The Elementary Theory of the Frobenius Automorphisms

Ehud Hrushovski [*]

February 5, 2004

Abstract

We lay down some elements of a geometry based on difference equations. Various constructions of algebraic geometry are shown to have meaningful analogs: dimensions, blowing-up, moving lemmas.

Analogy aside, the geometry of difference equations has two quite different functorial connections with ordinary algebraic geometry. On the one hand, a difference scheme is determined by geometric data, including principally a pro-algebraic scheme. On the other hand, for each prime power p^m, one has a functor into algebraic schemes over $\mathbb{F}_p$, where the structure endomorphism becomes Frobenius.

Transformal zero-cycles have a rich structure in the new geometry. In particular, the Frobenius reduction functors show that they encapsulate data described in classical cases by zeta or L-functions. A theory of rational and algebraic equivalence of 0-cycles is initiated, via a study of the transformal analog of discrete valuation rings.

The central application and motivation is the determination of the elementary theory of the class of Frobenius difference fields (algebraically closed fields of characteristic $p > 0$, enriched with a symbol for $x \mapsto x^{p^m}$. It coincides with the model companion ACFA of the theory of difference fields. In particular, a sentence P holds in $(\mathbb{F}_p^a, x \mapsto x^p)$ for almost all primes p if and only if it is true for a generic automorphism of a field of characteristic 0; i.e. true in (L, σ) for a co-meager set of $\sigma \in Aut(L)$, where L is the field of algebraic functions in denumerably many variables over $\mathbb{Q}$.

The proof requires a twisted version of the Lang-Weil estimates, related to Deligne's conjecture but less precise and valid more generally. It is proved on the basis of the preceding work on difference geometry.

Some applications are given, in particular to finite simple groups, and to the Jacobi bound for difference equations.

[*] Department of Mathematics, Hebrew University, Jerusalem, Israel; ehud@math.huji.ac.il

Contents

1 Introduction

1.1 The main theorem

A classical result of algebraic model theory ([Ax],[Fried-Jarden86]) shows that all elementary statements about finite fields are determined by three facts from geometry. The fundamental fact is Weil's Riemann Hypothesis for curves, entering via the Lang-Weil estimates [Lang-Weil]. The auxiliary statements are the Cebotarev density theorem, and the cyclic nature of the Galois group. All three facts are proved, within algebraic geometry, number theory and Galois theory respectively, by viewing a finite field as the fixed field of the Frobenius automorphism. One might therefore hope that the pertinent geometry could be used directly to derive the full elementary theory of the Frobenius maps. The elementary theory of finite fields would follow as a special case.

This hope turns out to be correct, except that the Lang-Weil estimates do not suffice, and need to be replaced by a more general principle. Let k be an algebraically closed field. If V is a variety over k and σ is an automorphism of k, we denote by V^σ the variety obtained from V by applying σ to the defining parameters. (In terms of schemes, if $f : V \to Spec(K)$ is a scheme over $Spec(K)$, then V^σ is the same scheme, but with f replaced by $f \circ \sigma^{-1}$.) ϕ_q denotes the Frobenius automorphism $x \mapsto x^q$. $\Phi_q \subset (X \times X^{\phi_q})$ denotes the graph of ϕ_q, viewed as a subvariety.

Theorem 1.1 *Let X be an affine variety over k, and let $S \subset (X \times X^{\phi_q})$ be an irreducible subvariety. Assume $\dim(S) = \dim(X) = d$, the maps $S \to X$, $S \to X'$ are dominant, and one is quasi-finite. Let $a = [S : X]/[S : X']_{insep}$. Then*

$$|S(k) \cap \Phi_q(k)| = aq^d + O(q^{d-\frac{1}{2}})$$

The numbers $[S : X]$,$[S : X']_{insep}$ refer to the degree (respectively purely inseparable degree) of the field extensions $k(S)/k(X)$, $k(S)/k(X')$. In particular if S is the graph of a separable morphism $X \to X$, we have $a = 1$.

The meaning of the O can be explained as follows. If $X \subset \mathbb{A}^m$, let $\bar{S}$ be the closure of $S \subset (\mathbb{A}^m \times \mathbb{A}^m) \subset \mathbb{P}^m \times \mathbb{P}^m$; then the O depends only on m and the bidegrees of $\bar{S}$. In particular, if q is large compared to these degrees, the theorem implies that $S(k) \cap \Phi_q(k) \neq \emptyset$.

We will give other versions of the statement, better suited to the inherent uniformity of the situation. Below (1.1A) we state it in the language of difference algebra. A algebro-geometric version will be given later (§11.4.)

Observe that X need not be defined over the fixed field of ϕ_q (or indeed over any finite field.) If X is defined over $GF(q)$, then $X = X^{\phi_q}$, and the diagonal (or the graph any other dominant morphism $X \to X$) becomes a possible choice of S. In the case of the diagonal, one obtains the Lang-Weil estimates.

When X and S descend to a finite field, the projection $S \to X$ is proper, and $S \to X'$ is quasi-finite, Theorem 1.1 follows from Deligne's conjecture ([Fujiwara97],[Pink92]) together with his theorem on eigenvalues of Frobenius. In general these last two assumptions ([Pink92] 7.1.1) cannot be simultaneously obtained in our context, as far as I can see.

Since X is affine, an arbitrary proper divisor can be removed, preserving the same estimate. We thus find that the set of points of $S \cap (X \times X^{\phi_q})$ is "asymptotically Zariski dense". Let us state a special case of this (when S is a morphism, it was used by Borisov and Sapir for a group theoretic application.)

Corollary 1.2 *Let X be an affine variety over $k = GF(q)$, and let $S \subset X^2$ be an irreducible subvariety over K^a. Assume the two projections $S \to X$ are dominant. Then for any proper subvariety W of X, for large enough m, there exist $x \in X(k^a)$ with $(x, \phi_q^m(x)) \in S$ and $x \notin W$.*

(k^a denotes the algebraic closure of k.)

Difference algebraic statement Here is the same theorem in the language of difference algebra, that will be used in most of this paper, followed by a completely algebraic corollary; we do not know any purely algebraic proof of this corollary.

A *difference field* is a field K with a distinguished endomorphism σ. K is *inversive* if σ is an automorphism. A subring of K, closed under σ, is called a *difference domain.* A difference ring R is *simple* if every difference ring homomorphism is injective or zero on R.

If q is a power of a prime p, let K_q be the difference field consisting of an algebraically closed field of characteristic p, together with the q-Frobenius automorphism $\phi_q(x) = x^q$. K_q will be called a *Frobenius difference field.*

Theorem 1.1A *Let $D \subset R$ be finitely generated difference domains. Assume there exists an embedding of D into an algebraically closed, inversive difference field L, such that $R \otimes_D L$ is a difference domain, of transcendence degree d over L. Then there exist $b, b' \in \mathbb{N}$, $0 < c \in \mathbb{Q}$, and $d \in D$ such that for any prime power $q \geq b$, and any homomorphism of difference rings $h : D \to K_q$ with $h(d) \neq 0$, h extends to a homomorphism $\bar{h} : R \to K_q$. Moreover, the number of different $\bar{h}$ is $cq^d + e$, where $e \leq b' q^{d-1/2}$.*

Corollary 1.3 *Let D be a finitely generated difference domain. Assume D is simple. Then D is a finite field, with a power of Frobenius.*

The assumption that D is a domain can be weakened to the condition: $ab = 0$ implies $a\sigma(b) = 0$. A difference ring satisfying this condition will be called *well-mixed.*

Model theoretic consequences With this in hand, the methods of Ax easily generalize to yield the first order theory of the Frobenius difference fields. Let T_∞ be the set of all first-order sentences θ such that for all sufficiently large q, $K_q \models \theta$.

Theorem 1.4 *T_∞ is decidable. It coincides with the model companion "ACFA" of difference fields.*

Here is a more precise presentation of the theorem (restricted to primes), that does not explicitly mention the axioms of ACFA. Fix a countable algebraically closed field L of infinite transcendence degree over $\mathbb{Q}$. Let $G = Aut(L)$; for $\sigma \in Aut(\mathbb{Q}^a)$, let $G_\sigma = \{\tau \in G : \tau|\mathbb{Q}^a = \sigma\}$. G_σ is a Polish space, and it makes sense to talk of meager and co-meager sets; we will mean "almost every" in this sense.

Let ϕ be a first order sentence in the language of difference rings. The following numbers are equal:

1) The Dirichlet density of the set of primes p such that $K_p \models \phi$.

2) The Haar measure of the set of $\sigma \in Aut(\mathbb{Q}^a)$, such that, for almost every $\tau \in G_\sigma$, $(L, \tau) \models \phi$.

The theory ACFA is described and studied in [Chatzidakis-Hrushovski]. The axioms state that the field is algebraically closed, and:

Let V be an absolutely irreducible variety over K, and let S be an irreducible subvariety of $V \times V^\sigma$ projecting dominantly onto V and onto V^σ. Then there exists $a \in V(K)$ with $(a, \sigma(a)) \in S$.

With different axioms, the existence of a model companion for difference fields was discovered earlier by Angus Macintyre and Lou Van den Dries, precisely in connection with the search for the theory of the Frobenius. See [Macintyre95]. In positive characteristic, we can also conclude:

Corollary 1.5 *Let $F = GF(p)^{\mathrm{alg}}$. For almost every automorphism σ of F (in the sense of Baire category,) $(F, \sigma) \models ACFA$.*

I.e. the set of exceptional automorphisms is meager.

A similar result holds, much more trivially, if $GF(p)^{\mathrm{alg}}$ is replaced by an algebraically closed field of infinitely countable transcendence degree over the prime field. However, an example of Cherlin and Jarden shows that a generic automorphism of $\bar{\mathbb{Q}}$ does not yield a model of ACFA, nor is this true for any other field of finite transcendence degree.

We can also consider the theory $T_{\mathrm{Frobenius}}$ consisting of those sentences true in *all* K_q, or in all K_p with p prime. The completions of this theory are just the completions of $ACFA$ together with the complete theories of the individual K_q. With some extra care, we can also conclude:

Theorem 1.6 *The theory of all K_p, p prime, is decidable.*

A similar result holds for the theory of all K_q, q an arbitrary prime power; or with q ranging over all powers of a given prime; etc.

1.2 Some applications

Finite simple groups The theory of definable groups has been worked out for $ACFA$, and has some suggestive corollaries for finite simple groups. This work has not yet been published; but an important special case (groups over the fixed field) contains the main ideas, see [Hrushovski-Pillay]. Here we give a preview.

We say that a family of finite groups indexed by prime powers q is *uniformly definable* if there exist first order formulas ϕ, ψ such that ϕ defines a finite subset of each K_q, ψ defines a group operation on ϕ, and the family consists of these groups for the various K_q. Examples of such families include $G_n(q)$, where n is fixed, and $G_n(q)$ is one of the families of finite simple groups (e.g. $PSL_n(q)$). All but two of these families are already definable over finite fields. However, the Ree and Suzuki families are not. Instead they are defined by the formula:

$$\sigma(x) = \phi(x), x \in G$$

where G is a certain algebraic group, and ϕ is an algebraic automorphism whose square is the Frobenius ϕ_2 or ϕ_3.

From Theorem 1.4, we obtain immediately:

Theorem 1.7 *For each fixed n, the first order theory of each of the classes $G_n(q)$ of finite simple groups is decidable.*

The remaining results use the theory of groups of finite S1-rank.

Theorem 1.8 *Every uniformly definable family of finite simple groups is (up to a finite set) contained in a finite union of families $G_n(q)$*

More precisely, for each q, there is a definable isomorphism in K_q between the group $\phi(K_q)$, and some $G_n(q)$; only finitely many values of n occur, and finitely many definable isomorphisms. The theorem is proved without using the classification of the finite simple groups. Note that it puts the Ree and Suzuki groups into a natural general context. (See also [Bombieri] in this connection.)

Theorem 1.9 *Let $G_n(q)$ be one of the families of finite simple groups. Then there exists an integer $r = r(n)$ such that for any q, every nontrivial conjugacy class of $G_n(q)$ generates $G_n(q)$ in at most r steps.*

This is a special case of more general results on generation of subgroups by definable subsets. One can indeed take $r(n) = 2\dim H_n$, where H_n is the associated algebraic group.

