|L) \cdot [K(A, \sigma(A)) : K(A)] = \text{ld}(M|L) \cdot \text{ld}(L|K), \end{aligned}$$

where the last estimate comes from Lemma 4.3.1 (iii).

It remains to show that $\text{ld}(M|K) \geq \text{ld}(M|L) \cdot \text{ld}(L|K)$. Note that if M is not σ -algebraic over K , then at least one of the extensions $L|K$, $M|L$ is not σ -algebraic by Lemma 4.1.6. In this case the theorem is true because both sides of the equality are equal to ∞ by Remark 4.3.6. So we can assume that M is σ -algebraic over K . In this case all three values $\text{ld}(M|K)$, $\text{ld}(M|L)$ and $\text{ld}(L|K)$ are finite. Let, as above, A be a finite σ -generating set of $L|K$ such that $\text{ld}(L|K) = [K(A, \sigma(A)) : K(A)]$. By Remark 4.3.7 we can find a finite σ -generating set B of $M|L$ such that $\text{ld}(M|L) = [K(A, B, \sigma(A), \sigma(B)) : K(A, B)]$. Let $D \subset B$ be a transcendence basis of $K(A, B)$ over $K(A)$ and $n = [K(A, B) : K(A, D)]$. (Because $K(A, B)|K(A, D)$ is a finitely generated algebraic field extension n is finite.) To simplify the notation we write $\theta_i(A)$ instead of $A, \sigma(A), \dots, \sigma^i(A)$ for $i \in \mathbb{N}$. Since $\text{ld}(L|K)$ is finite $K(\theta_{i+1}(A))|K(\theta_i(A))$ is algebraic for $i \in \mathbb{N}$. Therefore also $K(\theta_i(A))|K(A)$ is algebraic for $i \in \mathbb{N}$. Because D is algebraically independent over $K(A)$, this implies that D is also algebraically independent over $K(\theta_i(A))$. It now follows from Lemma 4.3.1 (iii) that

$$[K(\theta_i(A), D) : K(A, D)] = [K(\theta_i(A)) : K(A)] = \text{ld}(L|K)^i.$$

Considering the inclusions $K(A, D) \subset K(A, B) \subset K(\theta_i(A), B)$ we find that

$$\begin{aligned} [K(\theta_i(A), B) : K(A, B)] \cdot [K(A, B) : K(A, D)] &= [K(\theta_i(A), B) : K(A, D)] \geq \\ &\geq [K(\theta_i(A), D) : K(A, D)] = \text{ld}(L|K)^i. \end{aligned}$$

Therefore

$$[K(\theta_i(A), B) : K(A, B)] \geq \text{ld}(L|K)^i/n.$$

Considering the inclusions $K(A, B) \subset K(\theta_i(A), B) \subset K(\theta_i(A, B))$, using Lemma 4.3.1 (iii) and the fact that $[L(\theta_i(B)) : L(\theta_{i-1}(B))] \geq \text{ld}(M|L)$ we find that

$$\begin{aligned} \text{ld}(M|K)^i &= [K(\theta_i(A, B)) : K(A, B)] = \\ &= [K(\theta_i(A, B)) : K(\theta_i(A), B)] \cdot [K(\theta_i(A), B) : K(A, B)] \geq \\ &\geq [L(\theta_i(B)) : L(B)] \cdot \text{ld}(L|K)^i/n = \text{ld}(M|L)^i \text{ld}(L|K)^i/n \end{aligned}$$

Letting i tend to infinity this yields $\text{ld}(M|K) \geq \text{ld}(M|L) \cdot \text{ld}(L|K)$. \square

Let $L|K$ be a finitely σ -generated extension of σ -fields. If $K \subset L' \subset L$ is an intermediate σ -field which is also finitely σ -generated¹ over K , then $\text{ld}(L'|K) \leq \text{ld}(L|K)$ by Proposition 4.3.9. This shows that

$$\text{ld}(L|K) = \sup_{L'} \{\text{ld}(L'|K)\},$$

where L' ranges over all finitely σ -generated intermediate σ -field extensions of $L|K$ (including L). With the help of this formula we can extend the definition of the limit degree to arbitrary σ -field extensions.

Definition 4.3.10. *Let $L|K$ be an extension of σ -fields. We define the limit degree of $L|K$ by*

$$\text{ld}(L|K) = \sup_{L'} \{\text{ld}(L'|K)\},$$

where L' ranges over all finitely σ -generated intermediate σ -field extensions of $L|K$.

Proposition 4.3.9 can now be improved to:

Theorem 4.3.11. *Let $K \subset L \subset M$ be σ -fields. Then*

$$\text{ld}(M|K) = \text{ld}(M|L) \cdot \text{ld}(L|K).$$

Proof. Let A be a finite subset of L and B a finite subset of M . By Proposition 4.3.9

$$\text{ld}(M|K) \geq \text{ld}(K\langle A, B \rangle|K) = \text{ld}(K\langle A, B \rangle|K\langle A \rangle) \cdot \text{ld}(K\langle A \rangle|K).$$

As before we write $\theta_i(B)$ for $B, \dots, \sigma^i(B)$. For $i \gg 0$ we have

$$\text{ld}(K\langle A, B \rangle|K\langle A \rangle) = [K\langle A \rangle(\theta_i(B)) : K\langle A \rangle(\theta_{i-1}(B))] \geq [L(\theta_i(B)) : L(\theta_{i-1}(B))] = \text{ld}(L\langle B \rangle|L)$$

by Lemma 4.3.1 (iii).

Combining the above two equations yields

$$\text{ld}(M|K) \geq \text{ld}(L\langle B \rangle|L) \cdot \text{ld}(K\langle A \rangle|K).$$

This shows that $\text{ld}(M|K) \geq \text{ld}(M|L) \cdot \text{ld}(L|K)$.

It remains to prove that $\text{ld}(M|K) \leq \text{ld}(M|L) \cdot \text{ld}(L|K)$. To do this we can obviously assume that $\text{ld}(M|L)$ and $\text{ld}(L|K)$ are finite. Let $C \subset M$ be finite. We have to show that $\text{ld}(K\langle C \rangle|K) \leq \text{ld}(M|L) \cdot \text{ld}(L|K)$.

¹Eventually we will show that this is always the case. See Theorem 4.4.1

By definition $\text{ld}(L\langle C\rangle|L) \leq \text{ld}(M|L)$. So

$$[L(\theta_n(C)) : L(\theta_{n-1}(C))] \leq \text{ld}(M|L)$$

for some large enough $n \in \mathbb{N}$. It follows from Lemma 4.3.1 (iv) that there exists a finite set $A \subset L$ such that $[K(A, \theta_n(C)) : K(A, \theta_{n-1}(C))] = [L(\theta_n(C)) : L(\theta_{n-1}(C))]$. Then, using Lemma 4.3.1 (iii) again, we find that

$$\begin{aligned} [K\langle A\rangle(\theta_n(C)) : K\langle A\rangle(\theta_{n-1}(C))] &\leq [K(A, \theta_n(C)) : K(A, \theta_{n-1}(C))] = \\ &= [L(\theta_n(C)) : L(\theta_{n-1}(C))] \leq \text{ld}(M|L). \end{aligned}$$

This shows that $\text{ld}(K\langle A, C\rangle|K\langle A\rangle) \leq \text{ld}(M|L)$. Using Proposition 4.3.9 we can now evaluate $\text{ld}(K\langle A, C\rangle|K)$ in two ways. Firstly,

$$\text{ld}(K\langle A, C\rangle|K) = \text{ld}(K\langle A, C\rangle|K\langle A\rangle) \cdot \text{ld}(K\langle A\rangle|K) \leq \text{ld}(M|L) \cdot \text{ld}(L|K).$$

Secondly,

$$\text{ld}(K\langle A, C\rangle|K) = \text{ld}(K\langle A, C\rangle|K\langle C\rangle) \cdot \text{ld}(K\langle C\rangle|K).$$

It follows from the last two equations that $\text{ld}(K\langle C\rangle|K) \leq \text{ld}(M|L) \cdot \text{ld}(L|K)$ and we are done. \square

Exercise 4.3.12. Let K be a σ -field. Show that $\text{ld}(K^*|K) = 1$. More generally, let $L|K$ be an extension of σ -fields. Show that $\text{ld}(L^*|K^*) = \text{ld}(L^*|K) = \text{ld}(L|K)$.

Exercise 4.3.13. Let $L|K$ be a finitely σ -generated extension of σ -fields. Assume that the field extension $L|K$ is algebraic but not finite. Show that $\text{ld}(L|K) > 1$.

4.4 Finitely generated extensions of difference fields

A rather basic result in the theory of fields is that an intermediate field of a finitely generated field extension is finitely generated. The purpose of this section is to establish the σ -analog of this result.

Theorem 4.4.1. Let $K \subset L \subset M$ be σ -fields. If M is finitely σ -generated over K , then L is finitely σ -generated over K .

Proof. We start with a special case:

Claim: Let A be a finite subset of M which is σ -algebraically independent over L . If $L\langle A\rangle$ is finitely σ -generated over K , then also L is finitely σ -generated over K .

Let B be a finite σ -generating set of $L\langle A\rangle$ over K . The elements of B are quotients of σ -polynomials in A with coefficients in L . Let $C \subset L$ denote the (finite) set of all coefficients of these σ -polynomials. Then $L\langle A\rangle = K\langle B\rangle = K\langle C, A\rangle$. We will show that C is a σ -generating set for L over K . Let $d \in L$ and write $d = f(A)/g(A)$ for some σ -polynomials f, g with coefficients in $K\langle C\rangle$. Because A is σ -algebraically independent over $K\langle C\rangle \subset L$, we can compare coefficients in the equality $d \cdot g(A) = f(A)$, to find that $d \in K\langle C\rangle$. This proves the claim.

Next we will reduce to the case that $\sigma\text{-trdeg}(M|K) = 0$. So, we assume that the theorem holds under this condition. Let A be a σ -transcendence basis of $M|L$ and let B be a σ -transcendence basis of $L\langle A\rangle|K$. It follows from Theorem 4.2.14 that B is finite and that

$$\sigma\text{-trdeg}(M|K\langle B\rangle) = \sigma\text{-trdeg}(M|L\langle A\rangle) + \sigma\text{-trdeg}(L\langle A\rangle|K\langle B\rangle) = 0.$$

So we can apply the theorem to the σ -fields $K\langle B \rangle \subset L\langle A \rangle \subset M$ to deduce that $L\langle A \rangle$ is finitely σ -generated over $K\langle B \rangle$. But then $L\langle A \rangle$ is also finitely σ -generated over K . By the claim above, this implies that L is finitely σ -generated over K .

So from now we can assume that $\sigma\text{-trdeg}(M|K) = 0$. By Remark 4.3.6 this implies that $\text{ld}(M|K)$ is finite. By Theorem 4.3.11 also $\text{ld}(L|K)$ is finite and so there exists a finite set $A \subset L$ such that $\text{ld}(K\langle A \rangle|K) = \text{ld}(L|K)$. Consider the inclusions $K\langle A \rangle \subset L \subset M$. Note that $\text{ld}(L|K\langle A \rangle) = 1$. Moreover, if L is finitely σ -generated over $K\langle A \rangle$ then L is also finitely σ -generated over K . This shows that we can reduce to the case that $\text{ld}(L|K) = 1$.

From now we can assume that $\sigma\text{-trdeg}(M|K) = 0$ and $\text{ld}(L|K) = 1$. Let A be a finite σ -generating set of $M|K$. Since $\text{ld}(M|K) = \text{ld}(M|L)$, there exists an $m \geq 1$ such that

$$[K(\theta_i(A)) : K(\theta_{i-1}(A))] = [L(\theta_i(A)) : L(\theta_{i-1}(A))] = \text{ld}(M|L) \quad (4.1)$$

for $i \geq m$. We claim that $L \subset K(\theta_m(A))$. Suppose the contrary. Then there exists $b \in L$ such that $b \notin K(\theta_m(A))$. Because M is the union of all the fields $K(\theta_i(A))$ ($i \geq m$), there exists a $j > m$ such that $b \in K(\theta_j(A))$ but $b \notin K(\theta_{j-1}(A))$. Using Lemma 4.3.1 (iii) we find

$$\begin{aligned} [K(\theta_j(A)) : K(\theta_{j-1}(A))] &> [K(\theta_j(A)) : K(b, \theta_{j-1}(A))] \geq \\ &\geq [L(\theta_j(A)) : L(b, \theta_{j-1}(A))] = [L(\theta_j(A)) : L(\theta_{j-1}(A))]. \end{aligned}$$

This contradicts equation (4.1). So $L \subset K(\theta_m(A))$. An intermediate field of a finitely generated field extension is finitely generated. Therefore L is finitely generated as a field extension of K . A fortiori, L is finitely σ -generated over K . \square

4.5 The core and Babbitt's decomposition

Babbitt's decomposition is an important structure theorem for certain extensions $L|K$ of σ -fields such that the underlying field extension is algebraic. We start by recalling some elementary properties of algebraic field extension:

An algebraic field extension $L|K$ is called *normal* if it satisfies the following equivalent conditions:

- (i) L is the splitting field of a family $(f_i)_{i \in I}$ of univariate polynomials over K . That is, every f_i splits into linear factors over L and L is generated (as a field extension of K) by all the roots of the f_i in L .
- (ii) If \bar{K} is an algebraic closure of K containing L and $\tau: L \rightarrow \bar{K}$ a K -morphism, then $\tau(L) = L$.

An algebraic field extension $L|K$ is called *Galois* if it is normal and separable. For later use we record the following lemma.