Jacobi's bound for difference equations A linear difference or differential equation of order h, it is well-known, has a solution set whose dimension is at most h. The same is true for a nonlinear equation, once the appropriate definitions of dimension are made. Jacobi, in [Jacobi], proposed a generalization to systems of n differential equations in n variables. The statement is still conjectural today; an analogous conjecture for difference equations was formulated by Cohn. We prove this in §15. Our method of proof illustrates theorem 1.1: we show that our analog of the Lefschetz principle translates Cohn's conjecture to a very strong, but known, form of Bezout's theorem.

See §14 for proofs and some other applications.

1.3 Difference algebraic geometry

Take a system X of difference equations: to fix ideas, take the equations of the unitary group, $\sigma(u)u^t = 1$, where u is an $n \times n$ -matrix variable; or in one variable, take $\sigma(x) = x^2 + c$, for some constant c. Consider four approaches to the study of X.

It may be viewed as part of the category of objects defined by difference equations; i.e. part of an autonomous geometry of difference equations. See the discussion and examples in [Hrushovski01]. Here X is identified, in the first instance, with the set of solutions of the defining equations in a difference field. The rough geography of the category of definable sets and maps is worked out in [Chatzidakis-Hrushovski], [Chatzidakis-Hrushovski-Peterzil]. A great deal remains to be done, for example with respect to topology. The definition of difference schemes in Part I can be viewed as a small step in a related direction.

Alternatively, it is possible to view a difference equation in terms of the algebro-geometry data defining it. One can often represent the system as a correspondence: a pair (V, S), where V is an algebraic variety over a difference field k, S a subvariety of $V \times V^{\sigma}$; the equation is meant to be $X = \{x : (x, \sigma(x)) \in S\}$. See 6.2. The great strength of this approach is that results of algebraic geometry become directly accessible, especially when V is be smooth and complete, and the fibers of $S \to V, S \to V^{\sigma}$ of equal dimensions. It is generally not possible to meet these desiderata. Moreover, given a system (V, S) in this form, simple difference-theoretic questions (such as irreducibility) are not geometrically evident; and simple operations require radical changes in V. Nevertheless, we are often obliged to take this approach, in order to use techniques that are not available directly for difference varieties or schemes; especially the cohomological results used in §12. We will in fact stay with this formalism for as long as possible, particularly in §11, even when a shorter difference-theoretic treatment is visible; cf. the proof of the moving lemmas in §11.1. One obtains more explicit statements, and a contrast with those points that seem to really require a difference-theoretic treatment.

The third and most classical approach is that of dynamics. See Gromov's [Gromov] for a combination of dynamics with algebraic geometry in a wider context, replacing our single endomorphism with a finitely generated group; it also contains, in a different way than ours, reductions to finite objects. Here the space corresponding to (V, S) is

$$Y = \{(a_0, a_1, \ldots) : (a_n, a_{n+1}) \in S\}$$

This is a pro-algebraic variety; we can topologize it by the coarsest topology making the projections to algebraic varieties continuous, if the latter are given the discrete topology. Note that the points of X and Y are quite different; in particular it makes sense to talk of points of Y over a field (rather than a difference field.)

For the time being the work described here barely touches upon dynamical issues; though just under the surface, the proofs in [Chatzidakis-Hrushovski-Peterzil] and in §12, the definition of transformal degree, and other points do strongly involve iteration. It is also true that by showing density of Frobenius difference fields among all difference fields, we show density of periodic points (at least if one is willing to change characteristics); raising hopes for future contact.

A fourth point of view, introduced in §10, is to view the symbol σ as standing for a variable Frobenius. For each prime power $q = p^m$ we obtain an algebraic variety (or scheme) X_q, by interpreting σ as the Frobenius $x \mapsto x^q$ over a field of characteristic p. The main theorem can be read as saying that this approach is equivalent to the others.

It can be quite curious to see information, hidden in (V, S), released by the two apparently disparate processes of iteration, and intersection with Frobenius. Even simple dimension-theoretic facts can sometimes be seen more readily in this way. Cf. 16.1, or the proof of the Jacobi bound for difference equations in §15 (where I know no other way.)

Here is a the simplest example of a definition and a number from the various points of view. The following are essentially equivalent:

- X has transformal dimension 0, or equivalently finite SU-rank in the sense of [Chatzidakis-Hrushovski], or finite total order (§4);

- The map $S \to V$ is quasi-finite;
- The space Y is locally compact;
- Almost all X_q are finite.

If this is the case, and X is irreducible, the following numbers are essentially equal:

- the degree of the correspondence S;
- the entropy of Y;
- the number c such that X_q has asymptotic size cd^q, where $d = \dim(V)$.

The associated invariant of X needs to be defined in terms of difference schemes, rather than varieties, and will be described in §5.3.

"Essentially" indicates that these equivalences are not tautologies, and require certain conditions; for instance it suffices for $S \to V$ to be generically quasi-finite, if the exceptional infinite fibers do not meet X. The various technicalities will be taken up in §2- 6.

1.4 Description of the paper

§2 - 6 introduce a theory of difference schemes. As in algebraic geometry, they are obtained by gluing together basic schemes of the form $\mathrm{Spec}^\sigma A$, with A a difference ring. The $\mathrm{Spec}^\sigma A$ are more general than difference varieties in that A may have nilpotents, both in the usual sense, and elements a such that $\sigma(a) = 0$ while $a \neq 0$.

§2 develops the basic properties of a well-mixed difference ring. Any difference ring without zero-divisors, as well as any difference ring whose structure endomorphism is Frobenius, is well-mixed. Well-mixed rings admit a reasonable theory of localization; in particular, Proposition 3.8 shows that a difference scheme based on well-mixed rings makes sense intrinsically, without including the generating difference rings explicitly in the data.

In §4, the dimension and degree of a difference scheme is defined. In fact there are two dimensions: transformal dimension, and when that vanishes, total dimension. The transformal affine line has transformal dimension one, and contains many difference subschemes of transformal dimension zero and finite total dimension, notably the "short" affine line over the fixed field, of total dimension one. Here too, for well-mixed schemes we obtain a smoother theory; cf. 4.3, 4.30. We improve a theorem of Ritt and Raudenbush by proving Noetherianity of finitely generated difference rings with respect to radical, well-mixed ideals. We do not know whether the same theorem holds for all well-mixed ideals. (Ritt-Raudenbush proved it for perfect ideals; cf. [Cohn].)

§5 is concerned with the kernel of the endomorphism σ on the local rings, or with transformal nilpotents; it is necessary to show how their effect on total dimension is distributed across a difference scheme. §6 deals with the presentation of a difference scheme as a correspondence (V, S), and how in simple case one may test for e.g. irreduciblity.

Further basic material, not used in the main line of the paper, is relegated to §16. In §16.1, the transformal dimension of the difference scheme corresponding to (V, S) is determined in general; I do not know a proof that does not go through the main theorem of the paper. §16.2 looks at transformal projective space. §16.3 introduces two transformal versions of the blowing-up construction. They are shown to essentially coincide; as one (based on the short line, over the fixed field) has no Frobenius analog, this may be geometrically interesting.

The idea of a moving family of varieties, and of a limit of such a family, is characteristic of geometry. It is poorly understood model-theoretically, and has not so far been studied in difference algebraic geometry. In algebraic geometry, it is possible to give a completely synthetic treatment: given a family V_t of varieties over the affine t-line (except over 0), take the closure and let the limit be the fiber V_0. This works in part since an irreducible family over the line is automatically flat. In difference algebraic geometry, this still works over the short line, but not over the transformal affine line. It is possible to use blowing ups to remedy the situation, but the use of valuation theory appears to be much clearer.

Transformal valued fields are studied in §7. They differ from the valued difference fields of [Scanlon], primarily in that the endomorphism does not fix the value group: on the contrary, it acts as a rapidly increasing map. We are not concerned with the first order theory here, but rather with the structure of certain analogues of discrete valuation rings. Their value group is not $\mathbb{Z}$ but $\mathbb{Z}[\sigma]$, the polynomial ring over $\mathbb{Z}$ in a variable σ, with $\sigma > \mathbb{Z}$. A satisfactory structure theory can be obtained, especially for the completions. Extensions of such fields of finite transcendence degree can be described in terms of inertial and totally ramified extensions, and the completion process; there are no other immediate extensions. As a result, the sum of the inertial transcendence degree extension and the ramification dimension is a good invariant; cf. 7.37. This valuative dimension will be used in §8.

Given a family X_t of difference varieties, we define in §9 the specialization $X_{\to 0}$, a difference subvariety of the "naive" special fiber X_0. The total dimension of $X_{\to 0}$ is at most that of the generic fiber X_t; this need not be the case for X_0. Together with the material in §8, this lays the basis for a theory of rational and algebraic equivalence of transformal cycles. (It will be made explicit in a future work.)

The main discovery allowing the description of specializations in terms of transformal valuation rings is this: definable sets of finite total dimension are analyzable in terms of the residue field. The simplest case is of a definable set X_t in K, contained entirely in the valuation ring, and such that the residue map is injective on X_t. The equation $\sigma(x) = x$ for is immediately seen to have this property. In general, for an equation of transformal dimension zero, X_t is *analyzable* over the residue field (§8.2.) This model-theoretic notion is explained in §8.3. It means that sets of finite total dimension over the generic fiber can actually be viewed as belonging to the residue field; but in a somewhat sophisticated sense, that will be treated separately. Here we will content ourselves with the numerical consequence for Frobenius fields, bounding the number of points specializing to a given point of X_0 in terms of the valuative dimension. With this relation in mind, §8.4 contains a somewhat technical result, needed in the main estimates, bounding the valuative dimension in terms of directly visible data, when a difference scheme is presented geometrically as (V, S).

The proof of the main theorem begins in §11. In this section, it is reduced to an intersection problem on a smooth, complete algebraic variety V: one has correspondence $S \leq V \times V^{\sigma}$, and wants to estimate the number of points of $S \cap \Phi_q$, where Φ_q is a graph of Frobenius; or rather those points lying outside a proper subvariety W of V. The main difficulty is that, away from W, the projection $S \to V$ need not be quasi-finite. For instance, S may contain a "square" $C \times C$, and then any Frobenius will meet S in an infinite set. A moving lemma in §11.1 shows that, while this can perhaps not be avoided, when viewed as a cycle S can be moved to another, S', whose components do not raise a similar difficulty.

In §12 we obtain an estimate for the intersection number $S \cdot \Phi_q$ in the sense of intersection theory. See 12.3 for a quick proof in the case of projective space, and 12.4 for a proof for curves: indeed Weil's proof, using positivity in the intersection product on a surface, works in our case too.

For general varieties, in the absence of the standard conjectures, we use the cohomological representation and Deligne's theorem. According to the Lefschetz fixed point formula, the question amounts to a question on the eigenvalues of the composed correspondence $\Phi^{-1} \circ S$, acting on the cohomology $H^*(V)$ of V. When $V = V^\sigma$ and S, Φ commute, one can apply Deligne's theorem on the eigenvalues of Φ. But in general, Φ induces a map between two distinct spaces $H^*(V)$ and $H^*(V^\sigma)$; so one cannot speak of eigenvalues. Some trickery is therefore needed , in order to apply the results of Deligne. It is in these manipulations that the precision of Deligne's theorem is needed, rather than just the Lang-Weil estimates. It may suffice to know that every eigenvalue on H^i, $i < 2n$, has absolute value $\leq q^{2n-1/2}$; but because of possible cancellations in the trace (when e.g. all m'th roots of unity are eigenvalues for some m), this still seems beyond Lang-Weil.

It would be interesting to sharpen the statement of the theorem, so as to give a precise rather than asymptotic cohomological account of the size of Frobenius specializations of zero-dimensional difference schemes.

Having found the intersection number, we still do not know the number of (isolated) points of the intersection $S \cap \Phi_q(\backslash W)$; this is due to the possible existence of infinite components, mentioned above. Cf. [Fulton], and Kleiman's essay in [Seidenberg1980]. §13 is devoted to estimating the "equivalence" of these components. A purely geometric proof using the theory of [Fulton] appears to be difficult; the main trouble lies in telling apart the 0-dimensional distinguished subvarieties of the intersection, from the true isolated points (not embedded in larger components) that we actually wish to count. Instead, we use the methods of specialization of difference schemes, developed above. By the moving lemma, there exists a better-behaved cycle S_t on $V \times V^\sigma$, specializing to S. The intersection numbers corresponding to S_t, S are the same; and we can assume that the equivalence has been bounded for S_t. The problem is to bound the number of points of $S_t \cap \Phi_q$ specializing into W; as well as the number (with multiplicities) of points of $S_0 \cap \Phi_q$ to which no point of $S_t \cap \Phi_q$ specializes. The inverse image of W under the residue map is shown to be analyzable over the residue field, of dimension smaller than $\dim(V)$, proving this point.

This use of intersection theory does not flow smoothly, since the difference varieties do not really wish to remain restricted to $V \times V^\sigma$. They are often more naturally represented by subvarieties of, say, $V \times V^\sigma \times V^{\sigma^2}$. Even the decomposition into irreducible components of distinct transformal dimensions cannot be carried out in $V \times V^\sigma$; e.g. above, the closure of $\{x \in V \setminus W : (x, \sigma(x)) \in S\}$ is always a difference scheme of finite total dimension; but S is irreducible, and the closure cannot be represented on $V \times V^\sigma$. However, the intersection-theoretic and cohomological methods that we use do not seem to make sense when one intersects two $\dim(V)$-dimensional subvarieties of $V \times V^{\sigma} \times V^{\sigma^2}$; and so we must strain to push the data into $V \times V^{\sigma}$ (as in 8.4.) This is one of a number of places leading one to dream of a broader formalism.

The proofs of the various forms of the main theorem, the applications mentioned above, and a few others, are gathered in §14 and §15.

I would like to thank Ron Livne and Dan Abramovich for useful discussions of this problem, and Gabriel Carlyle for his excellent comments on the original manuscript. Warm thanks to Richard Cohn for information about the Jacobi bound problem.

Part I

Difference schemes

2 Well-mixed rings

The basic reference for difference algebra is [Cohn].

A *difference ring* is a commutative ring R with 1 and with a distinguished ring homomorphism $\sigma : R \to R$. We write $a^{\sigma} = \sigma(a)$, and $a^{\sum m_i \sigma^i} = \Pi_i \sigma^i(a)^{m_i}$.

A *well-mixed ring* is a difference ring satisfying:

$$ab = 0 \Rightarrow ab^{\sigma} = 0$$

Any difference ring R has a smallest difference ideal $I = rad_{wm}(R)$ such that R/I is well-mixed.

Given a difference ring R, a *difference ideal* is an ideal such that $x \in I$ implies $\sigma(x) \in I$. If in addition R/I is well-mixed we say that I is well-mixed. A *transformally prime ideal* is a prime ideal, such that, moreover, $x \in I \Leftrightarrow \sigma(x) \in I$. Every transformally prime ideal is well-mixed. A difference ring is a *difference domain* if the zero-ideal is a transformally prime ideal. [1]

A difference field is a difference domain, that is a field. (An endomorphism of a field is automatically injective.)

Lemma 2.1 *Let R be a difference ring, P an algebraically prime difference ideal, K the fraction field of R/P, $h : R \to K$ the natural map.*

1. *P is transformally prime iff there exists a difference field structure on K, extending the natural quotient structure on R/P.*

[1] Ritt and Cohn ([Cohn]) use the term *prime difference ideal* for what we call a transformally prime ideal. When as we will one considers difference ideals that are prime, but not transformally prime, this terminology becomes awkward. To avoid confusion, we will speak of *algebraically prime* difference ideal to mean difference ideals that are prime. We kept the use of *difference domain*, as being less confusing; though *transformally integral* might be better. At all events, difference rings with the weaker property of having no zero-divisors will be referred to as as *algebraically integral*. Cohn calls an ideal *mixed* if the quotient is well-mixed.

2. *Let $R' \leq R$ be a difference subring, $P' = (P \cap R')$. Let K' be the fraction field of R'/P', and suppose K is algebraic over K'. Then P is transformally prime iff P' is.*

Proof (1) If P is transformally prime, R/P is a difference domain, and σ extends naturally to the field of fractions: if $b \neq 0$ then $\sigma(b) \neq 0$, and one can define $\sigma(a/b) = \sigma(a)/\sigma(b)$. Conversely, if L is a difference field, and $R \to_h R/P \subset L$ is a homomorphism, if $a \notin P$ then $h(a) \neq 0$ so $h(\sigma(a)) = \sigma(h(a)) \neq 0$, and thus $\sigma(a) \notin P$.

(2) If P is transformally prime, clearly P' is too. If P' is transformally prime, then the multiplicative set $h(R' \setminus P') \subseteq K' \setminus (0) \subset K \setminus (0)$ is closed under σ, so σ extends to the localization of $h(R)$ by this set, i.e. to $M = K'(R/P)$. But since $K' \subseteq M \subseteq K$ and K is algebraic over K', M is a field, hence $M = K$. □

Notation 2.2 *Let I be a difference ideal.*

$$^{\sigma}\sqrt{(I)} = \{a \in R : a^{\sum_{0 \leq i \leq n} m_i \sigma^i} \in I, \text{ some } m_0, m_1, \ldots, m_n \geq 0\}$$

$$\sqrt{(I)} = \{a \in R : a^n \in I, \text{ some } n \in \mathbb{N}\}$$

Lemma 2.3 *Let R be a well-mixed ring. Then $\sqrt{(0)}$ is a radical, well-mixed ideal. $^{\sigma}\sqrt{(0)}$ coincides with*

$I = \{a \in R : a^{m\sigma^i} = 0 \text{ for some } m \geq 1, i \geq 0\} = \{a \in R : a^{\sigma^i} \in \sqrt{(0)} \text{ for some } i \geq 0\}$

and is a perfect ideal.

Proof The statement regarding $\sqrt{(0)}$ is left to the reader. If $a^{m\sigma^i} = 0$, and $m \leq m'$, $i \leq i'$, then $a^{m'\sigma^{i'}} = 0$. It follows that I is closed under σ and under multiplication, and addition (if $a^{m\sigma^i} = 0$ and $b^{m\sigma^i} = 0$ then $(a+b)^{2m\sigma^i} = 0$.) Moreover if $aa^{\sigma} \in I$, then $a^{(\sigma+1)m\sigma^i} = 0$. So $bb^{\sigma} = 0$ where $b = a^{m\sigma^i}$. As R is well-mixed, $b^{\sigma}b^{\sigma} = 0$. Thus $a^{2m\sigma^{i+1}} = 0$, so $a \in I$. □

Lemma 2.4 *Let R be a difference ring. Any intersection of well-mixed ideals is well-mixed.*

Proof : Clear.

We write $Ann(a/I)$ for $\{b \in R : ab \in I\}$, $Ann(a) = \{b \in R : ab = 0\}$.

Lemma 2.5 *Let R be a well-mixed ring, $a \in R$, $I(a)$ the smallest well-mixed ideal containing a. Then $Ann(a)$ is a well-mixed ideal. If $b \in I(a) \cap Ann(a)$, then $b^2 = 0$.*

Proof Let $b \in Ann(a)$. Then $\sigma(b) \in Ann(a)$ since R is well-mixed. Thus $Ann(a)$ is a difference ideal. Suppose $cb \in Ann(a)$. Then $abc = 0$. So $abc^{\sigma} = 0$, as R is well-mixed. Thus $bc^{\sigma} \in Ann(a)$. So $Ann(a)$ is well-mixed.

Let $c \in Ann(a)$. Then $a \in Ann(c)$. Since $Ann(c)$ is an well-mixed ideal, $I(a) \subset Ann(c)$. Thus if at the same time $c \in I(a)$, we have $c \in Ann(c)$, so $c^2 = 0$.

Lemma 2.6 *Let R be a well-mixed domain, $a \in R$, $J(a)$ the smallest algebraically radical, well-mixed ideal containing a. Then $Ann(a)$ is an algebraically radical, well-mixed ideal; and $J(a) \cap Ann(a) = (0)$.*

Proof In 2.5 it was shown that $Ann(a)$ is a well-mixed ideal. If $b^m \in Ann(a)$, then $b^m a = 0$, so $(ba)^m = 0$, so $ba = 0$. Thus $Ann(a)$ is algebraically radical. Let $c \in Ann(a)$. Then $a \in Ann(c)$. Since $Ann(c)$ is a well-mixed, algebraically radical ideal, $J(a) \subset Ann(c)$. Thus if at the same time $c \in I(a)$, we have $c \in Ann(c)$, so $c^2 = 0$, hence $c = 0$.

Lemma 2.7 *Let R be a well-mixed ring, $0 \neq a \in R$. For $p \in \operatorname{Spec}^\sigma R$, let a_p denote the image of a in the local ring R_p. Then $\{p \in \operatorname{Spec}^\sigma R : a_p \neq 0\}$ is a nonempty closed subset of* $\operatorname{Spec}^\sigma R$.

Proof By 2.6, $Ann(a)$ is a well-mixed ideal. Let p be a maximal well-mixed ideal containing $Ann(a)$. Then p is a prime ideal. For if $cd \in p$, then $p \subset Ann(c/p)$, so $p = Ann(c/p)$ (and then $d \in p$) or $R = Ann(c/p)$ (and then $c \in p$). Moreover as p is well-mixed, it is a difference ideal, so $p \subset \sigma^{-1}(p)$; again by maximality of p, $p = \sigma^{-1}(p)$. Thus p is a transformally prime ideal. If $d \notin p$, then $ad \neq 0$ since $d \notin Ann(a)$. So $a_p \neq 0$. □

Lemma 2.8 *Let R be a well-mixed difference ring, S a subring.*

1. *Let $b \in R$, $q = Ann(b) \cap S$. If $a \in q, \sigma(a) \in S$ then $\sigma(a) \in q$. Same for $Ann(b/\sqrt{(0)}) \cap S$.*
2. *If S is Noetherian and q is a minimal prime of S, the same conclusion holds.*
3. *Call an ideal p of R* cofinally minimal *if for any finite $F \subset R$ there exists a Noetherian subring S of R with $F \subset S$ and with $p \cap S$ a minimal prime of S. Then any cofinally minimal ideal is a difference ideal.*

Proof

1. By 2.3, 2.5, $Ann(b)$ and $Ann(b/\sqrt{(0)})$ are difference ideals of R. The statement about the intersection with S is therefore obvious.
2. Let S be a Noetherian subring of R. Then $\sqrt{(0)}_S = q_1 \cap \ldots \cap q_l$ for some minimal primes $q_1 \cap \ldots \cap q_l$; and $q = q_i$ for some i, say $q = q_1$. Let $b_i \in q_i$, $b_i \notin q$ for $i > 1$, and let $b = b_2 \cdot \ldots \cdot b_l$. Then $q = Ann(b/\sqrt{(0)}) \cap S$.
3. Let p be cofinally minimal, $a \in p$. Let $F = \{a, \sigma(a)\}$, and let S be a Noetherian subring of R containing F with $p \cap S$ a minimal prime of S. By (2), $\sigma(a) \in (p \cap S)$, so $\sigma(a) \in p$.

□

Remark 2.9 *Let k be a Noetherian commutative ring, R be a countably generated k-algebra. Then cofinally minimal ideals exist. Indeed if S is a finitely generated k-subalgebra of R, p_S any minimal prime ideal of S, then p_S extends to a cofinally minimal p with $p_S = p \cap S$.*

Proof Find a sequence $S = S_1 \subset S_2 \subset \ldots$ of finitely generated k-subalgebras of R, with $R = \cup_n S_n$. Find inductively minimal prime ideals p_n of S_n with $p_{n+1} \cap S_n = p_n$, $p_1 = p_S$; then let $p = \cup_n p_n$. Given p_n, we must find a minimal prime p_{n+1} of S_{n+1} with $p_{n+1} \cap S_n = p_n$. Let $r_1, \ldots, r_k$ be the minimal primes of S_{n+1}; then $\cap_i r_i$ is a nil ideal; so $\cap_i (r_i \cap S_n)$ is a nil ideal; thus $\cap_i (r_i \cap S_n) \subset p_n$. As p_n is prime, for some i, $r_i \cap S_n \subset p_n$; as p_n is minimal, $r_i \cap S_n = p_n$. □

Lemma 2.10 *Let R be a difference ring $I = \sqrt{I}$ a well-mixed ideal. Then I is the intersection of algebraically prime difference ideals.*

Proof We may assume $I = 0$. If $a \neq 0$, we must find an algebraically prime difference ideal q with $a \notin q$. By compactness, we may assume here that R is finitely generated. Find using 2.9 a cofinally minimal prime q with $a \notin q$. By 2.8 it is an algebraically prime difference ideal. □

Lemma 2.11 *Let R be a well-mixed ring, $a \in R$. Then $a \in p$ for all $p \in \operatorname{Spec}^\sigma R$ iff $a^n = 0$, some $n \in \mathbb{N}[\sigma]$.*

Proof One direction is immediate. For the other, assume $a^n \neq 0$ for $n \in \mathbb{N}[\sigma]$. Let I be a maximal well-mixed ideal with $a^n \notin I$, $n \in \mathbb{N}[\sigma]$. Clearly $\sqrt{I} = I$. So $I = \cap p_j$, p_j prime well-mixed ideals. If for each j, $a^{n_j} \in p_j$, then $a^{\sum n_j} \in I$, a contradiction. Thus some p_j contains no a^n, and $I \subset p_j$, so by maximality $I = p_j$. For the same reason $I = \sigma^{-1}(I)$. So $p = I \in \operatorname{Spec}^\sigma R$, and $a \notin p$. □

Difference polynomial rings Let R be a difference ring. A *difference monomial* over R is an expression rX^ν, where $\nu = \sum_{i=0}^m m_i \sigma^i$, $m_i \in \mathbb{N}$. The *order* of the monomial is the highest i with $m_i \neq 0$. A *difference polynomial* in one variable X (of order $\leq M$) is a formal sum of difference monomials over R (of order $\leq M$.) The difference polynomials form a difference R-algebra $R[X]_\sigma$.