Lemma 4.5.1. *Let $L|K$ and $M|K$ be field extensions contained in some field extension of K and let N denote the compositum of $L|K$ and $M|K$.*

- (i) *If $L|K$ is finite and Galois and $L \cap M = K$, then L and M are linearly disjoint over K .*
- (ii) *If $L|K$ and $M|K$ are finite Galois, then*

$$[N : K] = \frac{[L : K][M : K]}{[L \cap M : K]}$$

and $L \cap M|K$ is Galois.

If K is a σ -field, $f \in K[y]$ a univariate polynomial and $i \in \mathbb{N}$, we denote by

$$\sigma^i f \in K[y]$$

the polynomial obtained from f by applying σ^i to the coefficients of f . Note that if f is a separable polynomial, i.e., 1 lies in the ideal generated by f and its formal derivative, then also $\sigma^i f$ is separable.

Definition 4.5.2. Let $L|K$ be an extension of σ -fields. The core of $L|K$, denoted by

$$\text{Core}(L|K),$$

consists of all elements $a \in L$ which are separably algebraic over K and satisfy $\text{ld}(K\langle a \rangle|K) = 1$.

In other words, an element $a \in L$ which is separably algebraic over K lies in $\text{Core}(L|K)$ if and only if $K\langle a \rangle$ is a finite field extension of K . (Cf. Theorem 4.3.8.)

Remark 4.5.3. Let $L|K$ be an extension of σ -fields. Then $\text{Core}(L|K)$ is the largest separable algebraic σ -field extension of K inside L with limit degree one over K . If $L|K$ is finitely σ -generated, then $\text{Core}(L|K)$ is a finite field extension of K .

Proof. First of all, let us confirm that $\text{Core}(L|K)$ is a σ -field. Let $a \in \text{Core}(L|K)$ and let $f \in K[y]$ denote the minimal polynomial of a over K . Since $\sigma^i(a)$ is a root of $\sigma^i f$ and $\sigma^i f$ is separable for $i \geq 0$ we find that $K\langle a \rangle$ is a finite separable field extension of K . Therefore $K\langle a \rangle \subset \text{Core}(L|K)$ and $\text{Core}(L|K)$ is a σ -field extension of K . It is also clear that $\text{Core}(L|K)$ is a separable algebraic σ -field extension of K with limit degree one over K .

If $M \subset L$ is a separable algebraic σ -field extension of K with $\text{ld}(M|K) = 1$ and $a \in M$, then $\text{ld}(K\langle a \rangle|K) = 1$ and a is separable algebraic over K . Therefore $M \subset \text{Core}(L|K)$.

If $L|K$ is finitely σ -generated, then also $\text{Core}(L|K)|K$ is finitely σ -generated by Theorem 4.4.1. It follows from Theorem 4.3.8 that $L|K$ is finite. \square

For later use we record some lemmas.

Lemma 4.5.4. Let $L|K$ be an extension of σ -fields. Then

$$\text{Core}(L|\text{Core}(L|K)) = \text{Core}(L|K).$$

Proof. Let $a \in \text{Core}(L|\text{Core}(L|K))$. Since a is separably algebraic over $M := \text{Core}(L|K)$ and M is separably algebraic over K we see that a is separably algebraic over K . We have

$$\text{ld}(K\langle a \rangle|K) \leq \text{ld}(M\langle a \rangle|K) = \text{ld}(M\langle a \rangle|M) \cdot \text{ld}(M|K) = 1 \cdot 1 = 1.$$

Thus $a \in \text{Core}(L|K)$. \square

Lemma 4.5.5. Let $L|K$ be an algebraic extension of σ -fields with $\text{ld}(L|K) = 1$. Then, if K is invertive, also L is invertive.

Proof. Let $a \in L$ and set $M = K\langle a \rangle$. Since $\text{ld}(M|K) = 1$, we know from Theorem 4.3.8 that M is finitely generated as field extension of K . Therefore, $M|K$ is finite. As K is invertive

$$[M : K] = [\sigma(M) : \sigma(K)] = [\sigma(M) : K].$$

Consequently $M = \sigma(M)$ and $a \in \sigma(M) \subset \sigma(L)$. Thus $\sigma(L) = L$. \square

Lemma 4.5.6. *Let $L|K$ be an extension σ -fields and let $K^* \subset L^*$ denote the inverse closures. Then*

$$\text{Core}(LK^*|K^*) = \text{Core}(L|K)^*.$$

Proof. Let $M = K^*(a) \subset LK^*$ be a finite separable σ -field extension of K^* . We will show that $\sigma^n(a) \in \text{Core}(L|K)$ for some $n \in \mathbb{N}$. Since $\sigma(a)$ can be expressed as a polynomial in a with coefficients in K^* , we see that for $i \gg 0$, we can express $\sigma^{i+1}(a)$ as polynomial in $\sigma^i(a)$ with coefficients in K . Therefore $K(\sigma^i(a))$ is a σ -field extension of K for $i \gg 0$. Since a satisfies a separable polynomial over K^* also $\sigma^i(a)$ is separably algebraic over K . Therefore $\sigma^i(a) \in \text{Core}(L|K)$ for $i \gg 0$.

Since $\text{Core}(LK^*|K^*)$ is inverse by Lemma 4.5.5 and $\text{Core}(L|K) \subset \text{Core}(LK^*|K^*)$ we have $\text{Core}(LK^*|K^*) = \text{Core}(L|K)^*$ by Lemma 1.1.23. \square

The following lemma is a rather basic result in the theory of field extensions and you are probably already aware of it. In fact, a similar result was already used in the proof of Theorem 4.1.7.

Lemma 4.5.7. *Let K be a σ -field and L an algebraically closed field containing K . Then $\sigma: K \rightarrow K$ can be extended to $\sigma: L \rightarrow L$.*

Proof. Let A be a transcendence basis of $L|K$. Because $\sigma(K) \subset K$ we see that A is algebraically independent over $\sigma(K)$ and so we can define an isomorphism

$$\sigma: K(A) \rightarrow \sigma(K)(A) \subset K(A)$$

by setting $\sigma(a) = a$ for $a \in A$. Because L is algebraic over $K(A)$ we can reduce to the case that $L = \overline{K}$ is the algebraic closure of K . Let \mathcal{M} denote the set of all pairs (M, σ_M) where $K \subset M \subset \overline{K}$ is an intermediate field and $\sigma_M: M \rightarrow \overline{K}$ an extension of $\sigma: K \rightarrow K \subset \overline{K}$. We can define a partial order on \mathcal{M} by setting $(M, \sigma_M) \leq (M', \sigma_{M'})$ if $M \subset M'$ and $\sigma_{M'}$ extends σ_M . Clearly, every chain in \mathcal{M} has an upper bound in \mathcal{M} . Thus, by Zorn's lemma, there exists a maximal element (M, σ) in \mathcal{M} . We will show that $M = \overline{K}$. Suppose the contrary. Then there exists $a \in \overline{K}$ with $a \notin M$. Let f denote the minimal polynomial of a over M . Because \overline{K} is algebraically closed σf has a root b in \overline{K} . Then we can extend $\sigma: M \rightarrow \overline{K}$ to $\sigma: M(a) \rightarrow \overline{K}$ by sending a to b . This contradicts the maximality of (M, σ) . \square

Note that the extension of σ from K to L in the above lemma need not be unique.

⚠ 4.5.8. *Let K be a σ -field and L a field extension of K . It might not be possible to extend σ from K to L . Take for example $K = \mathbb{C}(z)$ with $\sigma(f(z)) = f(z+1)$ and $L = K(\sqrt{z})$. In fact, σ does not extend to any finite field extension L of $\mathbb{C}(z)$. (Cf. Proposition 4.5.22.)*

Lemma 4.5.9. *Let K be a σ -field and \overline{K} an algebraic closure of K . Let $M \subset \overline{K}$ be a normal extension of K and $i \in \mathbb{N}$. Choose an extension of $\sigma: K \rightarrow K$ to an endomorphism $\sigma: \overline{K} \rightarrow \overline{K}$. Then the field $K(\sigma^i(M)) \subset \overline{K}$ is normal over K and does not depend on the choice of the extension of σ . If $M|K$ is Galois then $K(\sigma^i(M))|K$ is Galois.*

Proof. Let $(f_j)_{j \in J}$ be a family of univariate polynomials over K such that M is the splitting field of $(f_j)_{j \in J}$. It suffices to show that $K(\sigma^i(M))$ is the splitting field of $(\sigma^i f_j)_{j \in J}$. We may assume that the f_j 's are normalized. We have $f_j = (y - a_1) \cdots (y - a_n)$ for some $a_1, \dots, a_n \in L$. Then $\sigma^i f_j = (y - \sigma^i(a_1)) \cdots (y - \sigma^i(a_n))$. It follows that $K(\sigma^i(M))$ is the splitting field for $(\sigma^i f_j)_{j \in J}$ over K . If $M|K$ is separable, the f_j 's can be chosen to be separable and we see that also $K(\sigma^i(M))|K$ is separable. \square

4.5.10. *The above lemma is not true without the condition that M is normal over K . For example, let $K = \mathbb{Q}$ (considered as a constant σ -field) and let $M \subset \overline{\mathbb{Q}}$ be an extension which is not normal. Then there exists a morphism $\sigma: M \rightarrow \overline{\mathbb{Q}}$ with $\sigma(M) \neq M$. Extend σ to $\sigma: \overline{\mathbb{Q}} \rightarrow \overline{\mathbb{Q}}$. Also the identity map is an extension of $\sigma: \mathbb{Q} \rightarrow \mathbb{Q}$. But $\text{id}(M) \neq \sigma(M)$.*

Let K be a σ -field and \overline{K} an algebraic closure of K . Let $M \subset \overline{K}$ be a normal extension of K and $i \in \mathbb{N}$. By Lemma 4.5.9, the field extension $K(\sigma^i(M)) \subset \overline{K}$ does not depend on the choice of an extension of σ to \overline{K} . Therefore, we will still denote it with

$$K(\sigma^i(M)),$$

even if no extension of σ to the algebraic closure has been fixed.

Corollary 4.5.11. *Let K be a σ -field and $M|K$ a field extension. Let $\sigma_1, \sigma_2: M \rightarrow M$ denote extensions of $\sigma: K \rightarrow K$. If $K \subset L \subset M$ is an intermediate field with $L|K$ normal, then L is stable under σ_1 if and only if L is stable under σ_2 .*

Proof. The relative algebraic closure of K in M is stable under σ_1 and σ_2 and we can extend σ_1 and σ_2 to the algebraic closure of K . Thus we can assume that M is the algebraic closure of K . Since $K(\sigma_1(L)) = K(\sigma_2(L))$ by Lemma 4.5.9 we have $\sigma_1(L) \subset L$ if and only if $\sigma_2(L) \subset L$. \square

Lemma 4.5.12. *Let $K \subset L \subset M$ be σ -fields such that K is inversive and $M|K$ is normal. (E.g. $M = \overline{K}$.) Then the normal closure of L over K in M is a σ -field, i.e., stable under σ .*

Proof. Let $N \subset M$ denote the normal closure of $L|K$. Let f be the minimal polynomial (over K) of an element a from L and let $a_1, \dots, a_n \in M$ denote the roots of f . It suffices to show that $\sigma(a_1), \dots, \sigma(a_n) \in N$. Because K is inversive, σf is irreducible (over K). Since σf has the root $\sigma(a) \in L \subset N$, this implies that all the roots of σf lie in N . But $\sigma(a_1), \dots, \sigma(a_n)$ are the roots of σf . \square

Corollary 4.5.13. *Let K be an inversive σ -field, $L|K$ an algebraic σ -field extension and N a normal closure of $L|K$. Then there exists an extension $\sigma: N \rightarrow N$ of $\sigma: L \rightarrow L$. Moreover, if $L|K$ is finitely σ -generated, then $N|K$ is finitely σ -generated.*

Proof. Let \overline{K} denote an algebraic closure of K containing L . By Lemma 4.5.7 we can extend $\sigma: L \rightarrow L$ to $\sigma: \overline{K} \rightarrow \overline{K}$ and by Lemma 4.5.12 $N \subset \overline{K}$ is stable under σ .

If $L|K$ is σ -generated by $a_1, \dots, a_n \in L$, then $N|K$ is σ -generated by all the conjugates of a_1, \dots, a_n over K . \square

Exercise 4.5.14. *Do we really need the assumption that K is inversive in Lemma 4.5.12?*

Exercise 4.5.15. *Let K be an inversive σ -field and extend σ to the algebraic closure \overline{K} of K . Show that \overline{K} is inversive.*

Lemma 4.5.16. *Let K be an inversive σ -field and $L|K$ a normal σ -field extension of K . Then $\text{Core}(L|K)|K$ is Galois.*

Proof. Let M be a finite σ -field extension of K contained in $\text{Core}(L|K)$ and let $N \subset L$ denote the normal closure of $L|K$. Then $N|K$ is a Galois σ -field extension by Lemma 4.5.12. Since $N|K$ is finite, $N \subset \text{Core}(L|K)$. Therefore $\text{Core}(L|K)$ is Galois. \square

Exercise 4.5.17. Consider $K = \mathbb{Q}(z)$ as σ -field via $\sigma(f(z)) = f(z+1)$. Extend σ to \overline{K} and choose a fourth root $\sqrt[4]{z} \in \overline{K}$ of z . Consider the σ -field extension $L := K\langle\sqrt[4]{z}\rangle$ of K . Show that the core of the normal closure of L over K is strictly larger than the normal closure of the core of $L|K$.

Lemma 4.5.18. Let K be an inversive σ -field and $L|K$ a normal field extension of K . Let $\sigma_1, \sigma_2: L \rightarrow L$ denote extensions of $\sigma: K \rightarrow K$. Then the underlying field of the core of $L|K$ with respect to σ_1 agrees with the underlying field of the core of $L|K$ with respect to σ_2 .