Transformal derivatives Let $F \in K[X]_\sigma$ be a difference polynomial in one variable. We may write $F(X) = \sum_\nu c_\nu X^\nu$, $\nu \in \mathbb{N}[\sigma]$; where $\{\nu : c_\nu \neq 0\}$, the *support* of F, is assumed finite.

Clearly $F(X+U) = \sum_\nu F_\nu(X) U^\nu$, where F_ν are certain (uniquely defined) polynomials.

Definition 2.12 *The F_ν will be called the* transformal derivatives *of F, and denoted $\partial_\nu x F = F_\nu$.*

Clearly $\partial_\nu x F = 0$ for all but finitely many ν. The $\partial_\nu x F$ satisfy rules analogous to those written down by Hasse. In particular (as can be verified at the level of monomials.)

$\partial_0(F) = F$

$\partial_\nu(F+G) = \partial_\nu x F + \partial_\nu x G$

$\partial_\nu(FG) = \sum_{\mu+\mu'=\nu} F_\mu(X) F_{\mu'}(Y)$

$\partial_{\sigma\nu}\sigma(F) = \sigma(\partial_\nu(F))$

The definitions in several variables are analogous.

3 Definition of difference schemes

3.1 Localization, rings of sections

Localization If R is a difference ring, $X \subset R$, $\sigma(X) \subset X$, $XX \subseteq X$, and $0 \notin X$, we consider the localization of R by X, and write

$R[X^{-1}] = \{\frac{a}{b} : a, b \in R, b \in X\}$

It admits a natural difference ring structure, with $\sigma(\frac{a}{b}) = \frac{\sigma a}{\sigma b}$. (This abuses notation slightly, since R need not inject into $R[X^{-1}]$; $\frac{a}{b}$ is taken to denote an element of $R[X^{-1}]$, the ratio of the images of a, b.)

In particular, if p is a transformally prime ideal, then the localization R_p is defined to be $R[(R \setminus p)^{-1}] = \{\dfrac{a}{b} : a, b \in R, b \notin p\}$

The *difference spectrum*, $\mathrm{Spec}^\sigma (R)$, is defined to be the set of transformally prime ideals. It is made into a topological space in the following way: a *closed* subset of $\mathrm{Spec}^\sigma (R)$ is the set of elements of $\mathrm{Spec}^\sigma (R)$ extending a given ideal I.

An ideal of R is *perfect* if $x\sigma(x) \in I \Rightarrow x, \sigma(x) \in I$. A perfect ideal is an intersection of transformally prime ideals. There is a bijective correspondence between closed sets in $\mathrm{Spec}^\sigma (R)$ and perfect ideals of R. (cf. [Cohn].)

A perfect ideal is well-mixed: modulo a perfect ideal, $ab = 0 \Rightarrow (ab^\sigma)(ab^\sigma)^\sigma = 0 \Rightarrow ab^\sigma = 0$.

We define a sheaf of difference rings on $\mathrm{Spec}^\sigma (R)$, called the structure sheaf and denoted $\mathcal{O}_{\mathrm{Spec}^\sigma R}$, as follows.

If U is an open subset of $\mathrm{Spec}^\sigma R$, a section of $\mathcal{O}_{\mathrm{Spec}^\sigma R}$, by definition, is a function f on U such that $f(p) \in R_p$, and such that for any $p \in U$, for some $b \notin p$ and $a \in R$, $f(q) = \frac{a}{b}$ for any $q \in U$, $b \notin q$.

We let $\mathcal{O}_{\mathrm{Spec}^\sigma R}(U)$ be the collection of all such functions; it is a difference ring with the pointwise operations. One verifies immediately that this gives a sheaf $\mathcal{O}_{\mathrm{Spec}^\sigma R}$.

The space $\mathrm{Spec}^\sigma (R)$ together with the sheaf $\mathcal{O}_{\mathrm{Spec}^\sigma R}$ is called the *affine difference scheme* determined by R.

If $Y = \mathrm{Spec}^\sigma R$, a ring of the form $\mathcal{O}_Y(U)$ will be called an *affine ring*.

Affine rings of global functions .

$c \in R$ is a σ-unit if c belongs to no transformally prime ideal. More generally, $c|^\sigma d$ if for every transformally prime ideal p, for some $b \notin p$, $c|db$.

It is easy to verify that a difference ring R without zero-divisors is an affine ring iff for all $c, d \in R$, if $c|^\sigma d$ then $c|d$.

Each of the two properties: well-mixed, affine, imply a property that one might call "residually local": if $a \in R$, and the image of a in every localization R_p at a transformally prime ideal vanishes, then $a = 0$. (For the well-mixed case, cf. 2.7.)

A *difference scheme* is a topological space X together with a sheaf $\mathcal{O}_X$ of well-mixed rings, locally modeled on affine difference schemes of well-mixed rings. In other words X has an covering by open sets U_i; and there are isomorphisms $f_i : \mathrm{Spec}^\sigma (R_i) \to (U_i, \mathcal{O}_X|U_i)$ for some family of well-mixed rings R_i.

A *morphism of difference schemes* is a morphism of locally ringed spaces, preserving the difference ring structure on the local rings.

If Y is a difference scheme, a difference scheme *over* Y is a difference scheme X together with a morphism $f : X \to Y$. If $Y = \mathrm{Spec}^\sigma D$, we also say that X is over D.

Gluing Assume given a space X with a covering by open sets U_i ; a family of well-mixed rings R_i, and homeomorphisms $f_i : \mathrm{Spec}^\sigma\,(R_i) \to U_i$; such that for any i, j, the map

$$f_j^{-1} f_i : \; f_i^{-1}(U_j) \subset \mathrm{Spec}^\sigma\,(R_i) \;\to\; f_j^{-1}(U_i) \subset \mathrm{Spec}^\sigma\,(R_j)$$

is induced by difference ring isomorphisms between the appropriate localizations of R_i, R_j. In this situation there is a unique difference scheme structure on X, such that $f_i : \mathrm{Spec}^\sigma\,(R_i) \to U_i$ is an isomorphism of difference schemes for each i (with U_i given the open subscheme structure.)

If there are finitely many U_i, each of the form $\mathrm{Spec}^\sigma\,(R_i)$ with R_i a finitely generated difference ring, (or difference D-algebra) we will say that X is of finite type (over D).

Remarks and definitions.

3.1

We will usually consider difference rings R that are finitely generated over a difference field, or over $\mathbb{Z}$, or localizations of such rings. Now by [Cohn], Chapter 3, Theorem V (p. 89), such rings have no ascending chains of perfect ideals. (See 4.26 for a stronger result.) So their difference spectra are Noetherian as topological spaces. (Their Cantor-Bendixon rank will be $< \omega^2$, not in general $< \omega$ as in the case of algebras.)

3.2

A special place is held by difference rings in positive characteristic p whose distinguished homomorphism is the Frobenius $x \mapsto x^q$, q a positive power of p. Note that for such rings, σ is injective iff the ring has no nilpotents, and surjective iff it is perfect. They are always well-mixed.

3.3

$\mathrm{Spec}^\sigma\,(R)$ may be empty: e.g. $R = \mathbb{Q}[X]/(X^2 - 1)$, $\sigma(X) = -X$.

However, this does not happen if R is well-mixed: by 2.3, R has a perfect ideal $I \neq R$; by [Cohn] (or cf. 2.8) a maximal proper perfect difference ideal is prime.

3.4

Let R be a ring, σ a ring endomorphism of R. Then σ acts on $\mathrm{Spec}\, R$, taking a prime p to $\sigma^*(p) = \sigma^{-1}(p)$. The transformally prime ideals are precisely the points of $\mathrm{Spec}\, R$ fixed by this map. (Note that though σ^* is continuous on $\mathrm{Spec}\, R$, the fixed point set is rarely closed.)

In some contexts it may be useful to consider not only points of $\mathrm{Spec}\, R$ fixed by σ, but also those with finite orbits, or even topologically recurrent orbits. (Cf. [Chatzidakis-Hrushovski-Peterzil]). For instance, when R is Noetherian, the closed subscheme of $\mathrm{Spec}^\sigma\, R$ corresponding to a difference ideal J may be empty even if $J \neq R$, while this is avoided if the wider definition is taken. In this paper we take a different approach, replacing arbitrary difference ideals by well-mixed ideals.

3.5

Let R be a difference ring, $\bar{R}$ be the ring of global sections of $\mathrm{Spec}^{\sigma}(R)$. There is a natural map $i : R \to \bar{R}$ It induces a map $i^* : \mathrm{Spec}^{\sigma} \bar{R} \to \mathrm{Spec}^{\sigma} R$.

There is also a natural map in the opposite direction, $\mathrm{Spec}^{\sigma} R \to \mathrm{Spec}^{\sigma} \bar{R}$:

$p \mapsto p^* = \{F \in \bar{R} : F(p) \in pR_p\}$

And $\bar{R}_{p^*} \to R_p$ is defined by :

$F/G \mapsto F(p)/G(p)$

p^* is the largest prime ideal of $\bar{R}$ restricting to p. For if $q \cap R = p$, and $F \in q$ and $F(p) \notin p$, say $F(p) = c/d$ with $d \notin p$, $dF \in q$, $dF(p) = c$, $c \notin pR_p$, but $c \in q \cap R$.

The image of the map * is a closed subset of $\mathrm{Spec}^{\sigma} \bar{R}$; it consists of the transformally prime ideals containing bF whenever $F \in R^*$ vanishes on every $p \in \mathrm{Spec}^{\sigma} R$ with $b \notin p$. I do not know whether the additional primes are of interest (or if they exist) in general. In the well-mixed context, they do not:

Lemma 3.6 *Let R be a well-mixed ring, $X = \mathrm{Spec}^{\sigma} R$ Noetherian (as a topological space), $f \in \mathcal{O}_X(X)$, $p, q \in \mathrm{Spec}^{\sigma} R$. Suppose $f(p) = 0 \in R_p$. Then there exists $b \in R$, $b \notin p$, with $bf(q) = 0 \in R_q$.*

Proof Say $f(q) = \frac{c}{d} \in R_q$. If $ce = 0, e \notin q$, then $f(q) = 0 \in R_q$ and we are done. If $bc = 0, b \notin p$, then $bf(q) = 0 \in R_q$ and we are done again. Thus we may assume $Ann(c) \subset p \cap q$.

Since R is well-mixed, so is $Ann(c)$; by 2.3, $I =^{\sigma} \sqrt{A}nn(c)$ is perfect. Thus $I = p_1 \cap \ldots \cap p_n$ for some transformally prime ideals p_i.