Proof. Let $N \subset L$ be a finite Galois σ -field extension of K contained in the core of $L|K$ with respect to σ_1 . Let f be an irreducible polynomial over K such that N is the splitting field of f . Since σ_2 must map the roots of f to the roots of σf and the roots of σf all lie in N we see that $\sigma_2(N) \subset N$. This shows that the core of $L|K$ with respect to σ_1 is contained in the core of $L|K$ with respect to σ_2 . By symmetry we are done. \square

Definition 4.5.19. Let K be an inversive σ -field. By Lemma 4.5.18 above, the underlying field of the core of $\overline{K}|K$ does not depend on the choice of the extension of σ to the algebraic closure \overline{K} of K . We simply call it the core of K and denote it with

$$\text{Core}(K).$$

4.5.20. The core of K is only a field extension and not a σ -field extension.

Clearly the core of K is only unique up to isomorphisms (since the algebraic closure is only unique up to isomorphisms).

Example 4.5.21. Let K be a constant σ -field. Then $\text{Core}(K)$ is the separable algebraic closure of K .

Proposition 4.5.22. Consider the field $\mathbb{C}(z)$ of rational functions in the variable x over \mathbb{C} as σ -field by virtue of $\sigma(f(z)) = f(z+1)$. Then $\text{Core}(\mathbb{C}(z)) = \mathbb{C}(z)$.

The proof uses some facts about “ramification”. Unfortunately it is out of reach for us. The same applies to the next proposition.

Proposition 4.5.23. Consider the field $\mathbb{C}(z)$ of rational functions in the variable z over \mathbb{C} as σ -field by virtue of $\sigma(f(z)) = f(qz)$, where $q \in \mathbb{C}^\times$ is not a root of unity. Then

$$\text{Core}(\mathbb{C}(z)) = \bigcup_{n \geq 1} \mathbb{C}(z^{\frac{1}{n}}).$$

The inclusion “ \supset ” is easy to see: Choose an n -th root $q^{\frac{1}{n}} \in \mathbb{C}$ of q . Then we can turn $\mathbb{C}(z^{\frac{1}{n}})$ into a σ -field extension of $\mathbb{C}(z)$ by setting $\sigma(z^{\frac{1}{n}}) = q^{\frac{1}{n}} z^{\frac{1}{n}}$.

Definition 4.5.24. A σ -field extension $L|K$ is called benign if it is of the form $L = [\sigma]_K M$, where M is a finite Galois extension of K .

4.5.25. Let K be a σ -field and M a finite Galois extension of K . Then, in general, $[\sigma]_K M$ need to be a field. Take for example, $K = \mathbb{Q}$ and $M = \mathbb{Q}(\sqrt{2})$.

Let $L|K$ be a field extension and let $(L_i)_{i \in I}$ be a family of intermediate fields. Recall that the family is called linearly disjoint over K if for every finite subset J of I the canonical map $\otimes_{j \in J} L_j \rightarrow L$ is injective.

Proposition 4.5.26. *Let K be a σ -field and M a finite Galois extension of K , contained in an algebraic closure \overline{K} of K . Choose $a \in M$ such that $M = K(a)$ and let $f = f(y)$ denote the minimal polynomial of a over K . The following statements are equivalent:*

- (i) *The K - σ -algebra $L := [\sigma]_K M$ is a field (i.e., $L|K$ is benign).*
- (ii) *For every $i \in \mathbb{N}$ we have $[K(\sigma^i(M)) : K] = [M : K]$ and the fields $(K(\sigma^i(M)))_{i \in \mathbb{N}}$ are linearly disjoint over K .*
- (iii) *For every $i \in \mathbb{N}$ the polynomial $\sigma^i f$ is irreducible over the splitting field of $f, \dots, \sigma^{i-1} f$.*
- (iv) *The σ -ideal $[f] \subset K\{y\}$ is prime.*
- (v) *The σ -ideal $[f] \subset K\{y\}$ is σ -prime.*

Proof. Note that M is the splitting field of f over K . The field $K(\sigma^i(M))$ is the splitting field of $\sigma^i f$ over K .

We recall the construction of $[\sigma]_K M$ from subsection 1.1.7: One sets $\sigma^i M = M \otimes_K K$ where the tensor product is formed by using $\sigma^i: K \rightarrow K$ on the right hand side. Consider $\sigma^i M = M \otimes_K K$ as K -algebra via the right factor and set $M_i = M \otimes_K \sigma M \otimes_K \dots \otimes_K \sigma^i M$. We have natural inclusions $M_i \hookrightarrow M_{i+1}$ and $[\sigma]_K M$ is the union of the M_i 's. We have $M = K[y]/(f)$ and consequently $\sigma^i M = K[y]/(\sigma^i f)$. Note that $a \otimes 1 \in M \otimes_K K = \sigma^i M$ is a root of $\sigma^i f$ and that $a \otimes 1$ generates $\sigma^i M$ as K -algebra. We have a surjective morphism

$$\psi_i: \sigma^i M \rightarrow K(\sigma^i(M))$$

of K -algebras given by $\psi_i(a \otimes 1) = \sigma^i(a)$.

(i) \Rightarrow (ii): Since $[\sigma]_K M$ is a field, also $\sigma^i M$ must be a field. It follows that ψ_i is an isomorphism. Since the degree of $\sigma^i M = K[y]/(\sigma^i f)$ over K is $[M : K]$ we see that $[K(\sigma^i(M)) : K] = [M : K]$. Since $M_i = M \otimes_K \dots \otimes_K \sigma^i M \simeq M \otimes_K \dots \otimes_K K(\sigma^i(M))$ is a field, it is clear that the family $(K(\sigma^i(M)))_{i \in \mathbb{N}}$ is linearly disjoint over K .

(ii) \Rightarrow (iii): Because $[K(\sigma^i(M)) : K] = [M : K]$ the morphism ψ_i is an isomorphism. Fix $i \in \mathbb{N}$ and let N denote the compositum over K of the fields $M, K(\sigma(M)), \dots, K(\sigma^{i-1}(M))$. Then N is the splitting field of $f, \dots, \sigma^{i-1} f$. By assumption, N is linearly disjoint from $K(\sigma^i(M))$ over K . Therefore $N \otimes_K K(\sigma^i(M))$ is a field. But as

$$N \otimes_K K(\sigma^i(M)) \simeq N \otimes_K K[y]/(\sigma^i f) = N[y]/(\sigma^i f)$$

this shows that $\sigma^i f$ is irreducible over N .

(iii) \Rightarrow (iv): It suffices to show that for every $i \in \mathbb{N}$ the ideal generated by $f, \sigma(f), \dots, \sigma^i(f)$ in $K[y, \dots, \sigma^i(y)]$ is prime. The cases $i = 0$ being trivial, we proceed by induction on i . By the induction hypothesis $N := K[y, \dots, \sigma^{i-1}(y)]/(f, \dots, \sigma^{i-1}(f))$ is a field. So N is the splitting field of $f, \dots, \sigma^{i-1} f$. We have $K[y, \dots, \sigma^i(y)]/(f, \dots, \sigma^i(f)) \simeq N[y]/(\sigma^i f)$. By assumption the latter is a field.

(iv) \Rightarrow (v): As $[f] \subset K\{y\}$ is prime, the quotient $K\{y\}/[f]$ is an integral domain. But since $K\{y\}/[f]$ is integral over K we see that $K\{y\}/[f]$ is a field. Because σ is automatically injective on a field, this shows that $[f]$ is reflexive.

(v) \Rightarrow (i): This is clear from $[\sigma]_K M = K\{y\}/[f]$. □

Remark 4.5.27. *If $L|K$ is benign, then $L|K$ is Galois and finitely σ -generated. Moreover, if $M|K$ is a finite Galois extension such that $L = [\sigma]_K M$. Then $\text{ld}(L|K) = [M : K]$.*

Definition 4.5.28. *Let $L|K$ be a finitely σ -generated extension of σ -fields such that $L|K$ is Galois. An element $a \in L$ is called a standard generator of $L|K$ if the following properties are satisfied:*

- $K\langle a \rangle = L$.
- $[K(a, \sigma(a)) : K(a)] = \text{ld}(L|K)$.
- The field extension $K(a)|K$ is Galois.

A standard generator a of $L|K$ is called minimal if $[K(a) : K] \leq [K(b) : K]$ for every standard generator b of $L|K$.

Lemma 4.5.29. *Let $L|K$ be a finitely σ -generated extension of σ -fields such that $L|K$ is Galois. Then there exists a standard generator of $L|K$.*

Proof. Let A be a finite σ -generating set of $L|K$. Enlarging A if necessary, we can assume that $K(A)|K$ is Galois. By the primitive element theorem there exists an element $a \in K(A)$ such that $K(a) = K(A)$. This implies that $K\langle a \rangle = L$. There exists an $n \in \mathbb{N}$ such that $[K(a, \dots, \sigma^{n+1}(a)) : K(a, \dots, \sigma^n(a))] = \text{ld}(L|K)$.

Because $K(a)|K$ is Galois, also $K(a, \dots, \sigma^n(a))|K$ is Galois. Again, by the primitive element theorem, there exists an element $b \in K(a, \dots, \sigma^n(a))$ such that $K(b) = K(a, \dots, \sigma^n(a))$. Then $K(b, \sigma(b)) = K(a, \dots, \sigma^{n+1}(a))$, therefore

$$[K(b, \sigma(b)) : K(b)] = [K(a, \dots, \sigma^{n+1}(a)) : K(a, \dots, \sigma^n(a))] = \text{ld}(L|K).$$

So b is a standard generator for $L|K$. □

Let $L|K$ be a finitely σ -generated extension of σ -fields such that $L|K$ is Galois and let a be a minimal standard generator of $L|K$. By Theorem 4.3.3 we have

$$[K(a) : K] \geq [K(a, \sigma(a)) : K(a)] = \text{ld}(L|K).$$

Lemma 4.5.30. *Let $L|K$ be a finitely σ -generated extension of σ -fields such that $L|K$ is Galois and let a be a minimal standard generator of $L|K$. Then $L|K$ is benign if and only if $[K(a) : K] = \text{ld}(L|K)$.*

Proof. Assume that $L|K$ is benign and let M be a finite Galois extension of K such that $L = [\sigma]_K M$. Choose $b \in M$ such that $M = K(b)$. Then b is a minimal standard generator of $L|K$ and $[K(b) : K] = \text{ld}(L|K)$.

Conversely, if a is a minimal standard generator of $L|K$ with $[K(a) : K] = \text{ld}(L|K)$. Then $L = [\sigma]_K M$ with $M = K(a)$. □

Definition 4.5.31. *An extension of σ -fields $L|K$ is called σ -radicial if for every $a \in L$ there exists an $n \in \mathbb{N}$ such that $\sigma^n(a) \in K$.*

Example 4.5.32. Let K be a σ -field. Then $K^*|K$ is σ -radicial.

Moreover, every σ -radicial extension of K is contained in K^* .

Lemma 4.5.33. *Let $L|K$ and $M|L$ be σ -radical σ -field extensions. Then $M|K$ is also σ -radicial.*

Proof. Obvious. □

Lemma 4.5.34. *Let $L|K$ be a σ -radicial extension of σ -fields and a_1, \dots, a_n elements in a σ -field extension of L such that $L\langle a_1, \dots, a_i \rangle$ is benign over $L\langle a_1, \dots, a_{i-1} \rangle$ with minimal standard generator a_i for $i = 1, \dots, n$. Then, for a given $m \in \mathbb{N}$, there exist integers $r_1, \dots, r_n \geq m$ such that $K\langle \sigma^{r_1}(a_1), \dots, \sigma^{r_i}(a_i) \rangle$ is benign over $K\langle \sigma^{r_1}(a_1), \dots, \sigma^{r_{i-1}}(a_{i-1}) \rangle$ with minimal standard generator $\sigma^{r_i}(a_i)$ for $i = 1, \dots, n$.*

Proof. We first treat the case $n = 1$. So $a := a_1$ is a minimal standard generator of $L\langle a \rangle$ over L . We claim that $K\langle \sigma^j(a) \rangle | K$ is benign with minimal standard generator $\sigma^j(a)$ for $j \gg 0$.

Let f denote the minimal polynomial of a over L . Because $L(a)$ is the splitting field of f , every root of f can be written as a polynomial in a with coefficients in L . Since $L|K$ is σ -radicial, every root of $\sigma^j f$ can be written as a polynomial in $\sigma^j(a)$ with coefficients in K for $j \gg 0$. Therefore $K\langle \sigma^j(a) \rangle$ is the splitting field of $\sigma^j f$ for $j \gg 0$. Because f is separable also $\sigma^j f$ is separable. Therefore $K\langle \sigma^j(a) \rangle | K$ is Galois.

Because L is σ -radicial over K there exists an integer i such that $\sigma^j f$ has coefficients in K for $j \geq i$. It suffices to show that $\sigma^{j+l} f$ is irreducible over $K\langle \sigma^j(a), \dots, \sigma^{j+l-1}(a) \rangle$ for $l \in \mathbb{N}$ and $j \geq i$. Suppose the contrary. Then $\sigma^{j+l} f$ is reducible over $L\langle a, \dots, \sigma^{j+l-1}(a) \rangle$. This contradicts the fact that $L\langle a \rangle | L$ is benign with minimal standard generator a . This finishes the case $n = 1$.

The general case now follows from the $n = 1$ case by induction, with $K\langle \sigma^{r_1}(a_1), \dots, \sigma^{r_{n-1}}(a_{n-1}) \rangle$ in place of K and $L\langle a_1, \dots, a_{n-1} \rangle$ in place of L . \square

Theorem 4.5.35 (Babbitt's Decomposition). *Let K be an inversive σ -field and L a finitely σ -generated σ -field extension of K such that $L|K$ is Galois. Then there exists a chain of intermediate σ -fields*

$$K \subset L_1 \subset L_2 \subset \dots \subset L_{n-1} \subset L_n \subset L$$

such that $L_1 = \text{Core}(L|K)$, L_{i+1} is benign over L_i for $i = 1, \dots, n-1$ and L is σ -radicial over L_n .