If $p_i \subset p \cap q$, then $f(p_i) = 0 \in R_{p_i}$ (since $p_i \subset p$) and $f(p_i) = c/d \in R_{p_i}$ (since $p_i \subset q$), hence $c = 0 \in R_{p_i}$, contradicting $Ann(c) \subset p_i$. Thus for each i, there exists $a_i \in p_i, a_i \notin q$ or else $b_i \in p_i, b_i \notin p$. Let a be the product of the a_i, b the product of the b_i. So $a \notin q, b \notin p, ab \in I$; thus $(ab)^n \in Ann(c)$, some $n \in \mathbb{N}[\sigma]f$. Replacing a, b by a^n, b^n, we obtain $a \notin q, b \notin p, ab \in Ann(c)$. So $abc = 0$, hence $bc = 0 \in R_q$, and thus $bf(q) = 0 \in R_q$. □

Lemma 3.7 *Let R be a well-mixed ring, with $X = \mathrm{Spec}^{\sigma} R$ Noetherian, $f \in \mathcal{O}_X(X)$, $p \in \mathrm{Spec}^{\sigma} R$. Then there exist $a, b \in R$, $b \notin p$, with $bf - a = 0$. (I.e. $bf(a) - a = 0 \in R_q$ for all $q \in \mathrm{Spec}^{\sigma} R$.)*

Proof By definition of $\mathcal{O}_X(X)$ for some $a, b \in R, b \notin p$, we have: $f(p) = \frac{a}{b} \in R_p$. If $bf - a$ satisfies the lemma, then so does f. So we may assume $f(p) = 0 \in R_p$. By 3.6, for each $q \in \mathrm{Spec}^{\sigma} R$ there exists $b_q \in R, b_q \notin p$, with $b_q f(q) = 0 \in R_q$. But by definition of $\mathcal{O}_X(X)$ again, for some c, d, d' with $d, d' \notin q$ we have $f(q') = \frac{c}{d} \in R_{q'}$ whenever $d' \notin q'$. The equation $b_q f(q) = 0 \in R_q$ means that there exists $d'' \notin q$ with $d'' b_q c = 0$. So $b_q f(q') = 0 \in R_{q'}$ whenever $dd'd'' \notin q'$ Let $U_q = \{q' : dd'd'' \notin q'\}$. By compactness of X, finitely many open sets U_q cover X; say for $q_1, \ldots, q_n$. Let $b = b_{q_1} \cdot \ldots b_{q_n}$. Then $bf(q') = 0 \in R_{q'}$ whenever $q' \in X$. □

If R is affine, without 0-divisors, then $i : R \to \bar{R}$ is an isomorphism; so $X \simeq \mathrm{Spec}^{\sigma} \bar{R}$. This latter statement is true more generally:

Proposition 3.8 *Let R be a well-mixed ring, $X = \mathrm{Spec}^\sigma R$, $\bar{R} = \mathcal{O}_X(X)$. Then $i^* : \mathrm{Spec}^\sigma \bar{R} \to X$ is an isomorphism of difference schemes.*

Proof By 2.7, $i : R \to \bar{R}$ is an embedding. Let $\bar{p} \in \mathrm{Spec}^\sigma \bar{R}$, $p = \bar{p} \cap R$. By 3.7, there exist $a, b \in R$, $b \notin p$, with $bf - a = 0 \in \bar{R}$. So $f \in \bar{p}$ iff $a \in p$. This shows that i^* is injective, and so induces a bijection of points; similarly 3.7, 3.5 show that i induces an isomorphism $R_p \to \bar{R}_{\bar{p}}$. □

3.2 Some functorial constructions

3.9

Let R be a finitely generated difference ring extension of an existentially closed difference field $\mathcal{U}$. The closed points of $\mathrm{Spec}^\sigma R$ are the kernels of the difference ring homomorphisms $f : R \to \mathcal{U}$; if R is provided with generators $(a_1, ..., a_n)$, then the closed point $\ker f$ corresponds to the point $(f(a_1), \dots, f(a_n)) \in \mathcal{U}^n$. For instance if say $R = \mathcal{U}[X, \sigma X, \dots]$, then the closed points of $\mathrm{Spec}^\sigma (R)$ become identified with the points of $\mathcal{U}$ (viewed as the affine line over $\mathcal{U}$.)

In general, if X,Y are difference schemes, we define a point of X with values in Y be a morphism $Y \to X$. The set of Y-valued points of X is denoted $X(Y)$.

The action of σ on $\mathrm{Spec}^\sigma U$ induces an action on the U-valued points of R "by composition". This should not be confused with the action of σ on $\mathrm{Spec}\, R$. In particular, $\mathrm{Spec}^\sigma R$ can be viewed as the set of points of $\mathrm{Spec}\, R$ fixed by σ. But this just means we are looking at difference ideals, and has nothing to do with the fixed field of $\mathcal{U}$.

This observation yields a construction of difference schemes, starting from an action on an ordinary (but non-Noetherian) scheme. I do not know whether or not every difference scheme is obtained in this way; at least it seems not to be the case via a natural functor adjoint to Fix^σ .

Definition 3.10 *Let Y be a scheme, not necessarily Noetherian, and let σ be a morphism of schemes $Y \to Y$. Define a difference scheme $\mathrm{Fix}^\sigma (Y)$ as follows. The underlying space is the set of points of Y fixed by σ. It is given the induced topology (it is in general neither open or closed however.) The structure sheaf is the one induced from that of Y, by taking the direct limit of $\mathcal{O}_Y(U)$ over all open subsets U of Y containing a given open subset of $\mathrm{Fix}^\sigma (Y)$.*

If $Y = \mathrm{Spec}\, R$ is affine, then R is the ring of global sections of the structure sheaf of Y; σ gives a map $\sigma : R \to R$; and one verifies easily that $\mathrm{Fix}^\sigma (Y) = \mathrm{Spec}^\sigma (R, \sigma)$.

Every closed sub-difference scheme X' of $X = \mathrm{Fix}^\sigma (Y)$ has the form $\mathrm{Fix}^\sigma (Y')$, Y' a closed subscheme of Y.

Most of the difference schemes we will encounter will admit projective embeddings; hence they can all be constructed as $\mathrm{Fix}^\sigma (Y)$ for some scheme Y.

A modification of the above approach does yield all difference schemes. Given a difference ring R, let $\mathrm{Spec}' R$ be the set of prime ideals of R, with the following topology: a basic open set has the form $\cap_n \sigma^{-n} G$, where G is a Zariski open set. Equivalently, a basic open set is the image of $\mathrm{Spec}' R'$, where R' is a localization of R by finitely many elements *as a difference* *ring.* Define a structure sheaf, and gluing; obtain a category that one might call *transformation schemes*. Extend the functor Fix^{σ} to this category, and obtain all difference schemes in the image (but a Noetherian difference scheme need not be the image of a Noetherian transformation scheme.)

Products, pullbacks, fiber products

3.11

Let $\mathcal{X}, \mathcal{Y}$ be difference schemes, with underlying spaces X, Y and sheaves of rings $\mathcal{O}_X, \mathcal{O}_Y$. Let $\mathcal{O}_{X,p}$, $\mathcal{O}_{Y,q}$ denote the stalks at points p, q. Let $R_{p,q} = \mathcal{O}_{X,p} \otimes \mathcal{O}_{Y,q}$ be the tensor product, with the natural difference ring structure. Let i_X, i_Y denote the maps $(Id \otimes 1) : \mathcal{O}_{X,p} \to R_{p,q}$, resp. $(1 \otimes Id) : \mathcal{O}_{Y,q} \to R_{p,q}$.

We define the product $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$. As a set, we let Z be the disjoint union over $(p, q) \in X \times Y$ of

$Z_{p,q} = \{r \in \mathrm{Spec}^{\sigma} R_{p,q} : (i_X)^{-1}(r) = p\mathcal{O}_{X,p}, (i_Y)^{-1}(r) = q\mathcal{O}_{Y,p}\}$

If $r \in Z_{p,q}$, we let $pr_X(r) = p$, $pr_Y(r) = q$, $pr(r) = (p, q)$. Let R_r be the localization of $R_{p,q}$ at r.

Begin with open $U \subset X$ and $V \subset Y$, and

$f \in \mathcal{O}_X(U) \otimes \mathcal{O}_Y(V)$

f defines a function F on $pr^{-1}(U \times V)$; $F(r) \in R_r$ is the image of f under the natural homomorphism $\mathcal{O}_X(U) \otimes \mathcal{O}_Y(V) \to R_r$. A function such as F will be called a basic regular function. A set such as

$W_F = \{r \in pr^{-1}(U \times V) : F(r) \notin rR_r\}$

will be called a basic open set. Topologize Z using the sets W_F as a basis for the topology.

Given an open $W \subset Z$, we let $\mathcal{O}_Z(W)$ be the set of functions on W that agree at a neighborhood W of each point with a quotient F'/F, where F, F' are basic regular functions and $W \subset W_F$.

It is easy to check that if $\mathcal{X} = \mathrm{Spec}^{\sigma} R$, $\mathcal{Y} = \mathrm{Spec}^{\sigma} S$, then $\mathcal{Z}$ is isomorphic to $\mathrm{Spec}^{\sigma} (R \otimes S)$.

Similarly, if $\mathcal{X}, \mathcal{Y}$ are difference schemes over $\mathcal{U}$, we can define $\mathcal{X} \times_{\mathcal{U}} \mathcal{Y}$. Alternatively it can be defined as a closed difference subscheme of $\mathcal{X} \times \mathcal{Y}$, see (13) below. If ambiguity can arise as to the map $h : \mathcal{Y} \to \mathcal{U}$, we will write $\mathcal{X} \times_{\mathcal{U}, \langle} \mathcal{Y}$.

Notation 3.12 *Let X be a difference scheme over Y. If y is a point of Y with values in a difference field L, $y : \mathrm{Spec}^{\sigma} L \to Y$, we let $X_y = X \times_y \mathrm{Spec}^{\sigma} L$*

3.13

Let R be a difference ring. A difference-module is a module together with an additive $\sigma : A \to A$, such that $\sigma(a)(\sigma(m)) = \sigma(am)$ for $a \in R$, $m \in M$. A sheaf of difference modules over a sheaf of difference rings is defined as in [Hartshorne], II.5, adding the condition that the sheaf maps respect σ. If X is a difference scheme, a sub(pre)sheaf of $\mathcal{O}_X$, viewed as a (pre)sheaf of difference modules over itself, is called a difference ideal (pre)sheaf. Thus a sheaf $\mathcal{I}$ on X, such that $\mathcal{I}(U)$ is a difference ideal in $\mathcal{O}_X(U)$. We similarly take over the definition of a quasi-coherent sheaf.

Given a difference ideal sheaf $\mathcal{I}$, the stalk $\mathcal{I}_p$ of $\mathcal{I}$ at a point $p \in X$ is a difference ideal of the local difference ring $\mathcal{O}_{X,p}$. define the associated closed subscheme $\mathcal{Y}$ as follows. The underlying set is

$Y = \{p \in X : \mathcal{I}_p \neq \mathcal{O}_{X,p}\}$

with the topology induced from X. Let $\mathcal{O}_Y(U)$ be the ring of maps F on U such that $F(p) \in \mathcal{O}_{X,p}/\mathcal{I}_p$, and U admits a covering by open sets U' such that $F|U'$ is represented by an element of $\mathcal{O}_X(U')$.

Note that $\mathcal{I}$ can be retrieved from Y as a subscheme of X; $\mathcal{I}(U)$ is the kernel of the map $\mathcal{O}_X(U) \to \mathcal{O}_Y(U)$.

3.14

Let $f : X \to Y$ be a morphism of difference schemes, and let $\mathcal{J}$ be a difference ideal sheaf on Y. Then $f^*\mathcal{J}$ defined as in [Hartshorne] II.5 is a difference ideal sheaf on X. The closed subscheme associated with $f^*\mathcal{J}$ is the *pullback* of the closed subscheme associated with $\mathcal{J}$.

3.15

Let $f : X \to Z$ and $g : Y \to Z$ be morphisms of difference schemes. Then we obtain a map $(f, g) : X \times Y \to (Z \times Z)$. Let Z' be the diagonal closed subscheme of $Z \times Z$ (defined by the obvious equations). The pullback of Z' via (f, g) is called the *fiber product* and denoted $X \times_Z Y$.

3.16

Let X be a difference scheme, W an open subset of the set of points of X. Define $\mathcal{O}_W$ to be the restriction of $\mathcal{O}_X$ to W. Then $(W, \mathcal{O}_W)$ is a difference scheme. (The question is local, so we may assume $X = \mathrm{Spec}^\sigma R$; and also, since W is the union of open sets of the form $\{p \in X : a \notin p\}$, we may assume W is of this form. But then $(W, \mathcal{O}_W) = \mathrm{Spec}^\sigma R[a^{-1}, a^{-\sigma}, \ldots].$)

Difference scheme associated to an algebraic scheme For any commutative ring R, there exists a ring homomorphism $h : R \to S$ into a difference ring S, with the universal property for such morphisms: any ring homomorphism on R into a difference ring factors uniquely as gh, with g a difference ring homomorphism. This universal ring is denoted by $S = [\sigma]R$. If D is a difference domain, the same construction in the category of D-algebras is denoted $[\sigma]_D R$. For instance, $[\sigma]\mathbb{Q}[X_0] = \mathcal{U}[X_0, X_1, \ldots]$ with $\sigma(X_i) = X_{i+1}$.