Proof. The proof works by induction on $\text{ld}(L|K)$. If $\text{ld}(L|K) = 1$ then $\text{Core}(L|K) = L$ and we are done. So we can assume that $\text{ld}(L|K) > 1$. By Lemma 4.5.5 the core $\text{Core}(L|K)$ is inversive. Replacing K with $\text{Core}(L|K)$, we can assume that $\text{Core}(L|K) = K$ (Lemma 4.5.4).

First, we also assume that $L|K$ contains no intermediate σ -field M such that $M|K$ is Galois and $1 < \text{ld}(M|K) < \text{ld}(L|K)$. We will show that $L|K$ is benign.

Let a be a minimal standard generator of $L|K$. Then $K\langle a \rangle$ and $K\langle \sigma(a) \rangle$ are Galois extensions of K . By Lemma 4.5.1

$$[K\langle a, \sigma(a) \rangle : K] = \frac{[K\langle a \rangle : K] \cdot [K\langle \sigma(a) \rangle : K]}{[K\langle a \rangle \cap K\langle \sigma(a) \rangle : K]}.$$

On the other hand

$$[K\langle a, \sigma(a) \rangle : K] = [K\langle a, \sigma(a) \rangle : K\langle a \rangle] \cdot [K\langle a \rangle : K] = \text{ld}(L|K) \cdot [K\langle a \rangle : K].$$

Because K is inversive, $[K\langle \sigma(a) \rangle : K] = [K\langle a \rangle : K]$. This, together with the last two equations yields

$$[K\langle a \rangle : K] = \text{ld}(L|K) \cdot [K\langle a \rangle \cap K\langle \sigma(a) \rangle : K].$$

Because $\text{ld}(L|K) > 1$ we see that $[K\langle a \rangle \cap K\langle \sigma(a) \rangle : K] < [K\langle a \rangle : K]$.

Let b be a primitive element of $K\langle a \rangle \cap K\langle \sigma(a) \rangle | K$. Because $b \in K\langle \sigma(a) \rangle$ and K is inversive, there exists $c \in K\langle a \rangle$ such that $\sigma(c) = b$. Since K is inversive,

$$[K\langle c \rangle : K] = [K\langle b \rangle : K] < [K\langle a \rangle : K].$$

Because $K\langle b \rangle = K\langle a \rangle \cap K\langle \sigma(a) \rangle$ is Galois over K (Lemma 4.5.1) also $K\langle c \rangle | K$ is Galois, which implies that $K\langle c \rangle | K$ is Galois. So, by assumption, $\text{ld}(K\langle c \rangle | K)$ equals 1 or $\text{ld}(L|K)$.

Suppose that $\text{ld}(K\langle c\rangle|K) = \text{ld}(L|K)$. Then $[K(c, \sigma(c)) : K(c)] \geq \text{ld}(L|K)$ and therefore

$$[K(c, \sigma(c)) : K] \geq [K(c, \sigma(c)) : K(c)] \cdot [K(c) : K] \geq \text{ld}(L|K) \cdot [K(a) \cap K(\sigma(a)) : K].$$

But $K(c, \sigma(c)) \subset K(a)$, so

$$[K(c, \sigma(c)) : K] \leq [K(a) : K] = \text{ld}(L|K) \cdot [K(a) \cap K(\sigma(a)) : K].$$

Thus $K(c, \sigma(c)) = K(a)$. But then $K\langle c\rangle = L$ and c is a standard generator for $L|K$. Since $[K(c) : K] < [K(a) : K]$, this contradicts the fact that a is a minimal standard generator of $L|K$.

Hence $\text{ld}(K\langle c\rangle|K) = 1$. But then $c \in K$, $K(a) \cap K(\sigma(a)) = K$ and $[K(a) : K] = \text{ld}(L|K)$. Thus $L|K$ is benign by Lemma 4.5.30.

Now we assume that there exists an intermediate σ -field $K \subset M \subset L$ such that $M|K$ is Galois and $1 < \text{ld}(M|K) < \text{ld}(L|K)$. To be able to apply the induction hypothesis we need to have an inversive base σ -field. So let $M^* \subset L^*$ denote the inversive closures of M and L respectively. Because L is Galois over M , also M^*L is Galois over M^* . The core of M^*L over M^* is a finite separable field extension of M^* . Choose $a \in \text{Core}(M^*L|M^*)$ such that $\text{Core}(M^*L|M^*) = M^*(a)$. Note that $\sigma(a)$ can be expressed as a polynomial in a with coefficients in M^* . Thus, for $i \gg 0$, $\sigma^{i+1}(a)$ can be expressed as a polynomial in $\sigma^i(a)$ with coefficients in M . From Lemma 4.5.5 we know that $M^*(a)$ is inversive and therefore $M^*(a) = M^*(\sigma^i(a))$ for $i \in \mathbb{N}$. Consequently, replacing a by $\sigma^i(a)$ for some suitable $i \in \mathbb{N}$ if necessary, we can assume that $a \in L$ and that $\sigma(a) \in M(a)$, so that $M(a) \subset L$ is a σ -field.

Let $A \subset L$ denote the set of all roots of the minimal polynomial of a over K . Since M is Galois over K we see that $M(A) \subset L$ is the normal closure of $M(a)|K$. By Lemma 4.5.12 $M(A)$ is a σ -field. We have

$$\text{ld}(M\langle A\rangle|K) = \text{ld}(M\langle A\rangle|M) \cdot \text{ld}(M|K) = \text{ld}(M|K) < \text{ld}(L|K).$$

In summary, we see that $M\langle A\rangle \subset L$ is a finitely σ -generated Galois σ -field extension of K with $\text{Core}(M\langle A\rangle|K) = K$ and $\text{ld}(M\langle A\rangle|K) < \text{ld}(L|K)$. Applying the induction hypothesis to $M\langle A\rangle|K$ yields a sequence of intermediate σ -fields

$$K \subset L_1 \subset \cdots \subset L_n \subset M\langle A\rangle,$$

where the extensions $L_1|K$ and $L_{i+1}|L_i$ ($1 \leq i \leq n-1$) are benign and $M\langle A\rangle|L_n$ is σ -radical.

We would also like to apply the induction hypothesis to the extension $M^*L|M^*\langle A\rangle$. Let's check the assumptions: We have $M^*\langle A\rangle = M^*(A)$ and so $M^*\langle A\rangle$ is inversive by Lemma 4.5.5. Because $L|K$ is Galois also $M^*L|M^*\langle A\rangle$ is Galois. With the help of Lemma 4.3.1 (ii) it is easy to see that $\text{ld}(M^*L|M^*\langle A\rangle) \leq \text{ld}(L|M\langle A\rangle)$. Therefore

$$\text{ld}(M^*L|M^*\langle A\rangle) \leq \text{ld}(L|M\langle A\rangle) \leq \text{ld}(L|M) < \text{ld}(L|K).$$

Moreover, since $\text{Core}(M^*L|M^*) = M^*(a)$ we have $\text{Core}(M^*L|M^*\langle A\rangle) = M^*\langle A\rangle$. Applying the induction hypothesis to $M^*L|M^*\langle A\rangle$ yields a sequence of σ -fields

$$M^*\langle A\rangle \subset M_1 \subset \cdots \subset M_m \subset M^*L$$

such that the extensions $M_1|M^*\langle A\rangle$ and $M_i|M_{i-1}$ ($2 \leq i \leq m$) are benign and $M^*L|M_m$ is σ -radical. Let $a_1 \in M_1$ be a minimal standard generator of $M_1|M^*\langle A\rangle$ and for $i = 2, \dots, m$

let $a_i \in M_i$ be a minimal standard generator of $M_i|M_{i-1}$. There exists $l \in \mathbb{N}$ such that $\sigma^l(a_i) \in L$ for $i = 1, \dots, m$.

Because the extensions $M\langle A\rangle|L_n$ and $M^*\langle A\rangle|M\langle A\rangle$ are σ -radicial we know (Lemma 4.5.33) that $M^*\langle A\rangle|L_n$ is σ -radicial. Applying Lemma 4.5.34 to the σ -radicial extension $M^*\langle A\rangle|L_n$, we find that there exist $r_1, \dots, r_m \geq l$ such that $L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_i}(a_i)\rangle$ is benign over $L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_{i-1}}(a_{i-1})\rangle$.

We have now constructed a sequence of benign σ -field extensions

$$K \subset L_1 \subset \dots \subset L_n \subset L_n\langle\sigma^{r_1}(a_1)\rangle \subset \dots \subset L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_m}(a_m)\rangle$$

inside L . It remains to see that L is σ -radicial over $L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_m}(a_m)\rangle$. But $M_m = M^*\langle A\rangle\langle a_1, \dots, a_m\rangle$ is σ -radicial over $L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_m}(a_m)\rangle$ and M^*L is σ -radicial over M_m . Therefore M^*L is σ -radicial over $L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_m}(a_m)\rangle$, a fortiori L is σ -radicial over $L_n\langle\sigma^{r_1}(a_1), \dots, \sigma^{r_m}(a_m)\rangle$. \square

Exercise 4.5.36. * Does one really need the σ -radical extension in Babbitt's decomposition?

4.6 Compatibility

In this section we deal with (in)compatible extensions of difference fields. There is no analog of this topic in classical commutative algebra or differential algebra. (See exercise 4.6.2.)

Definition 4.6.1. Let $L_1|K$ and $L_2|K$ be σ -field extensions. Then $L_1|K$ and $L_2|K$ are called compatible if there exists a σ -field extension M of K and K - σ -embeddings $L_1 \hookrightarrow M$ and $L_2 \hookrightarrow M$. Otherwise $L_1|K$ and $L_2|K$ are called incompatible.

The main result of this section is the compatibility theorem, which states that the compatibility of two σ -field extensions is determined by the compatibility of the cores.

Exercise 4.6.2. Show that any two field extensions $L_1|K$ and $L_2|K$ are compatible in the sense that there exists a field extension $M|K$ and K -embeddings $L_1 \hookrightarrow M$ and $L_2 \hookrightarrow M$.

Example 4.6.3. Let $K = \mathbb{Q}$, $L_1 = \mathbb{Q}(\sqrt{2})$ and $L_2 = \mathbb{Q}(\sqrt{2})$ where σ is the identity on L_1 and acting on L_2 by $\sigma(\sqrt{2}) = -\sqrt{2}$. Then $L_1|K$ and $L_2|K$ are incompatible.

Proof. In essence, this is because every field extension of K can only contain two roots of 2 and either σ is acting trivially on them or non-trivially. We can not have both. \square

Remark 4.6.4. Let $L_1|K$ and $L_2|K$ be σ -field extensions. Then $L_1|K$ and $L_2|K$ are compatible if and only if $\text{Spec}^\sigma(L_1 \otimes_K L_2)$ is non-empty.

Proof. If $\mathfrak{p} \in \text{Spec}^\sigma(L_1 \otimes_K L_2)$, then $k(\mathfrak{p})$ is a σ -field extension of K which comes equipped with K - σ -embeddings $L_1 \hookrightarrow k(\mathfrak{p})$ and $L_2 \hookrightarrow k(\mathfrak{p})$. Conversely, if M is a σ -field extension of K with K - σ -embeddings $\phi_1: L_1 \rightarrow M$ and $\phi_2: L_2 \rightarrow M$, then the kernel of

$$L_1 \otimes_K L_2 \longrightarrow M, f \otimes g \mapsto \phi_1(f)\phi_2(g)$$

is a σ -prime ideal of $L_1 \otimes_K L_2$. \square

Definition 4.6.5. An extension of σ -fields $L|K$ is called universally compatible if $L|K$ is compatible with any other σ -field extension of K .

The following three lemmas will provide examples of universally compatible σ -field extension.

We will need to use the following result from algebra: Let $L|K$ be a field extension such that K is relatively separably algebraically closed² in L , (i.e., every element in L which is separably algebraic over K belongs to K) then $R \otimes_K L$ has a unique minimal prime ideal for every K -algebra R which is an integral domain. I.e., the radical of R is prime.

Lemma 4.6.6. *Let $L|K$ be an extension of σ -fields such that K is relatively separably algebraically closed in L . Then $L|K$ is universally compatible.*

Proof. Let $M|K$ be an extension of σ -fields. Because K is relatively separably algebraically closed in L the nilradical \mathfrak{p} of $L \otimes_K M$ is prime. Clearly \mathfrak{p} is a σ -ideal. By Lemma 1.2.16 the reflexive closure of \mathfrak{p} is a σ -prime ideal of $L \otimes_K M$. Thus $L|K$ and $M|K$ are compatible by Remark 4.6.4. \square

A German translation of “benign” is “gutartig”. The following lemma justifies this naming.