Let $\mathcal{C}$ be the category of reduced, irreducible affine schemes over k, $\mathcal{C}'$ the category of reduced, irreducible affine schemes over k with a distinguished endomorphism compatible with that of k.

If k is an existentially closed difference field (a model of ACFA), the map $\mathrm{Spec}^\sigma\, [\sigma]_k R \to \mathrm{Spec}\, R$ induces a homeomorphism on the closed points.

This gives a functor from $\mathcal{C}$ to $\mathcal{C}'$, adjoint to the forgetful functor $\mathcal{C}' \to \mathcal{C}$. Composing with the Spec^σ functor from $\mathcal{C}'$ to difference schemes, we obtain a natural functor from affine varieties over D to affine difference varieties over D. Call $\mathrm{Spec}^\sigma\, [\sigma]_D R$ the difference scheme associated with $\mathrm{Spec}\, R/\mathrm{Spec}\, D$.

The functor $\mathcal{C} \to \mathcal{C}'$ above does not naturally extend to a similar functor on irreducible schemes to schemes with endomorphisms. (In the affine category, the image of an open subset of a scheme may not be an open subset of the image; rather a countable intersection of open subsets.) Similarly, while irreducible affine difference schemes are the same as irreducible affine schemes with endomorphisms, the notion of gluing is different, reflecting the fact that the endomorphism is viewed as part of the algebraic structure.

These differences cancel out; the composed functor, taking an affine variety to the associated affine difference scheme, does extend to a functor on schemes. (By gluing.) For projective schemes, we will see the associated difference scheme can also be described in another way, defining the difference analog of the functor Proj.

3.3 Components

For ordinary algebraic schemes, we have two notions, of a *reduced* and an *irreducible* scheme. For difference schemes a third type of reducibility arises.

A σ-nilpotent is an element a of a difference ring R such that some product of elements $b_i = \sigma^{m(i)}(a)$ vanishes. R is *perfectly reduced* if it has no σ-nilpotents. Equivalently, 0 is a perfect ideal.

R is *transformally reduced* if $\sigma(a) = 0$ implies $a = 0$ in R.

By way of contrast, if a difference ring R, viewed simply as a ring, has no nilpotent elements (is an integral domain, has $\operatorname{Spec} R$ irreducible) we will say that R is *algebraically reduced* (algebraically integral, algebraically irreducible.)

Lemma 3.17 *Let R be a difference ring, S is a subset closed under multiplication and under σ. Let $R[S^{-1}]$ be the localization. If R is well-mixed / perfectly reduced / algebraically reduced / algebraically irreducible, then so is $R[S^{-1}]$.*

Proof Straightforward verification.

Definition 3.18 *A difference scheme X is* irreducible *if the underlying topological space is irreducible, i.e. it is not the union of two proper closed subsets. X is* well-mixed / perfectly reduced/ algebraically reduced / algebraically irreducible *if for any open U, the rings $\mathcal{O}_X(U)$ have the property. X is* transformally integral *if X is perfectly reduced and irreducible.*

Notation 3.19 *If X is a difference scheme, then the ideals $red_{wm}(R_p)$ (smallest well-mixed ideals of the local rings) generate an ideal sheaf on the structure sheaf $\mathcal{O}_X$; it defines a closed subscheme, the well-mixed subscheme X_{wm} of X. We may similarly defined the underlying perfectly reduced subscheme, and the somewhat thicker $X_{wm,red}$ (underlying algebraically reduced, well-mixed subscheme.)*

Every Noetherian topological space X is a union of finitely many *irreducible components* X_i. The X_i are defined by the following: they are closed subsets, no one contained in another, and $X = \cup_i X_i$.

Definition 3.20

Let X be a difference scheme, and let Z be an irreducible component of X, defined by a prime ideal sheaf $\mathcal{P} \subset \mathcal{O}_X$. Define an ideal sheaf $\mathcal{I}$ as follows:

$$\mathcal{I}(U) = \{a \in \mathcal{O}_X(U) : ab = 0, \text{ some } b \notin \mathcal{P}(U)\}$$

The *sub-difference scheme of X supported on Z* is defined to have underlying space Z, and structure sheaf $(\mathcal{O}_X/\mathcal{I})|Z$.

If Y is the sub-difference scheme of X supported on Z, then the corresponding perfectly reduced scheme will be called the perfectly reduced subscheme supported on Z; and similarly for other reduction notions (such as well-mixed, defined below.)

Definition 3.21 *Let X be a difference scheme over a difference field k . We will say that a property holds of X* absolutely *if it holds of $X \times_{\text{Spec}^\sigma k} \text{Spec}^\sigma K$ for any difference field extension K of k.*

We attach here some lemmas on analogs of separability.

Definition 3.22 *1. A difference domain D is* inversive *if σ_D is an automorphism of D.*

2. *Let D be a difference domain. Up to isomorphism over D, there exists a unique inversive difference domain D^{inv} containing D, and such that $D^{inv} = \cup_m D^{\sigma^{-m}}$. It is called the inversive hull of D.*

3. *Let $K \subset L$ be difference fields. L is* transformally separable *over K if L is linearly disjoint over K from K^{inv}.*

4. *Recall also the classical definition of Weil: Let $K \subset L$ be fields. L/K is a regular extension if L is linearly disjoint from K^{alg} over K.*

Lemma 3.23 *Let D be a difference domain, p a transformally prime ideal of D. Then there exists a unique prime ideal p^{inv} of D^{inv} such that $p^{\text{inv}} \cap D = p$. The natural map* $\text{Spec}^\sigma D^{inv} \to \text{Spec}^\sigma D$ *is a bijection at the level of points.*

Proof If p^{inv} exists, then clearly $a \in p^{\text{inv}}$ iff $\sigma^n(a) \in p$ for some n. So define p' in this way: $p' = \cup_{n \in \mathbb{N}} \sigma^{-n}(p)$. It is easy to check that p' is prime, $\sigma^{-1}(p') = p'$, and $p' \cap D = p$.

Lemma 3.24 *Let X be a algebraically integral difference scheme over a difference field k*

1. *If for some algebraically closed field $L \supset k$, $X \otimes_k L$ is an integral domain , then X is absolutely algebraically integral.*

2. *Let K be the inversive hull of k. If $X \otimes_k K$ is transformally reduced, then X is absolutely transformally reduced.*

3. *If for some algebraically closed difference field $L \supset k$, $X \otimes_k L$ is an integral domain, then X is absolutely transformally integral.*

Proof The question reduces to the case $X = \text{Spec}^\sigma D$, D a k-difference algebra and a domain. (1) Let $D' = D \otimes_k K$. Whether or not D' is a domain clearly depends only on the k-algebra structures of D and K, and in particular the choice of difference structure on K is irrelevant. Moreover for any field extension L of k, we may find L' extending K and containing a copy of L; then by algebra $D \otimes_k L'$ is a domain, hence so is $D \otimes_k L$.

(2) Let $D' = D \otimes_k K$. Let L' be a difference field extending K. Let $e = \sum d_i \otimes c_i$ be a nonzero element of $D' \otimes_K L'$, with $d_i \in D, c_i \in L'$. We may choose the c_i linearly independent over K'. It follows (using inversivity) that the $\sigma(c_i)$ are linearly independent over K' . But then $\sigma(e) = \sum \sigma(d_i) \otimes \sigma(c_i) \neq 0$. Now if L is a difference field extension of k, the inversive hull L' of L extends K. If $0 \neq e \in D \otimes_k L$, then the image of e in $D \otimes_k L'$ is nonzero, since we are dealing with vector spaces. Thus $\sigma(e) \neq 0$ as required.

(3) Follows from (1) and (2).

Remark 3.25

Note that $[\sigma]\mathbb{Q}[\sqrt{2}]$ is not a domain. However let R be a k- algebra and assume $R \otimes_k k^{alg}$ is a domain; then Then $[\sigma]_k R$ is also a domain. More generally, suppose D is a difference domain, subring of an algebraically closed difference field L. Let R be a D-algebra, and suppose $R \otimes_D L$ is a domain. Then $R^* = [\sigma]_D R$ is a domain.

Definition 3.26

Let k be a difference field. An *affine difference variety* over k is a transformally integral affine difference scheme of finite type over k.

4 Dimensions

4.1 Total dimension and transformal dimension

Two types of dimensions are naturally associated with difference equations. If one thinks of sequences (a_i) with $\sigma(a_i) = a_{i+1}$, the *transformal dimension* measures, intuitively,the eventual number of degrees of freedom in choosing a_{i+1}, given the previous elements of the sequence. The *total dimension* measures the sum of all degrees of freedom in all stages. These correspond to transformal transcendence degree and order in [Cohn], but we need to generalize them to difference schemes (i.e. mostly from difference fields to difference rings.)

The dimensions we consider here and later will take extended natural number values $(0, 1, \ldots, \infty)$. We will sometimes define the dimension as the the maximal integer with some property; meaning ∞ if no maximum exists.

Since these dimensions are defined in terms of integral domains, they give the same value to a difference ring R and to R/I, where I is the smallest well-mixed ideal of R.

If k is a difference field and K is a difference field extension, a subset $B = \{b_1, \ldots, b_m\}$ of K is *transformally independent over* k if $b_1, \sigma(b_1), \ldots, b_2, \sigma(b_2), \ldots, b_m, \sigma(b_m), \ldots$ are algebraically independent over k. The size of a maximal k- transformally independent subset of K is called the transformal dimension.

Let k be a difference field, and let R be a difference k-algebra. Consider the set Ξ of triples (h, D, L), with D a algebraically integral difference k- algebra, $h : R \to D$ is a surjective homomorphism of difference k-algebras, and L is the field of fractions of D. Let Ξ' be the set of $(h, D, L) \in \Xi$ where in addition D is a difference domain. In general, L is only a field; but if $(h, D, L) \in \Xi'$, then L carries a canonical difference field structure.

The *transformal dimension* of R over k is the supremum over all $(h, D, L) \in \Xi'$ of the transformal dimension of L over k.

The *reduced total dimension* of R over k is the supremum over all $(h, D, L) \in \Xi'$ of the transcendence degree of L over k.

The *total dimension* of R over k is the supremum over $(h, D, L) \in \Xi$ of the transcendence degree of L over k.

Let X be a difference scheme over $\mathrm{Spec}^\sigma k$. The transformal (resp. total, reduced total) dimension of X over k is the maximum transformal (resp. total, reduced total) dimension of $\mathcal{O}_X(U)$ over k (U an open affine subset of X.)

If $f : X \to Y$ is a morphism of difference schemes, the transformal (resp. total) dimension of X over Y is the supremum of the corresponding dimension of $X \times_Y \mathrm{Spec}^\sigma k$ over $\mathrm{Spec}^\sigma k$, where k is a difference field and $\mathrm{Spec}^\sigma k \to Y$ is a k-valued point of Y. (cf. 4.9 for the coherence of these definitions, when $Y = \mathrm{Spec}^\sigma k$.)

Lemma 4.1 *Let k be a difference field, R a well-mixed difference k-algebra. Then the transformal dimension of R over k is the maximal n such that the transformal polynomial ring in n-variables $k[X_1, \ldots, X_n]_\sigma$ embeds into R over k.*

Proof If the transformal dimension is $\geq n$, let $(h, D, L) \in \Xi'$ show it; then D contains a copy of $k[X_1, \ldots, X_n]_\sigma$. We have $X_i = h(Y_i)$ for some $Y_i \in R$; and $k[Y_1, \ldots, Y_n]_\sigma$ must also be a copy of the transformal polynomial ring. Conversely, assume $k[X_1, \ldots, X_n]_\sigma \simeq_k S \leq R$. Let I be a maximal well-mixed ideal with $I \cap S = (0)$ (note the ideal (0) has this property.)
Claim I is an algebraically prime ideal: if $ab \in I$, $b \notin I$, then $a \in I$.
Proof First suppose $a \in S$. Then $Ann(a/I)$ is a well-mixed ideal (2.5), and $Ann(a/I) \cap S = (0)$ since S is an integral domain. As $b \notin I$, $Ann(a/I) \neq I$. By maximality, $1 \in Ann(a/I)$, i.e. $a \in I$.

Now in general: $Ann(b/I)$ is a well-mixed ideal. If $c \in Ann(b/I) \cap S$, then $bc \in I$, so by the case just covered, as $c \in S$, $c \in I$; so $c = 0$. Thus $Ann(b/I)$ is a well-mixed ideal meeting S trivially, and $1 \notin Ann(b/I)$; so again by maximality, $Ann(b/I) = I$; thus $a \in I$.
Claim I is transformally prime.
Proof Let $J = \sigma^{-1}(I)$. Then J is well-mixed (if $\sigma(ab) \in I$, then $\sigma(a\sigma(b)) = \sigma(a)\sigma^2(b) \in I$.) Also $J \cap S = (0)$. As $I \subset J$, we have $I = J$.

Now R/I shows that the transformal dimension of R is $\geq n$. □

Lemma 4.2 *Let k be a difference field k, R a difference k-algebra. Then R, $R/rad_{wm}(R)$ have the same total dimension. If R is well-mixed and finitely generated, then the following numbers are equal:*

1. *The total dimension $t.dim(R)$ of R.*
2. *The maximal transcendence degree t_2 over k of the fraction field of a quotient of R by a prime ideal.*
3. *$t_3 =$ the maximal n such that the polynomial ring over k in n variables embeds into R.*
4. *$t_4 =$ the maximal Krull dimension of a finitely generated k-subalgebra of R.*

Proof Any difference ring homomorphism of R into a difference ring without zero divisors factors through $R/rad_{wm}(R)$, so the first point is clear. Now assume R is well-mixed, of total dimension n. Clearly $n \leq t_2 \leq t_4$, while $t_4 \leq t_3$ since one can lift a a transcendence basis from the quotient ring to a set of elements of R, necessarily independent over k. Finally, $t_3 \leq t.dim(R)$: R contains a free polynomial ring R_1 of rank t_3. By 2.9, there exists a cofinally minimal (prime) ideal p of R with $p \cap R_1 = (0)$; by 2.8, p is a difference ideal. The quotient R/p clearly has field of fractions with transcendence degree $\geq n$. □

In what follows, the total dimension over F of a well-mixed difference F-algebra B is defined to be the supremum over all B' of the total dimension of B', where B' is a finitely generated difference F-subalgebra of B. Observe that B has the same total dimension as $B/\sqrt{(0)}$. If B is not well-mixed, define the total dimension to be that of B/J, with J the smallest well-mixed ideal of B.

Proposition 4.3 *Let k be a difference field, and X a well-mixed difference scheme of finite type over k. The following conditions are equivalent:*

1. *X has finite total dimension over* $\mathrm{Spec}^\sigma k$
2. *X has finite reduced total dimension over* $\mathrm{Spec}^\sigma k$
3. *X has transformal dimension zero over* $\mathrm{Spec}^\sigma k$

Proof Evidently (1) $\Rightarrow$ (2) $\Rightarrow$ (3) . To prove that (3) $\Rightarrow$ (1) , we may assume $X = \mathrm{Spec}^\sigma R$, R a finitely generated difference k-algebra.

By 4.1, R does not contain a copy of the difference-polynomial ring in one variable $k[t, t^\sigma, \ldots]$.

Let $r_1, \ldots, r_n$ be generators for R, $r_{ij} = \sigma^j(r_i)$. Let Ξ be the set of triples (h, D, L) as in the definition of total dimension. For each i, and each $(h, D, L) \in \Xi$, for some m, there exists a nontrivial polynomial relation $F(h(r_i), \ldots, h(r_{i,m})) = 0$, $F \in k[X_0, \ldots, X_m]$.

By compactness, there exists m and finitely many $F_1, \ldots, F_r \in k[X_0, \ldots, X_m]$ such that for all i and all $(h, D, L) \in \Xi$, for some $j \leq r$, $F_j(h(r_i), \ldots, h(r_{i,m})) = 0$. By the claim below, $tr.deg._k(L) \leq mn$, finishing the proof.

Claim Let k be a difference field, $R = k[a, a^\sigma, \ldots]$ a difference k-algebra with no zero divisors, L the field of fractions of R. Write $a_i = \sigma^i(a)$, and assume $F(a_0, \ldots, a_m) = 0$, $0 \neq F \in k[X_0, \ldots, X_m]$. Then $tr.deg._k(L) \leq (m)$.

Proof It suffices to show that $a_k \in k(a_0, \ldots, a_{k-1})^a$ for all $k \geq m$. We have $F^{\sigma^{k-m}}(a_{k-m}, \ldots, a_k) = 0$. Let l be least such that for some $0 \neq G \in k[X_0, \ldots, X_l]$ we have $G(a_{k-l}, \ldots, a_k) = 0$. Then $a_{k-l}, \ldots, a_{k-1}$ are algebraically independent over k. (If $H(a_{k-l}, \ldots, a_{k-1}) = 0$ then $H^\sigma(a_{k-(l-1)}, \ldots, a_k) = 0$, lowering the value of l.) So $a_k \in k(a_{k-l}, \ldots, a_{k-1})^a \subset k(a_0, \ldots, a_{k-1})^a$. □

Corollary 4.4 *If R is a k-difference algebra with generators $a_0, \ldots, a_n$, and R is an integral domain with field of fractions L, $L_0 = k(\{\sigma^j(a_i) : i \leq n, j \leq d\})$, and $tr.deg._k L_0 \leq d$, then $tr.deg._k L \leq d$.*

Proof By the Claim of 4.3, we have $\sigma^j(a_i) \in (L_0)^a$ for each i and all j. □

Corollary 4.5 *Let k be a difference field, R be a well-mixed difference k-algebra. Assume R is finitely generated as a k-algebra. Then the total dimension of R (as a difference k-algebra) equals the Krull dimension of R. If this dimension is 0, then* $\dim_k R < \infty$.

Proof As R is a finitely generated k-algebra, the equality follows follows from Lemma 4.2. A finitely generated k-algebra of Krull dimension 0 is finite dimensional over k (explicitly: let I be the nil ideal of R; $I = \cap_{i=1}^n p_i$, with p_i prime. R/p_i is a finite field extension of k; so $\dim_k R/p_i < \infty$, hence also $\dim_k R/I < \infty$. Now R is Noetherian, so I is finitely generated; thus I^l/I^{l+1} is a finite-dimensional R/I-space for each k; so $\dim_k I^l/I^{l+1} < \infty$ for each $l \leq r$. It follows that $\dim_k R < \infty$.) □

Lemma 4.6 *Let X be a well-mixed difference scheme of finite type over a difference field k. Let W be a subscheme of X, $\bar{W}$ the closure of W in X. Let $Z = \bar{W} \setminus W$. Then the transformal (total) dimension of Z and of W is at most that of X. If W has finite total dimension, then Z has smaller total dimension. (Hence $W, \bar{W}$ have the same dimension.)*

Here W is a closed subscheme of an open subscheme of X. The closure of W in X is the smallest well-mixed subscheme of X containing W.

Proof The statement regarding transformal dimension is obvious; the one for total dimension follows, say, from 4.2 (4). (A similar statement is true for pro-algebraic varieties.)

Example 4.7 *One cannot expect a strict inequality for transformal dimension.* Consider the subschemes of $\mathbb{A}^3$ defined by $xy = \sigma(x)z$ vs. $x = 0$, or just any 0-dimensional scheme vs. a point.

Remark 4.8

If X is a difference scheme over Y, $\bar{Y}$ is a closed subscheme of Y, and $\bar{X}$ is the pullback to X, then the transformal (total) dimension of $\bar{X}$ over $\bar{Y}$ is bounded by that of X over Y. Any closed subscheme of X will also obviously have dimension $\leq m$ over Y.

Lemma 4.9 *Let k be a difference field, R a difference k-algebra, of total dimension l. Let K be a difference field extension of k. Then $R \otimes_k K$ has total dimension $\leq l$ over K.*

Proof Let (h, D, L) be: an algebraically integral difference K-algebra D, a surjective homomorphism $h : R \otimes_k K \to D$, with L the field of fractions of D. We must show that $tr.deg._K(L) \leq l$. Let $h' = h|R$, $D' = h'(R)$, $L' =$ the field of fractions of D' within L. Then in L, L is the field amalgam of K,L'. Thus $tr.deg._K(L) \leq tr.deg._k(L') \leq l$. □

Remark 4.10

If K/k is a regular field extension, or if R/k is a contained in a regular extension, equality holds in 4.9.

In general it may not, because of possible incompatibility within k^{alg}: if $k = \mathbb{Q}$, $a \in R, b \in K$ with $a^2 = 2, \sigma(a) = a, b^2 = 2, \sigma(b) = -b$, then regardless of the total dimension of R, the total dimension of $R \otimes_{\mathbb{Q}} K$ equals zero.

If one counts total dimension only with respect to difference domains lying within a given universal domain (model of ACFA), and this universal domain contains K, then again this dimension is base-change invariant.

Lemma 4.11 *Let K be a difference field, R a difference K-algebra, of total dimension l. Then $R \otimes_K K^{inv}$ has total dimension l over K^{inv}.*

Proof The total dimension of $R \otimes_K K^{inv}$ is no smaller than that of R, by Lemma 4.9. Thus we must show that it equals at least l. We may pass to to a quotient of R demonstrating that R has total dimension $\geq l$; i.e. we may assume R is an integral domain, whose field of fractions has K-transcendence degree $\geq l$. Next, if the characteristic is $p > 0$, we may assume K is perfect. For let L be the perfect closure of K. The perfect closure of R shows that $R \otimes_K L$ has total dimension $\geq l$. By the perfect case, $R \otimes_K L^{inv}$ has total dimension $\geq l$. By 4.9, $R \otimes_K K^{inv}$ does too. So assume K is perfect.

If R/K is a regular field extension, then $R \otimes_K K^{inv}$ is an integral domain, and the assertion is clear.

Let $K_1 = R \cap K^{alg}$. Then R is a regular extension of K_1. As the total dimension of R (viewed as a K_1-algebra) is still $\geq l$, so is the total dimension of $R \otimes_{K_1} {K_1}^{inv}$, hence by 4.9, also of $R \otimes_{K_1} K_1 K^{inv}$. As K_1 is algebraic over K, in the definition of total dimension, it matters not if the transcendence degree is computed over K or over K_1. So as a K-algebra, $R \otimes_{K_1} K_1 K^{inv}$ has total dimension $\geq l$. Now $R \otimes_{K_1} K_1 K^{inv}$ is a homomorphic image of $R \otimes_K K_1 K^{inv}$, which is in turn a homomorphic image of $R \otimes_K K^{inv} \otimes_K K_1$. Thus this difference ring too has total dimension $\geq l$, and by 4.9, so does $R \otimes_K K^{inv}$. □

Note that the reduced total dimension can certainly go down upon base change to K^{inv}, even if $k = k^{alg}$; e.g. $R = k(t)$, $K = k(\sigma(t)) \subset R$.

Embedding in the transformal line A difference domain R is *twisted-periodic* if for some n, p,m, $R \models (\forall x) \sigma^n(x) = x^{p^m}$. If a subring R of a difference field K is not twisted-periodic, then R^n is Ritt-Cohn dense in K^n. (cf. [Cohn]. Note that if every element of R satisfies a difference equation of total dimension n, then every element of K satisfies one of total dimension $\leq 2n$. In one dimension, polarize.)

Below, a morphism $f : X \to Y$ of difference schemes over a field K will be said to be *point - injective* if for every difference field extension K' of K, f induces an injective map $X(K') \to Y(K')$.

Lemma 4.12 *Let R be a difference domain, with field of fractions K; let $r : R \to k$ be a surjective homomorphism into a difference field k, with k not twisted-periodic. Assume R is a valuation ring. Let $X \subset \mathbb{A}^n$ be an difference scheme over R, of finite total dimension. Then there exist $c_1, \ldots, c_n \in R$, such that if $H(x) = \sum c_i x_i$, $h(x) = \sum \bar{c}_i x_i$, then H and h are point- injective.*

Proof If $\sum c_i x_i = \sum c_i y_i$, $x, y \in X$ distinct, then $c = (c_1, \ldots, c_n)$ has transformal dimension $< n$ over $K(x, y)$. Thus $\{c : (\exists x \neq y \in X) c \cdot x = c \cdot y\}$ has transformal dimension $\leq (n-1) + 2 trans.dim(X) = n - 1$. So for any c, outside a proper difference subscheme of

$\mathbb{A}^n$, $H(x) = \sum c_i x_i$ is injective. □

Corollary 4.13 *Let X be a projective difference scheme of finite total dimension over a field k, with k not twisted-periodic. Then there exists a linear projection to $\mathbb{P}^1$, defined over k, point-injective on X.*

Proof Say $X \subset \mathbb{P}^n$. The set of hyperplanes passing through some point of X is contained in a difference scheme of transformal dimension $n-1$. So there exists a hyperplane defined over k and avoiding X. Thus we may assume $X \subset \mathbb{A}^n$. Now 4.12 applies.