Lemma 4.6.7. *Let $L|K$ be a benign extension of σ -fields. Then $L|K$ is universally compatible.*

Proof. Let M be a σ -field extension of K . Let \overline{M} denote an algebraic closure of M . By Lemma 4.5.7 we can extend σ from M to \overline{M} . It suffices to find a K - σ -embedding of L into M . Let $a \in L$ be a minimal standard generator of $L|K$ and let f denote the minimal polynomial of a over K . Then f has a root $b \in \overline{M}$. Because $L \simeq K\{y\}/[f]$ (Cf. Proposition 4.5.26.) we can define a K - σ -morphism

$$L \rightarrow \overline{M}, a \mapsto b.$$

\square

Lemma 4.6.8. *Let $L|K$ be a σ -radicial extension of σ -fields. Then $L|K$ is universally compatible.*

Proof. Let $M|K$ be σ -field extension. Since $L|K$ is σ -radicial the inductive closure of L agrees with the inductive closure of K , but the inductive closure of K is contained in the inductive closure of M . (Cf. Lemma 1.1.23.) Thus the inductive closure of M contains isomorphic copies of $L|K$ and $M|K$. \square

Lemma 4.6.9. *Let $L_1|K$ and $L_2|K$ be incompatible σ -field extensions. Then there exist intermediate σ -fields $K \subset M_1 \subset L_1$ and $K \subset M_2 \subset L_2$ such that $M_1|K$ and $M_2|K$ are finitely σ -generated and $M_1|K$ and $M_2|K$ are incompatible.*

Proof. By Remark 4.6.4 and Proposition 1.2.22 the extensions $L_1|K$ and $L_2|K$ are incompatible if and only if 1 lies in the perfect closure of the zero ideal of $L_1 \otimes_K L_2$. The shuffling process (see Subsection 1.2.2) shows that only finitely many elements f_1, \dots, f_n from $L_1 \otimes_K L_2$ are needed to verify that $1 \in \{0\}$. Let $A \subset L_1$ and $B \subset L_2$ be a finite sets such that every f_i can be written in the form $\sum_j a_j \otimes b_j$ with $a_j \in A$ and $b_j \in B$. Then 1 lies in the perfect closure of the zero ideal of $K\langle A \rangle \otimes_K K\langle B \rangle$. So $K\langle A \rangle|K$ and $K\langle B \rangle|K$ are incompatible. \square

²An extension of fields $L|K$ such that K is relatively separably algebraically closed in L is sometimes called *primary*.

Lemma 4.6.10. *Let $M|K$ be an extension of σ -fields and let*

$$K \subset L_1 \subset \cdots \subset L_n$$

be a chain of σ -fields such that $L_i|L_{i-1}$ is universally compatible for $i = 2, \dots, n$. If $M|K$ is compatible with $L_1|K$, then $M|K$ is compatible with $L_n|K$.

Proof. By assumption there exists a σ -field extension $E_1|K$ containing isomorphic copies of $M|K$ and $L_1|K$. Since $L_2|L_1$ is universally compatible there exists a σ -field extension E_2 of L_1 containing isomorphic copies of $E_1|L_1$ and $L_2|L_1$. Inductively, we find that there exists a σ -field extension E_n of L_{n-1} containing isomorphic copies of $E_{n-1}|L_{n-1}$ and $L_n|L_{n-1}$. Then $E_n|K$ contains isomorphic copies of $M|K$ and $L_n|K$. \square

Theorem 4.6.11 (Compatibility theorem). *Let $L_1|K$ and $L_2|K$ be σ -field extensions. Then $L_1|K$ and $L_2|K$ are compatible if and only if $\text{Core}(L_1|K)|K$ and $\text{Core}(L_2|K)|K$ are compatible.*

Proof. Obviously $\text{Core}(L_1|K)|K$ and $\text{Core}(L_2|K)|K$ are compatible if $L_1|K$ and $L_2|K$ are compatible. To prove the converse, it suffices³ to prove the following

Claim: If $L_1|K$ is compatible with $\text{Core}(L_2|K)|K$ then $L_1|K$ and $L_2|K$ are compatible.

So we assume that $L_1|K$ is compatible with $\text{Core}(L_2|K)|K$. As the first step we will show that we can assume without loss of generality that K is inersive: Let $K^* \subset L_1^*$ denote the inersive closures. Since $L_1|K$ and $\text{Core}(L_2|K)|K$ are compatible there exists a σ -field extension M of K containing $L_1|K$ and $\text{Core}(L_2|K)|K$. Then M^* contains $L_1K^*|K^*$ and $\text{Core}(L_2|K)^*|K^*$. But $\text{Core}(L_2|K)^* = \text{Core}(L_2K^*|K^*)$ by Lemma 4.5.6. Therefore $L_1K^*|K^*$ and $\text{Core}(L_2K^*|K^*)|K^*$ are compatible. Thus we can apply the claim to deduce that $L_1K^*|K^*$ and $L_2K^*|K^*$ are compatible. But then also $L_1|K$ and $L_2|K$ are compatible. So from now on we will assume that K is inersive.

As the next step we will show that we can assume without loss of generality that $\text{Core}(L_2|K) = K$. By assumption $L_1|K$ and $\text{Core}(L_2|K)|K$ are compatible. Thus there exists a σ -field extension containing L_1 and $\text{Core}(L_2|K)$ and we can form the compositum $L_1 \text{Core}(L_2|K)$. If the extensions $L_1 \text{Core}(L_2|K)|\text{Core}(L_2|K)$ and $L_2|\text{Core}(L_2|K)$ are compatible, then also the extensions $L_1|K$ and $L_2|K$ are compatible. Note that $\text{Core}(L_2|\text{Core}(L_2|K)) = \text{Core}(L_2|K)$ by Lemma 4.5.4 and that $\text{Core}(L_2|K)$ is inersive by Lemma 4.5.5. So from now on we will assume that $\text{Core}(L_2|K) = K$.

Let $M \subset L_2$ denote the separable algebraic closure of K in L_2 . Then M is a σ -field extension of K . By Lemma 4.6.10 and Lemma 4.6.6 the extensions $L_1|K$ and $L_2|K$ are compatible if and only if the extensions $L_1|K$ and $M|K$ are compatible. Moreover, $\text{Core}(M|K) = \text{Core}(L_2|K) = K$. This shows that we can assume that L_2 is separably algebraic over K .

By Lemma 4.6.9 we can assume without loss of generality that $L_2|K$ is finitely σ -generated.

The next step is to show that we can assume without loss of generality that $L_2|K$ is Galois. Let N denote a normal closure of $L_2|K$. Then $N|K$ is Galois and by Corollary 4.5.13 we can extend σ to N . Moreover $N|K$ is finitely σ -generated.

We can extend $\sigma: L_1 \rightarrow L_1$ to $\sigma: \overline{L_1} \rightarrow \overline{L_1}$. Then there exists a K -embedding $\phi: \text{Core}(N|K) \hookrightarrow \overline{L_1}$. Since $\text{Core}(N|K)|K$ is Galois by Lemma 4.5.16 it follows from Corollary 4.5.11 that $\phi(\text{Core}(N|K)) \subset \overline{L_1}$ is stable under σ .

Since $\text{Core}(N|K) \cap L_2 = \text{Core}(L_2|K) = K$ and $\text{Core}(N|K)|K$ is finite Galois we know from Lemma 4.5.1 (i) that $\text{Core}(N|K)$ and L_2 are linearly disjoint over K . We will now

³To obtain the theorem from the claim apply the claim twice.

change the σ on N without changing σ on L_2 . We define the new σ on $\text{Core}(N|K)$ by transport of structure via ϕ . So $\sigma := \phi^{-1}\sigma\phi$ on $\text{Core}(N|K)$. From now on we consider $\text{Core}(N|K)$ as difference field extension of K via this new endomorphism. Because $\text{Core}(N|K)|K$ and $L_2|K$ are linearly disjoint over K we see that $\text{Core}(N|K)L_2 = \text{Core}(N|K) \otimes_K L_2$ is naturally a σ -field extension of K (Lemma 1.1.11). Since N is the normal closure of $\text{Core}(N|K)L_2$ over K we can extend $\sigma: \text{Core}(N|K)L_2 \rightarrow \text{Core}(N|K)L_2$ to a new $\sigma: N \rightarrow N$ (Corollary 4.5.13). By Lemma 4.5.18 the underlying field of the core of $N|K$ with respect to the old σ agrees with the underlying field of the core of $N|K$ with respect to the new σ . (So the notation $\text{Core}(N|K)$ is non-ambiguous.) The advantage of the new σ on N is that $\phi: \text{Core}(N|K) \rightarrow \overline{L_1}$ is a K - σ -morphism. In other words, $L_1|K$ and $\text{Core}(N|K)|K$ are compatible.

To sum up, we are now in the following situation: K is inversive, $N|K$ is a Galois σ -field extension, finitely σ -generated over K . The σ -field extensions $L_1|K$ and $\text{Core}(N|K)|K$ are compatible. If $L_1|K$ and $N|K$ are compatible, then also $L_1|K$ and $L_2|K$ are compatible (since $L_2 \subset N$). Therefore it suffices to show that $L_1|K$ and $N|K$ are compatible.

In other words, it suffices to prove the claim under the assumptions that K is inversive and that $L_2|K$ is Galois and finitely σ -generated over K . But this are precisely the assumptions needed for Babbitt's decomposition. Therefore the claim follows from Lemma 4.6.10 (using Lemmas 4.6.7 and 4.6.8). \square

Corollary 4.6.12. *Let $L|K$ be a σ -field extension such that $\text{Core}(L|K) = K$. Then $L|K$ is universally compatible.*

Proof. Since clearly $K|K$ is compatible with every σ -field extensions of K , this follows from Theorem 4.6.11. \square

Corollary 4.6.13. *Let K be an inversive σ -field with $\text{Core}(K) = K$. Then any two σ -field extensions of K are compatible. Moreover, if $\sigma_1, \sigma_2: \overline{K} \rightarrow \overline{K}$ are extensions of $\sigma: K \rightarrow K$, then (\overline{K}, σ_1) and (\overline{K}, σ_2) are isomorphic as σ -field extensions of K .*

Proof. Let $L_1|K$ and $L_2|K$ be σ -field extensions. Since the core of $L_1|K$ can be embedded into $\text{Core}(K)$ we see that $\text{Core}(L_1|K) = K$. Therefore $L_1|K$ and $L_2|K$ are compatible by Theorem 4.6.11.

Since (\overline{K}, σ_1) and (\overline{K}, σ_2) are compatible σ -field extensions of K there exists a σ -field M containing K -isomorphic copies of (\overline{K}, σ_1) and (\overline{K}, σ_2) . But inside M the algebraic closure is unique. Therefore (\overline{K}, σ_1) and (\overline{K}, σ_2) are isomorphic as σ -field extensions of K . \square

Note that by Proposition 4.5.22, the above corollary applies to the σ -field $K = \mathbb{C}(z)$, where $\sigma(f(z)) = f(z+1)$.

Exercise 4.6.14. *Let $L|K$ and $M|K$ be σ -field extensions. Assume that one of them is finitely σ -generated. Then there exists an integer $d \geq 1$ such that $L|K$ and $M|K$ are compatible as σ^d -field extensions.*

Chapter 5

Difference kernels

In this chapter we return to the study of systems of algebraic difference equations. The basic idea here is to study a σ -prime ideal \mathfrak{p} in a σ -polynomial ring over a σ -field by looking at the σ -polynomials in \mathfrak{p} up to a given order.

5.1 The difference degree

Let k be a σ -field and $k\{y\} = k\{y_1, \dots, y_n\}$ the σ -polynomial ring in the σ -variables $y = (y_1, \dots, y_n)$ over k . For $d \in \mathbb{N}$ we denote by $k\{y\}[d]$ the k -algebra of σ -polynomials of order not greater than d . That is,

$$k\{y\}[d] := k[y, \sigma(y), \dots, \sigma^d(y)] \subset k\{y\}.$$

We also set $k\{y\}[-1] := k$. For any set $A \subset k\{y\}$ we define

$$A[d] := A \cap k\{y\}[d].$$

For any ring R we denote by $\dim(R)$ the Krull dimension of R .

Theorem 5.1.1. *Let k be a σ -field and \mathfrak{p} a prime σ -ideal of $k\{y\} = k\{y_1, \dots, y_n\}$. For $i \in \mathbb{N}$ set*

$$d_i = \dim(k\{y\}[i]/\mathfrak{p}[i]).$$

Then there exist integers $d, e \in \mathbb{N}$ such that

$$d_i = d(i + 1) + e$$

for $i \gg 0$. Moreover, $d = \sigma\text{-trdeg}(k(\mathfrak{p}^)|k)$.*

Proof. Set $R := k\{y\}/\mathfrak{p}$ and let $a \in R^n$ denote the image of y . Then R is an integral k - σ -algebra. To simplify the notation we set

$$k[i, \dots, j] := k[\sigma^i(a), \sigma^{i+1}(a), \dots, \sigma^j(a)] \subset R$$

and

$$k(i, \dots, j) := k(\sigma^i(a), \sigma^{i+1}(a), \dots, \sigma^j(a)) \subset \text{Quot}(R)$$

for $i, j \in \mathbb{N}$ with $i \leq j$. We also abbreviate $d_i(l) := \dim(k[l, \dots, l+i])$ for $i, l \in \mathbb{N}$. So $d_i = d_i(0)$ for $i \in \mathbb{N}$. The map

$$\sigma: k[\sigma^l(a), \dots, \sigma^{l+i}(a)] \longrightarrow \sigma(k)[\sigma^{l+1}(a), \dots, \sigma^{l+1+i}(a)]$$

is surjective and so

$$\begin{aligned} d_i(l) &= \dim \left(k[\sigma^l(a), \dots, \sigma^{l+i}(a)] \right) \geq \dim \left(\sigma(k)[\sigma^{l+1}(a), \dots, \sigma^{l+1+i}(a)] \right) \geq \\ &\geq \dim \left(k[\sigma^{l+1}(a), \dots, \sigma^{l+1+i}(a)] \right) = d_i(l+1). \end{aligned}$$

Thus, for a fixed $i \in \mathbb{N}$ the sequence $(d_i(l))_{l \in \mathbb{N}}$ is non-increasing and therefore stabilizes at some value $e(i) \in \mathbb{N}$. For $i \in \mathbb{N}$ let $\beta(i) \in \mathbb{N}$ denote the smallest integer such that $d_i(\beta(i)) = d_i(\beta(i) + 1) = \dots = e(i)$ and define recursively a sequence $(\alpha(i))_{i \in \mathbb{N}}$ by $\alpha(0) := \beta(0)$ and

$$\alpha(i) := \max\{\alpha(i-1), \beta(i)\}$$

for $i \geq 1$.

We will next show that the sequence $(\varepsilon(i))_{i \in \mathbb{N}}$ defined by $\varepsilon(i) := e(i+1) - e(i)$ is also non-increasing: For l large enough we have

$$\begin{aligned} \varepsilon(i+1) &= e(i+2) - e(i+1) = \dim(k[l, \dots, l+2+i]) - \dim(k[l, \dots, l+1+i]) = \\ &= \text{trdeg}(k(l, \dots, l+2+i)|k(l, \dots, l+1+i)) \leq \\ &\leq \text{trdeg}(k(l+1, \dots, l+2+i)|k(l+1, \dots, l+1+i)) = \varepsilon(i). \end{aligned}$$

Therefore the sequence $(\varepsilon(i))_{i \in \mathbb{N}}$ eventually stabilizes at some value $d \in \mathbb{N}$. This implies that there exist $i_0 \in \mathbb{N}$ and $c \in \mathbb{Z}$ such that for $i \geq i_0$

$$e(i) = di + c.$$

Next we will show that $\varepsilon(i) = \varepsilon(i+1)$ implies that $\alpha(i+1) = \alpha(i+2)$. Suppose that $\alpha(i+1) < \alpha(i+2)$ and set $l := \alpha(i+2) - 1 \geq \alpha(i+1) \geq \alpha(i)$.

We have two towers of field extensions:

$$\begin{array}{ccc} k(l, \dots, l+2+i) & & k(l+1, \dots, l+3+i) \\ \left| \delta \right. & & \left| \varepsilon(i+1) \right. \\ k(l, \dots, l+1+i) & & k(l+1, \dots, l+2+i) \\ \left| \varepsilon(i) \right. & & \left| \varepsilon(i) \right. \\ k(l, \dots, l+i) & & k(l+1, \dots, l+1+i) \\ & \swarrow e(i) \quad \searrow e(i) & \\ & k & \end{array}$$

As already used above we have $\delta \leq \varepsilon(i)$. So if $\varepsilon(i) = \varepsilon(i+1)$ it follows that

$$\text{trdeg}(k(l, \dots, l+2+i)|k) \leq \text{trdeg}(k(l+1, \dots, l+3+i)|k).$$

But by choice of l , we must have

$$\text{trdeg}(k(l, \dots, l+2+i)|k) = d_{i+2}(l) > d_{i+2}(l+1) = \text{trdeg}(k(l+1, \dots, l+3+i)|k);$$

a contradiction.

Since the sequence $(\varepsilon(i))_{i \in \mathbb{N}}$ eventually stabilizes, this shows that also the sequence $(\alpha(i))_{i \in \mathbb{N}}$ stabilizes at some value $l_0 \in \mathbb{N}$. For $l \geq l_0$ and $i \geq i_0$ we have

$$d_i(l) = e(i) = di + c.$$

For $i \geq l_0$ we have

$$d_i = \text{trdeg}(k(0, \dots, i)|k) = \text{trdeg}(k(0, \dots, i)|k(l_0, \dots, i)) + \text{trdeg}(k(l_0, \dots, i)|k).$$

Since

$$\text{trdeg}(k(0, \dots, i)|k(l_0, \dots, i)) = \text{trdeg}\left(k(a, \sigma(a), \dots) \Big| k(\sigma^{l_0}(a), \sigma^{l_0+1}(a), \dots)\right) =: c'$$

for $i \gg 0$ it follows that

$$d_i = c' + di + c = d(i+1) + e$$

for $i \gg 0$ where $e = c + c' - d$.

It remains to show that $d = \sigma\text{-trdeg}(k(\mathfrak{p}^*)|k)$ and that $e \in \mathbb{N}$. For this we will use the following lemma.

Lemma 5.1.2. *Let k be a σ -field and $R \subset S$ an inclusion of k - σ -algebras such that R is a σ -domain and S an integral domain. Let $a \in S$ be σ -algebraically dependent over R . Then there exists a finite subset B of R and $N \in \mathbb{N}$ such that*

$$\text{trdeg}(k(B, \dots, \sigma^i(B), a, \dots, \sigma^i(a))|k(B, \dots, \sigma^i(B))) \leq N$$

for all $i \in \mathbb{N}$.

Proof. Let K denote the quotient field of R . Note that K is a σ -field. We work inside the quotient field of S (which need not be a σ -field). By assumption, there exists $l \in \mathbb{N}$ and a non-zero polynomial f with coefficients in R such that $f(a, \dots, \sigma^l(a)) = 0$.

Now let $l \in \mathbb{N}$ denote the smallest integer with the property that $f(\sigma^j(a), \dots, \sigma^{j+l}(a)) = 0$ for some $j \in \mathbb{N}$ and some non-zero polynomial f with coefficients in R . Let $B \subset R$ denote the (finite) set of coefficients of f . Let $i \in \mathbb{N}$ and let $\sigma^i f$ denote the polynomial obtained from f by applying σ^i to the coefficients. Then $\sigma^i f(\sigma^{j+i}(a), \dots, \sigma^{j+l+i}(a)) = 0$ and $\sigma^{j+i}(a), \dots, \sigma^{j+l+i-1}(a)$ are algebraically independent over R . This shows that $\sigma^{j+l+i}(a)$ is algebraic over $k(\sigma^i(B), \sigma^{j+i}(a), \dots, \sigma^{j+l+i-1}(a))$ and we find that $k(B, \dots, \sigma^i(B), a, \dots, \sigma^{j+i+l}(a))$ is algebraic over $k(B, \dots, \sigma^i(B), a, \dots, \sigma^{j+l}(a))$. As in the previous sections let us write $\theta_i(B)$ for $B, \dots, \sigma^i(B)$. We have

$$\begin{aligned} \text{trdeg}(k(\theta_i(B), \theta_i(a))|k(\theta_i(B))) &\leq \text{trdeg}(k(\theta_i(B), \theta_{j+i+l}(a))|k(\theta_i(B))) = \\ &= \text{trdeg}(k(\theta_i(B), \theta_{j+i+l}(a))|k(\theta_i(B), \theta_{j+l}(a))) + \text{trdeg}(k(\theta_i(B), \theta_{j+l}(a))|k(\theta_i(B))) = \\ &= \text{trdeg}(k(\theta_i(B), \theta_{j+l}(a))|k(\theta_i(B))) \leq j + l + 1. \end{aligned}$$

□

End of proof of Theorem 5.1.1. As before, let $a = (a_1, \dots, a_n)$ denote the image of y in $R = k\{y\}/\mathfrak{p}$. Let $m \in \mathbb{N}$ denote the largest integer such that among a_1, \dots, a_n there exists m elements which are σ -algebraically independent over k . To simplify the notation we assume that a_1, \dots, a_m are σ -algebraically independent over k . Then for $j = m+1, \dots, n$ the element $a_j \in R$ is σ -algebraically dependent over $k\{a_1, \dots, a_m\}$. It follows from Lemma 5.1.2 that there exists a finite subset B of $k\{a_1, \dots, a_m\}$ and $N \in \mathbb{N}$ such that

$$\text{trdeg}(k(\theta_i(B), \theta_i(a_{m+1}, \dots, a_n))|k(\theta_i(B))) \leq N$$

for $i \in \mathbb{N}$. Enlarging B if necessary, we can assume that $a_1, \dots, a_m \in B$. If $B \subset k[a_1, \dots, a_m, \dots, \sigma^j(a_1), \dots, \sigma^j(a_m)]$ then

$$\text{trdeg}(k(\theta_i(B))|k(\theta_i(a_1, \dots, a_m))) \leq \text{trdeg}(k(\theta_{i+j}(a_1, \dots, a_m))|k(\theta_i(a_1, \dots, a_m))) = jm.$$

So we have

$$\begin{aligned}
d_i &= \text{trdeg}(k(\theta_i(a))|k) = \text{trdeg}(k(\theta_i(a))|k(\theta_i(a_1, \dots, a_m))) + \text{trdeg}(k(\theta_i(a_1, \dots, a_m))|k) \leq \\
&\leq \text{trdeg}(k(\theta_i(B), \theta_i(a_{m+1}, \dots, a_n))|k(\theta_i(a_1, \dots, a_m))) + m(i+1) = \\
&= \text{trdeg}(k(\theta_i(B), \theta_i(a_{m+1}, \dots, a_n))|k(\theta_i(B))) + \text{trdeg}(k(\theta_i(B))|k(\theta_i(a_1, \dots, a_m))) + m(i+1) \\
&\leq N + jm + m(i+1) = m(i+1) + N + jm.
\end{aligned}$$

Since N, j and m are fixed and only i varies, this shows that $d \leq m$.

The kernel of the canonical map $\phi: k\{y\}/\mathfrak{p} \rightarrow k\{y\}/\mathfrak{p}^*$ consists of all $b \in R = k\{y\}/\mathfrak{p}$ such that $\sigma^j(b) = 0$ for some $j \in \mathbb{N}$ (Lemma 1.2.15). But σ is injective on $k\{a_1, \dots, a_m\}$ and so ϕ is injective on $k\{a_1, \dots, a_m\}$. This shows that $m \leq \sigma\text{-trdeg}(k(\mathfrak{p}^*)|k)$ and we find that $d \leq \sigma\text{-trdeg}(k(\mathfrak{p}^*)|k)$.

The reverse estimate is easy: Abbreviate $d' = \text{trdeg}(k(\mathfrak{p}^*)|k)$. Again, to simplify the notation we may assume that the images of $y_1, \dots, y_{d'}$ are a σ -transcendence basis of $k(\mathfrak{p}^*)|k$. Then $a_1, \dots, a_{d'} \in R$ must be σ -algebraically independent over k . Therefore

$$d_i = \text{trdeg}(k(\theta_i(a))|k) \geq \text{trdeg}(k(\theta_i(a_1, \dots, a_{d'}))|k) = d'(i+1).$$

This shows that $d \geq d'$, so $d = d'$, and that $e \in \mathbb{N}$. □


Definition 5.1.3. Let k be a σ -field and $\mathfrak{p} \subset k\{y\} = k\{y_1, \dots, y_n\}$ a prime σ -ideal. Let d and e be as in Theorem 5.1.1. The polynomial $\omega_{\mathfrak{p}}(t) = d(t+1) + e$ is called the dimension polynomial of \mathfrak{p} . The number $\sigma\text{-deg}(\mathfrak{p}) = e$ is called the σ -degree of \mathfrak{p} . If \mathfrak{p} is σ -prime and $X = \mathbb{V}(\mathfrak{p}) \subset \mathbb{A}_k^n$ denotes the corresponding σ -variety, we set $\omega_X(t) = \omega_{\mathfrak{p}}(t)$ and $\sigma\text{-deg}(X) = \sigma\text{-deg}(\mathfrak{p})$.

So if X is an irreducible σ -variety, the dimension polynomial $\omega_X(t)$ of X is of the form

$$\omega_X(t) = \sigma\text{-dim}(X)(t+1) + \sigma\text{-deg}(X).$$

The naming σ -degree is justified by the following analogy with the usual degree. Let k be an algebraically closed field and \mathcal{X} an irreducible d -dimensional affine variety embedded in \mathbb{A}_k^n . The intersection of \mathcal{X} with $n-d$ generic hyperplanes consists of finitely many points. The number of these points is called the degree of \mathcal{X} (or more precisely the degree of \mathcal{X} in \mathbb{A}_k^n). For example, if $\mathcal{X} = \mathbb{V}(f)$ for some irreducible polynomial f , then the degree of \mathcal{X} , is the degree of f .

Now, as we have seen in Section 4.1, the order of a σ -algebraic σ -field extension is analogous to the degree of an algebraic field extension. So the idea is that intersecting X with $n-d$ generic difference hyperplanes gives a σ -variety Y , such that the order of the σ -function field of Y agrees with the σ -degree of X . Note that if $d = \sigma\text{-dim}(X) = 0$, then, by definition, the σ -degree of X equals $\text{ord}(k\langle X \rangle|k)$. This is why the σ -degree is sometimes also called the order.

 **5.1.4.** The σ -degree of X depends on the embedding of X in \mathbb{A}_k^n . (See the exercise below.)

Exercise 5.1.5. Let $X \subset \mathbb{A}_k^n$ and $Y \subset \mathbb{A}_k^m$ be isomorphic σ -varieties. Show that $\sigma\text{-deg}(X)$ and $\sigma\text{-deg}(Y)$ need not be equal.

Corollary 5.1.6. Let k be a σ -field and R a finitely σ -generated k - σ -algebra. Then every ascending chain of prime σ -ideals of R is finite.

Proof. Since R can be written as a quotient of a σ -polynomial ring, it suffices to treat the case that $R = k\{y_1, \dots, y_n\}$. Let $\mathfrak{p}_1 \subset \mathfrak{p}_2 \subset \dots$ be an ascending chain of prime σ -ideals in R . This gives rise to a decreasing sequence of dimension polynomials $\omega_{\mathfrak{p}_1}(t) \geq \omega_{\mathfrak{p}_2}(t) \geq \dots$. Here we write $\omega_{\mathfrak{p}_1}(t) \geq \omega_{\mathfrak{p}_2}(t)$ to signify that $\omega_{\mathfrak{p}_1}(i) \geq \omega_{\mathfrak{p}_2}(i)$ for $i \gg 0$. Such a sequence must be finite. But if $\omega_{\mathfrak{p}_m}(t) = \omega_{\mathfrak{p}_{m+1}}(t)$, then also $\mathfrak{p}_m = \mathfrak{p}_{m+1}$. \square

Exercise 5.1.7. Let k be a σ -field and R a finitely σ -generated k - σ -algebra. Let $\mathfrak{p} \subset R$ be a prime σ -ideal. Show that there exists an $m \in \mathbb{N}$ such that $\mathfrak{p}^* = \{f \in R \mid \sigma^m(f) \in \mathfrak{p}\}$.

5.2 Prolongations and realizations of difference kernels

As above, let k be a σ -field, $k\{y\} = k\{y_1, \dots, y_n\}$ the σ -polynomial ring in the σ -variables y_1, \dots, y_n over k and $d \in \mathbb{N}$. Now $k\{y\}[d]$ is not a σ -ring, but $\sigma: k\{y\} \rightarrow k\{y\}$ induces a morphism of rings

$$\sigma: k\{y\}[d-1] \rightarrow k\{y\}[d].$$

Definition 5.2.1. Let k be a σ -field and $d \in \mathbb{N}$. A prime ideal \mathfrak{q} of $k\{y\}[d]$ is called a σ -kernel of length d in $k\{y\} = k\{y_1, \dots, y_n\}$ if

$$\sigma^{-1}(\mathfrak{q}) = \mathfrak{q}[d-1],$$

where $\sigma: k\{y\}[d-1] \rightarrow k\{y\}[d]$

Example 5.2.2. A difference kernel of length zero is simply a prime ideal of $k[y_1, \dots, y_n]$.

If $\mathfrak{p} \subset k\{y\}$ is a σ -prime ideal, then $\mathfrak{p}[d]$ is a σ -kernel because \mathfrak{p} is reflexive. It is now natural to ask whether every σ -kernel is of this form. If $\mathfrak{q} \subset k\{y\}[d]$ is a σ -kernel, we can consider the perfect closure $\mathfrak{p} := \{\mathfrak{q}\} \subset k\{y\}$ of \mathfrak{q} . But it is unclear if \mathfrak{p} is prime and if $\mathfrak{p}[d] = \mathfrak{q}$. Of course we have $\mathfrak{p}[d] \supset \mathfrak{q}$, but in principle this could be a proper inclusion. In fact, with out current knowledge¹, we might even have $1 \in \mathfrak{p}$.

Definition 5.2.3. Let \mathfrak{q} be a σ -kernel of length d in $k\{y\}$. A σ -prime ideal \mathfrak{p} of $k\{y\}$ is called a realization of \mathfrak{q} if $\mathfrak{q} \subset \mathfrak{p}$. The realization is called regular if $\mathfrak{p}[d] = \mathfrak{q}$.

Let \mathfrak{q} be a σ -kernel of length d in $k\{y\}$. Because $\sigma^{-1}(\mathfrak{q}) = \mathfrak{q}[d-1]$ we have an induced injective map

$$\sigma: k\{y\}[d-1]/\mathfrak{q}[d-1] \longrightarrow k\{y\}[d]/\mathfrak{q}$$

which extends to the quotient fields, i.e., $\sigma: k(\mathfrak{q}[d-1]) \rightarrow k(\mathfrak{q})$. So² if we denote by $a = (a_1, \dots, a_n)$ the image of y in $k(\mathfrak{q})$ then

$$\sigma: k(a, \dots, \sigma^{d-1}(a)) \longrightarrow k(a, \dots, \sigma^d(a)).$$

The following definition is motivated by Theorem 5.1.1.

Definition 5.2.4. Let \mathfrak{q} be a σ -kernel of length d in $k\{y\}$. The σ -dimension of \mathfrak{q} is

$$\sigma\text{-dim}(\mathfrak{q}) = \text{trdeg}(k(\mathfrak{q})|k(\mathfrak{q}[d-1])) = \text{trdeg}(k(a, \dots, \sigma^d(a))|k(a, \dots, \sigma^{d-1}(a))).$$

In other words,

$$\sigma\text{-dim}(\mathfrak{q}) = \dim(k\{y\}[d]/\mathfrak{q}) - \dim(k\{y\}[d-1]/\mathfrak{q}[d-1]).$$

To show that every σ -kernel has a realization we will use the following notion of prolongation of a σ -kernel.

¹We will see soon that this can not happen.

²In the literature often $k(\mathfrak{q})$ with this additional structure is called a difference kernel.

Definition 5.2.5. Let \mathfrak{q} be a σ -kernel of length d in $k\{y\}$. A σ -kernel \mathfrak{q}' of length $d + 1$ in $k\{y\}$ is called a prolongation of \mathfrak{q} if

$$\mathfrak{q}'[d] = \mathfrak{q}.$$

The prolongation is called generic if $\sigma\text{-dim}(\mathfrak{q}) = \sigma\text{-dim}(\mathfrak{q}')$. A regular realization $\mathfrak{p} \subset k\{y\}$ of \mathfrak{q} is called a principal realization if $\mathfrak{p}[i + 1]$ is a generic prolongation of $\mathfrak{p}[i]$ for every $i \geq d$.

We will need to use the following basic fact about base extension: Let $K|k$ be a field extension and let \mathfrak{p} be a prime ideal of $k[y] = k[y_1, \dots, y_n]$. If $\mathfrak{P} \subset K[y]$ is a minimal prime ideal of the ideal of $K[y]$ generated by \mathfrak{p} , then $\mathfrak{P} \cap k[y] = \mathfrak{p}$ and $\dim(K[y]/\mathfrak{P}) = \dim(k[y]/\mathfrak{p})$.

Lemma 5.2.6. Let \mathfrak{q} be a σ -kernel of length d in $k\{y\}$. Then there exists a generic prolongation of \mathfrak{q} .

Proof. Let $\mathfrak{p}_1 \subset k(a, \dots, \sigma^{d-1}(a))[\sigma^d(y)]$ denote the prime ideal consisting of all polynomials in $k(a, \dots, \sigma^{d-1}(a))[\sigma^d(y)]$ such that replacing $\sigma^d(y)$ with $\sigma^d(a)$ gives zero. (So $k(a, \dots, \sigma^{d-1}(a))[\sigma^d(y)]/\mathfrak{p}_1 \simeq k(a, \dots, \sigma^{d-1}(a))[\sigma^d(a)]$.) We have an isomorphism

$$\sigma: k(a, \dots, \sigma^{d-1}(a))[\sigma^d(y)] \rightarrow \sigma(k)(\sigma(a), \dots, \sigma^d(a))[\sigma^{d+1}(y)].$$

So $\sigma(\mathfrak{p}_1)$ is a prime ideal of $\sigma(k)(\sigma(a), \dots, \sigma^d(a))[\sigma^{d+1}(y)]$. Let $\mathfrak{p}_2 \subset k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$ denote a minimal prime ideal above the ideal of $k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$ generated by $\sigma(\mathfrak{p}_1)$. Then $\mathfrak{p}_2 \cap \sigma(k)(\sigma(a), \dots, \sigma^d(a))[\sigma^{d+1}(y)] = \sigma(\mathfrak{p}_1)$. This shows that we have a natural injective map

$$\sigma: k(a, \dots, \sigma^{d-1}(a))[\sigma^d(y)]/\mathfrak{p}_1 \rightarrow k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]/\mathfrak{p}_2.$$

Thus, writing $\sigma^{d+1}(a)$ for the image of $\sigma^d(y)$ in $k(\mathfrak{p}_2)$, we have a natural map

$$\sigma: k(a, \dots, \sigma^d(a)) \longrightarrow k(a, \dots, \sigma^{d+1}(a)).$$

So, if we let $\mathfrak{q}' \subset k[y, \dots, \sigma^{d+1}(y)]$ denote the prime ideal of all polynomials which vanish if we replace $\sigma^i(y)$ with $\sigma^i(a)$ for $i = 0, \dots, d + 1$, then \mathfrak{q}' is a kernel of length $d + 1$ and clearly it is a prolongation of \mathfrak{q} . Since

$$\begin{aligned} \sigma\text{-dim}(\mathfrak{q}) &= \text{trdeg}(k(a, \dots, \sigma^d(a))|k(a, \dots, \sigma^{d-1}(a))) = \dim(k(a, \dots, \sigma^{d-1}(a))[\sigma^d(y)]/\mathfrak{p}_1) \\ &= \dim(\sigma(k)(\sigma(a), \dots, \sigma^d(a))[\sigma^{d+1}(y)]/\sigma(\mathfrak{p}_1)) = \dim(k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]/\mathfrak{p}_2) \\ &= \text{trdeg}(k(a, \dots, \sigma^{d+1}(a))|k(a, \dots, \sigma^d(a))) = \sigma\text{-dim}(\mathfrak{q}') \end{aligned}$$

we see that \mathfrak{q}' is a generic prolongation of \mathfrak{q} . □

Remark 5.2.7. With the notation of the above lemma: The prime ideals of $k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$, which are minimal above the ideal of $k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$ generated by $\sigma(\mathfrak{p}_1)$, are in one-to-one correspondence with the generic prolongations of \mathfrak{q} . If \mathfrak{q}'' is a prolongation of \mathfrak{q} , then there exists a generic prolongation \mathfrak{q}' of \mathfrak{q} with $\mathfrak{q}' \subset \mathfrak{q}''$. Moreover, $\sigma\text{-dim}(\mathfrak{q}'') \leq \sigma\text{-dim}(\mathfrak{q})$ and there are no inclusions between generic prolongations.

Proof. Let \mathfrak{q}'' be a prolongation of \mathfrak{q} . Note that $k(\mathfrak{q}) = k(a, \dots, \sigma^d(a))$ is naturally contained in $k(\mathfrak{q}'')$. Let b denote the image of $\sigma^{d+1}(y)$ in $k(\mathfrak{q}'')$. So $k(\mathfrak{q}'') = k(a, \dots, \sigma^d(a), b)$.

Let $\mathfrak{p}'' \subset k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$ denote the prime ideal of all polynomials that vanish if we substitute b for $\sigma^{d+1}(y)$. Since we have a well-defined map

$$\sigma: k(a, \dots, \sigma^d(a)) \rightarrow k(a, \dots, \sigma^d(a), b)$$

with $\sigma(\sigma^d(a)) = b$, we see that $\sigma(\mathfrak{p}_1)$ is contained in \mathfrak{p}'' . This shows that there exists a prime ideal $\mathfrak{p}_2 \subset k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$, which is minimal above the ideal generated by $\sigma(\mathfrak{p}_1)$ in $k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$ and satisfies $\mathfrak{p}_2 \subset \mathfrak{p}''$. So

$$\begin{aligned} \sigma\text{-dim}(\mathfrak{q}'') &= \dim(k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]/\mathfrak{p}'') \leq \\ &\leq \dim(k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]/\mathfrak{p}_2) = \sigma\text{-dim}(\mathfrak{q}') = \sigma\text{-dim}(\mathfrak{q}), \end{aligned}$$

where $\mathfrak{q}' \subset \mathfrak{q}''$ is defined as in the proof of Lemma 5.2.6. If \mathfrak{q}'' is generic, then $\sigma\text{-dim}(\mathfrak{q}) = \sigma\text{-dim}(\mathfrak{q}'')$. So $\mathfrak{p}'' = \mathfrak{p}_2$.

Assume now that \mathfrak{q}' and \mathfrak{q}'' are generic prolongations of \mathfrak{q} with $\mathfrak{q}' \subset \mathfrak{q}''$. We have to show that $\mathfrak{q}' = \mathfrak{q}''$. Let \mathfrak{p}' and \mathfrak{p}'' denote the prime ideals of $k(a, \dots, \sigma^d(a))[\sigma^{d+1}(y)]$ corresponding to \mathfrak{q}' and \mathfrak{q}'' respectively. Since $\mathfrak{q}' \subset \mathfrak{q}''$ we also have $\mathfrak{p}' \subset \mathfrak{p}''$. But there are no inclusions between prime ideals which are minimal above a given ideal. So $\mathfrak{p}' = \mathfrak{p}''$. This implies $\mathfrak{q}' = \mathfrak{q}''$. \square

Corollary 5.2.8. *Let \mathfrak{q} be a σ -kernel in $k\{y\}$. Then there exists a principal realization of \mathfrak{q} .*

Proof. By Lemma 5.2.6 there exists a chain $\mathfrak{q} \subset \mathfrak{q}_1 \subset \mathfrak{q}_2 \subset \dots$ of generic prolongations. The union is a principal realization of \mathfrak{q} . \square

Exercise 5.2.9. *Let k be a σ -field and $f \in k\{y_1, \dots, y_n\}$ a non-constant σ -polynomial. Show that f has a solution in a σ -field extension of k . (Hint: You can assume that f is irreducible.)*

Lemma 5.2.10. *Let \mathfrak{q} be a σ -kernel of length d in $k\{y\}$ and let \mathfrak{p} be a principal realization of \mathfrak{q} . Then $\sigma\text{-trdeg}(k(\mathfrak{p})|k) = \sigma\text{-dim}(\mathfrak{q})$ and $\sigma\text{-deg}(\mathfrak{p}) = \text{trdeg}(k(\mathfrak{q})|k) - (d+1) \cdot (\sigma\text{-dim}(\mathfrak{q}))$.*

Proof. For $i \geq d$, we have

$$\begin{aligned} \text{trdeg}(k(\mathfrak{p}[i])|k) &= \text{trdeg}(k(\mathfrak{q})) + (i - d) \cdot (\sigma\text{-dim}(\mathfrak{q})) = \\ &= (\sigma\text{-dim}(\mathfrak{q}))(i + 1) + \text{trdeg}(k(\mathfrak{q})|k) - (d + 1) \cdot (\sigma\text{-dim}(\mathfrak{q})). \end{aligned}$$

Thus the claim follows from Theorem 5.1.1. \square

Lemma 5.2.11. *Let \mathfrak{q} be a σ -kernel in $k\{y\}$. If \mathfrak{p} is a principal realization of \mathfrak{q} , then \mathfrak{p} is a minimal prime ideal above $\{\mathfrak{q}\}$.*

Proof. Since $\mathfrak{q} \subset \mathfrak{p}$ we have $\{\mathfrak{q}\} \subset \mathfrak{p}$. Thus there exists a prime ideal $\{\mathfrak{q}\} \subset \mathfrak{p}' \subset \mathfrak{p}$ which is minimal above $\{\mathfrak{q}\}$. (Of course \mathfrak{p}' is σ -prime.) Assume that \mathfrak{q} has length d . Since $\mathfrak{p}'[d] \subset \mathfrak{p}[d] = \mathfrak{q}$, we see that $\mathfrak{p}'[d+1]$ is a prolongation of \mathfrak{q} . By Remark 5.2.7 there exists a generic prolongation \mathfrak{q}'' of \mathfrak{q} with $\mathfrak{q}'' \subset \mathfrak{p}'[d+1] \subset \mathfrak{p}[d+1]$. As there are no inclusions between generic prolongations we must have $\mathfrak{q}'' = \mathfrak{p}[d+1]$ and $\mathfrak{p}'[d+1] = \mathfrak{p}[d+1]$. Inductively we find $\mathfrak{p}'[i] = \mathfrak{p}[i]$ for $i \geq d$. So $\mathfrak{p} = \mathfrak{p}'$ is minimal above $\{\mathfrak{q}\}$. \square

Example 5.2.12. *A prime component of $\{\mathfrak{q}\}$ need not be a principal realization of \mathfrak{q} . In fact, a prime component of \mathfrak{q} need not even be a regular realization of \mathfrak{q} . (See Example 5.2.15.)*

One can however show that for every regular realization \mathfrak{p} of \mathfrak{q} , there exists a principal realization \mathfrak{p}' of \mathfrak{q} with $\mathfrak{p}' \subset \mathfrak{p}$.

Corollary 5.2.13. *Let \mathfrak{q} be a σ -kernel in $k\{y\}$. Then \mathfrak{q} has only finitely many principal realizations.*

Proof. This is clear from Lemma 5.2.11, since there are only finitely many prime ideals minimal above $\{\mathfrak{q}\}$ (Corollary 3.4.3). \square

5.2.1 Application to difference varieties

Recall that the effective order of a σ -polynomial was defined in Subsection 1.1.4.

Theorem 5.2.14. *Let k be a σ -field and $f \in k\{y_1, \dots, y_n\}$, $f \notin k$ an irreducible σ -polynomial such that $\text{Eord}(f) = \text{ord}(f)$. Then $\mathbb{V}(f) \subset \mathbb{A}_k^n$ has an irreducible component X such that $\sigma\text{-dim}(X) = n - 1$ and $\sigma\text{-deg}(X) = \text{ord}(f)$.*

Proof. Let $d \in \mathbb{N}$ denote the order of f , and set $\mathfrak{q} = (f) \subset k[y, \dots, \sigma^d(y)]$. We claim that \mathfrak{q} is a σ -kernel. Since f is irreducible, \mathfrak{q} is a prime ideal. Clearly, $\mathfrak{q}[d-1] = \{0\}$. So we have to show that $\sigma^{-1}(\mathfrak{q}) = \{0\}$. Let $g \in k[y, \dots, \sigma^{d-1}(y)]$, with $\sigma(g) \in \mathfrak{q}$. We have to show that $g = 0$. Suppose $g \neq 0$. Then $\text{Eord}(\sigma(g)) \leq d - 1$. But since $\text{Eord}(f) = d$, every non-zero element of $\mathfrak{q} = (f)$ has effective order d . So $\text{Eord}(\sigma(g)) = d$; a contradiction. Thus \mathfrak{q} is a σ -kernel. By Corollary 5.2.8 there exists a principal realization $\mathfrak{p} \subset k\{y\}$ of \mathfrak{q} . We know from Lemma 5.2.11 that \mathfrak{p} is a prime component of $\{\mathfrak{q}\} = \{f\}$. Therefore $X := \mathbb{V}(\mathfrak{p})$ is an irreducible component of $\mathbb{V}(f)$. We have $\sigma\text{-dim}(\mathfrak{q}) = n - 1$. So it follows from Lemma 5.2.10 that $\sigma\text{-dim}(X) = \sigma\text{-trdeg}(k(\mathfrak{p})|k) = \sigma\text{-dim}(\mathfrak{q}) = n - 1$ and

$$\begin{aligned} \sigma\text{-deg}(X) &= \sigma\text{-deg}(\mathfrak{p}) = \text{trdeg}(k(\mathfrak{q})|k) - (d+1) \cdot (\sigma\text{-dim}(\mathfrak{q})) = \\ &= (d+1)n - 1 - (d+1)(n-1) = d = \text{ord}(f). \end{aligned}$$

\square

Example 5.2.15. The claim of the above theorem does not hold for all irreducible components. Let k be a σ -field and let y denote a σ -variable over k . The σ -polynomial $f = y\sigma^2(y) + \sigma(y) \in k\{y\}$ is irreducible and satisfies $\text{ord}(f) = \text{Eord}(f) = 2$. Clearly $0 \in k$ is a solution of f . We claim that $X := \mathbb{V}(y) \subset \mathbb{V}(f)$ is an irreducible component of $\mathbb{V}(f)$. Let $\mathfrak{p} \subset k\{y\}$ denote a σ -prime σ -ideal which is a prime component of $\{f\}$ such that $\{f\} \subset \mathfrak{p} \subset \{y\}$. We have

$$y\sigma(f) - f = y(\sigma(y)\sigma^3(y) + \sigma^2(y)) - y\sigma^2(y) - \sigma(y) = \sigma(y)(y\sigma^3(y) - 1).$$

Thus, either $\sigma(y) \in \mathfrak{p}$ or $y\sigma^3(y) - 1 \in \mathfrak{p}$. If $y\sigma^3(y) - 1 \in \mathfrak{p} \subset \{y\}$ then $1 \in \{y\}$; a contradiction. Therefore $\sigma(y) \in \mathfrak{p}$ and it follows that $\mathfrak{p} = \{y\}$ is a prime component of $\{f\}$. I.e., X is an irreducible component of $\mathbb{V}(f)$. Clearly $\sigma\text{-deg}(X) = 0 \neq 2 = \text{ord}(f)$.

Nomenclature

- $[\sigma]_k R$ k - σ -algebra associated with the k -algebra R , page 14
- $[F]$ difference ideal generated by F , page 9
- $[L : K]$ degree of the field extension $L|K$, page 51
- \mathbb{A}_k^n difference affine (n -)space over k , page 27
- $\text{Core}(K)$ core of the σ -field K , page 60
- $\text{Core}(L|K)$ the core of the σ -field extension $L|K$, page 57
- $\dim(R)$ Krull dimension of the ring R , page 69
- $\text{Eord}(f)$ effective order of the σ -polynomial f , page 11
- ϕ^* dual morphism of $\phi: R \rightarrow S$, page 24
- $\mathbb{I}(X)$ vanishing ideal of the difference variety X , page 27
- $\mathfrak{a} : S$ saturation of the ideal \mathfrak{a} with respect to the multiplicatively closed set S , page 17
- \mathfrak{a} difference ideal, page 9
- \mathfrak{a}^* reflexive closure of the difference ideal \mathfrak{a} , page 18
- $\text{Im}(f)$ image of morphism f , page 32
- $\text{ld}(L|K)$ limit degree of the σ -field extension $L|K$, page 52
- \leq a ranking, page 35
- \mathbb{N} natural numbers (containing 0), page 6
- $\mathbb{N}_{\geq 1}$ natural numbers without 0, page 6
- $\omega_{\mathfrak{p}}(t)$ dimension polynomial of the prime σ -ideal \mathfrak{p} , page 72
- $\omega_X(t)$ dimension polynomial of the irreducible σ -variety X , page 72
- $\text{ord}(f)$ order of the σ -polynomial f , page 11
- $\text{ord}(L|K)$ order of the σ -field extension $L|K$, page 44
- \overline{K} algebraic closure of the field K , page 58
- $\mathfrak{p}, \mathfrak{q}$ prime ideals, prime σ -ideals or σ -prime ideals, page 16
- $\text{Quot}(R)$ quotient field of the integral domain R , page 6

$\text{rk}(f)$ rank of the σ -polynomial f , page 36
 σ endomorphism, page 7
 $\sigma^{-1}(N)$ inverse image of N under σ , page 6
 $\sigma\text{-deg}(\mathfrak{p})$ σ -degree of the prime σ -ideal \mathfrak{p} , page 72
 $\sigma\text{-deg}(X)$ σ -degree of the irreducible σ -variety X , page 72
 $\sigma\text{-dim}(\mathfrak{q})$ σ -dimension of the σ -kernel \mathfrak{q} , page 73
 $\sigma\text{-dim}(X)$ σ -dimension of the σ -variety X , page 50
 Seq_S ring of sequences, page 8
 $\sigma\text{-field}_k$ category of σ -field extensions of the σ -field k , page 26
 $\sqrt{\mathfrak{a}}$ radical of the ideal \mathfrak{a} , page 19
 $\text{Spec}^\sigma(R)$ difference spectrum of the σ -ring R , page 22
 $\sigma\text{-trdeg}(L|K)$ σ -transcendence degree of the σ -field extension $L|K$, page 49
 Θy the set consisting of all $\sigma^i(y_j)$, page 35
 $\theta_i(A)$ $A, \sigma(A), \dots, \sigma^i(A)$, page 54
 $\text{trdeg}(L|K)$ transcendence degree of the field extension $L|K$, page 44
 $\mathbb{V}(F)$ difference variety defined by F , page 26
 $\mathbb{V}_K(F)$ solutions of F in K , page 26
 $\mathcal{V}(F)$ closed subset of the difference spectrum defined by F , page 23
 $\{F\}$ perfect difference ideal generated by F , page 18
 $\{F\}_{wm}$ smallest radical well-mixed σ -ideal containing F , page 21
 $\sigma^i f$ polynomial obtained from the polynomial f by applying σ^i to the coefficients, page 57
 $A[d]$ difference polynomials in A of order not greater than d , page 69
 $D^\sigma(f)$ basic open subset, page 23
 $f(X)$ functorial image of the morphism f , page 32
 f^* morphism dual to f , page 29
 $F^{\{i\}}$ i -th step in the shuffling procedure, page 19
 FG set of all products of elements from F and G , page 6
 I_F multiplicatively closed σ -stable set generated by all the initials of elements from F , page 37
 I_f initial of the σ -polynomial f , page 36

$k(\mathfrak{p})$ residue field of the prime ideal \mathfrak{p} , page 6
 k, K, L, M difference fields, page 7
 $K\langle A \rangle$ σ -field extension of K generated by A , page 11
 $k\langle X \rangle$ σ -function field of the irreducible σ -variety X , page 50
 $k\{X\}$ σ -coordinate ring of the σ -variety X , page 27
 R, S difference rings, page 7
 $R\{1/f\}$ Difference localization at f , page 10
 $R\{A\}$ R - σ -algebra generated by A , page 11
 $R\{y\}$ σ -polynomial ring over R in the σ -variables $y = (y_1, \dots, y_n)$, page 10
 R^* inversive closure of the σ -ring R , page 12
 $R_{\mathfrak{p}}$ localization at the prime ideal \mathfrak{p} , page 18
 R_i $R \otimes_k {}^{\sigma}R \otimes_k \cdots \otimes_k {}^{\sigma^i}R$, page 14
 $S^{\#}$ k -algebra obtained from the k - σ -algebra S by forgetting σ , page 15
 u_f leader of the σ -polynomial f , page 35
 X, Y σ -varieties, page 26
 ${}^{\sigma^i}R$ extension of scalars of the k -algebra R via $\sigma^i: k \rightarrow k$, page 14

Index

- σ -algebraic, 44
- σ -algebraically dependent, 47
- σ -algebraically dependent on A , 47
- σ -algebraically independent, 47
- σ -coordinate ring, 27
- σ -degree, 72
- σ -dimension of a kernel, 73
- σ -domain, 16
- σ -field extension, 7
- σ -function field, 50
- σ -ideal, 9
- σ -kernel, 73
- σ -maximal, 16
- σ -polynomial map, 28
- σ -polynomial ring, 10
- σ -prime, 16
- σ -radicial, 62
- σ -stable, 6
- σ -subring, 7
- σ -subvariety, 27
- σ -transcendence basis, 48
- σ -transcendence degree, 49
- σ -transcendental, 44
- σ -unit, 20
- σ -variety, 26
- σ -dimension, 50
- m -basis, 40

- autoreduced, 36

- basis, 40
- benign, 60

- characteristic set, 37
- closed embedding, 32
- Cohn topology, 23
- compatible, 65
- core, 60

- dimension polynomial, 72
- dominant, 32

- effective order, 11

- finitely σ -presented, 39

- generic point, 23
- generic prolongation of a σ -kernel, 74

- incompatible, 65
- initial, 36
- inversive σ -ring, 11
- inversive closure, 12
- irreducible, 23
- irreducible components, 42
- irreducible components of X , 43

- leader, 35
- limit degree, 52

- morphism of k - σ -varieties, 28

- order, 11, 44

- perfect, 16
- perfectly reduced, 30
- prime components, 43
- prolongation of a σ -kernel, 74

- ranking, 35
- realization of a σ -kernel, 73
- reflexive, 16
- reflexive closure, 18
- regular realization of a σ -kernel, 73
- Ritt difference ring, 41

- shuffling procedure, 19

- topological space of a σ -variety, 31

- universally compatible, 65

- well-mixed, 16