(Improve this to: X is isomorphic to a difference subscheme of $\mathbb{P}^1$.)

4.2 Dimensions of generic specializations

Definition 4.14 *Let X be a difference scheme over a difference domain D with field of fractions K. A property will be said to hold* generically *of X if it is true of $X \times_{\mathrm{Spec}^\sigma D} \mathrm{Spec}^\sigma K$.*

For example, if R is a difference D-algebra, the generic *transformal (total) dimension of R over D is the transformal (total) dimension of $R \otimes_D K$ over K.*

Lemma 4.15 *Let D be a difference domain, R a finitely generated difference D-algebra, of generic transformal (total) dimension d over D. Then for some dense open $U \subset \mathrm{Spec}^\sigma(D)$, for all $p \in U$, if $h : D \to K$ is a map into a difference field with kernel p, then $R \otimes_D K$ has transformal (total) dimension $\leq d$ over K.*

Proof It is easy to give an effective argument (cf. 10.8.) We give a qualitative one here. Suppose the lemma is false; then there exist difference ring homomorphisms $h_i : R \to D_i \subset L_i$ such that $h_i(R) = D_i$ is a domain with field of fractions L_i, $h_i(D)$ has field of fractions K_i, and L_i is a difference field of transformal dimension $> d$ over K_i, (respectively, $tr.deg._{K_i} L_i > d$); and such that $p_i = ker(h_i|D)$ approaches the generic point of D. Let F be a finite set of generators of R as a D-algebra. Then for each i, for some $r_0, \ldots, r_d \in F$ (respectively: $F \cup \ldots \cup \sigma^d(F)$) the images $h_i(r_0), \ldots, h_i(r_d)$ are transformally independent over K_i. (In the total dimension case, use 4.4.) We may assume it is always the same sequence $r_0, \ldots, r_d$ (by refining the limit.) Taking an ultraproduct we obtain $h_* : R \to D_* \subset L_*$, injective on D, with the analogs of the above properties; in particular $h_*(r_0), \ldots, h_*(r_d)$ are transformally (resp. algebraically) independent over K_*, contradicting the assumption on generic dimension. □

Remark 4.16

(1) One can find a difference domain D', $D \subset D'$, D' a finite integral extension of D, and a dense open $U \subset \mathrm{Spec}^\sigma(D)$, such that for all $p \in U$, if $h : D \to L$ is a map into a difference field with kernel p, and h extends to a difference ring homomorphism $h' : D' \to L$, then $R \otimes_D L$ has transformal dimension precisely d over L. The proof of this fact, that we will not require, uses 6.4 and 6.2 below.

(2) The base extension D' in (1) cannot be avoided:

for instance let $D = \mathbb{Z}$, $R = D[X, \sigma(X), \ldots]/(X^2 + \sigma(X)^2, \sigma^2(X) - X)$. Then R has transformal dimension 1 over the field of fractions of D. But if p is any rational prime with $p = 1 \bmod 4$, and $h : R \to L$ is a map into a difference field of characteristic p, then the fixed field of (L, σ) has an element i with $i^2 = -1$, so $h(X) = \pm ih(\sigma X)$. Applying σ, $h(\sigma X) = \pm ih(\sigma^2 X)$ so $h\sigma^2(X) = -hX$, yielding $X = 0$. It follows that $R \otimes_D (Z/pZ)$ has transformal dimension 0.

Lemma 4.17 *Let D be a difference domain, with field of fractions of D of transcendence degree $\geq k$ over a difference field k. Let R be a D-difference algebra, of generic total dimension m. Then R has total dimension $\geq k + m$ over k.*

Proof There exists a surjective homomorphism of transformal K-algebras $h : R \otimes_D K \to D'$, D' has field of fractions K', $tr.deg._K K' = m$. So $tr.deg_k K' \geq k + m$, and $h(R)$ generates K' as a field. □

Lemma 4.18 *Let Y be a difference scheme with no strictly decreasing chain of perfectly reduced subschemes, and let X be a difference scheme of finite type over Y (via $f : X \to Y$.) Let m be an integer. Then:*

1. *there exists a perfectly reduced difference subscheme $D_m(f) = D_m(X/Y)$ of Y, such that X_a has total dimension $\geq m$ if a is a generic point of (any component of) $D_m(X/Y)$, and has total dimension $< m$ (or is empty) if $a \notin D_m(X/Y)$.*
2. *The above properties characterize $D_m(X/Y)$ uniquely.*
3. *If X has finite total dimension d, then $D_m(X/Y)$ has total dimension $\leq d - m$.*
4. *If equality hold in (3), then $f^{-1}D_m(X/Y)$ contains a weak component of X. (cf. 4.29).*

Proof The uniqueness (2) is clear, since if Z, Z' are two candidates, each generic point of Z must lie on Z', and vice versa. To show existence (1), and (3), we use Noetherian induction on Y. We may assume Y is perfectly reduced. We may also assume Y is irreducible, since if it has a number of components Y_j, then we can let $D_m(X/Y) = \cup_j D_m(f^{-1}Y_j/Y_j)$. If X/Y has generic total dimension $\geq m$, let $D_m(X/Y) = Y$. By 4.17, Y has total dimension $\leq d - m$ in this case. Otherwise, Y has total dimension $\leq m - 1$. By 4.15, there exists a proper closed difference subscheme Y' of Y such that for $a \in Y \setminus Y'$, X_a has total dimension $\leq m - 1$. Now $D_m(f^{-1}Y'/Y')$ exists by Noetherian induction, and we can let $D_m(X/Y) = D_m(f^{-1}Y'/Y')$.

Finally, (4) is clear since $f^{-1}D_m(X/Y)$ is a difference subscheme of X of the same total dimension d (by the first part of (1).) □

Caution: This is not preserved under base change $Y' \to Y$. $D_m(X'/Y') = \emptyset$, $D_m(X/Y) = Y$ is possible.

From a logical point of view, ACFA does not eliminate quantifiers. $D_0(X/Y)$ for instance is therefore not the projection, but rather the difference-scheme theoretic closure of the projection.

4.3 Difference schemes and pro-algebraic varieties

Let X be a difference subscheme of an algebraic variety V over k; more precisely, of the difference scheme $[\sigma]_k V$. Let $V_n = V \times \ldots \times \ldots \times V^{\sigma^n}$. (So $V = V_0$.)

We associate with X a sequence of algebraic subschemes $X[n]$ of V_n; $X[n]$ will be called the n-th -order *weak Zariski closure of* X. We describe these locally on V. Let U be an open affine subset of V. So U can be identified with $\operatorname{Spec} A$, with $A = \mathcal{O}_V(U)$, a k-algebra. The inclusion $X \to [\sigma]_k V$ gives a k-algebra homomorphism $[\sigma]_k A \to \mathcal{X}(U)$, with kernel I (so I is a difference ideal of $\operatorname{Spec}^\sigma [\sigma]_k A$, and $X \cap U \simeq \operatorname{Spec}^\sigma ([\sigma]_k A)/I$.) Let A_n be the sub k-algebra of A generated by $A \cup \ldots \cup \sigma^n(A)$; so $\operatorname{Spec} A_n = U \times \ldots \times \ldots \times U^{\sigma^n}$. Finally define $(X \cap U)_n = \operatorname{Spec} A_n/(I \cap A_n)$.

We also let X_ω be the projective limit of the $X[n]$; it can be viewed as a scheme, or as a pro-(scheme of finite type). In particular, $[\sigma]_k V_\omega = \Pi_{n \geq 0} V^{\sigma^n}$. Note that $[\sigma]_k V_\omega$ is isomorphic to the scheme $\Pi_{n \geq 1} V^{\sigma^n}$ (by an isomorphism intertwining with $\sigma : k \to \sigma(k)$). At the same time we have the projection $r : \Pi_{n \geq 0} V^{\sigma^n} \to \Pi_{n \geq 1} V^{\sigma^n}$.

Remark 4.19 *A subscheme Y of $\Pi_{n \geq 0} V^{\sigma^n}$ has the form X_ω for some difference subscheme X of $[\sigma]_k V$ iff Y contains $rY^{\sigma^{-1}}$ as a scheme*

Proof Locally, Y is defined by an ideal I. I is a difference ideal iff $I \subset \sigma^{-1}(I \cap R^\sigma)$.

Definition 4.20 *X is* weakly Zariski dense in V *if the 0-th-order weak Zariski closure of X equals V.*

It is quite possible that X be weakly Zariski dense in V, while the set of points of X is not Zariski dense in V. For instance, when $V = \mathbb{A}^1$, the subscheme X cut out by $\sigma(x) = 0$ has this property.

If X is algebraically integral, each $X[n]$ is an irreducible algebraic variety over k, of dimension $\leq n \dim V$. The natural map $X[n+1] \to X[n]$ is dominant.

4.4 Transformal degree

Let notation be as above: k is a difference field, $X[n]$ is the n'th-order weak Zariski closure of X.

Lemma and Definition 4.21 *Assume X is an algebraically integral difference subscheme of an algebraic variety V over k. There exist integers a, b such that for all sufficiently large $n \in \mathbb{N}$,*

$$\dim (X[n]) = a(n+1) + b$$

a is the transformal dimension of X, while b is called the dimension growth degree *or* transformal degree. *If $a = 0$ then b is the total dimension of X.*

Proof One proof uses intersections with generic hyperplanes. We can assume V is projective. We define the notion of a generic hyperplane section X'; it is the intersection of X with a linear equation, whose coefficients are generic in the transformal sense. Show that if X has positive transformal dimension, then $\dim X'[n] = \dim X[n] - 1$ for large n. In this way reduce to the case of transformal dimension 0, where we must show that $\dim (X_n)$ is bounded (hence eventually constant), with eventual value equal to the total dimension. This is clear by Proposition 4.3 and Lemma 4.2.

Here is a more direct proof. Let $(c_0, c_1, \ldots)$ be such that $(c_0, \ldots, c_n)$ is a generic point over k of $X[n]$, in some field extension of k. (This makes sense since X is algebraically integral; so $X[n]$ is an irreducible variety over k; and $X[n+1]$ projects dominantly to $X[n]$.) So $(c_0, \ldots, c_n) \to (c_l, \ldots, c_{n+l})$ is a (Weil) specialization over k, and thus (with $k[l,n] = k(c_l, \ldots, c_{n+l})$)

$a(l,n) = tr.deg._k k[l,n]$ is non-increasing with l. So for each n, for some $\beta(n)$, for $l \geq \beta(n)$, $a(l,n) = a(l+1,n) = \ldots =_{def} e(n)$. Take the least possible value of $\beta(n+1)$, subject to $\beta(n+1) \geq \beta(n)$.

Now $\epsilon(n) = e(n+1) - e(n)$ is non-increasing: for large enough l, $e(n+1) - e(n) = tr.deg._{k[l,n]} k[l, n+1]$; so

$$e(n+2) - e(n+1) = tr.deg._{k[l,n+1]} k[l, n+2] \leq tr.deg._{k[l+1,n]} k[l+1, n+1] = e(n+1) - e(n)$$

Thus $\epsilon(n)$ is eventually constant, with value a. So for all $n \geq \gamma$, for all $l \geq \beta(n)$, $a(l,n) = \alpha(n) + an$ for some $\alpha(n)$.

Note that $\beta(n) = \beta(n+1)$ unless $e(n) > e(n+1)$. Thus $\beta(n)$ also stabilizes at some maximal β.

It follows that for $l \geq \beta, n \geq \gamma$, $a(l,n) = \alpha + an$.

Finally, $tr.deg._{k(c_\beta, c_{\beta+1}, \ldots)} k(c_0, c_1, \ldots) = \delta = tr.deg._k(c_\beta, c_{\beta+1}, \ldots, c_L) k(c_0, \ldots, c_L)$, for some sufficiently large $L \geq \beta + \gamma$.

So for $n \geq L$, $tr.deg._k k(c_0, \ldots, c_n) = \alpha + \delta + an$. □

Example 4.22 *The eventual dimension growth formula of 4.21 need not hold for all n; the differences* $\dim(X[n+1]) - \dim(X[n])$ *need not be monotone.*

Let $D = \mathbb{Q}[z]_\sigma$, $R = D[x]$ with $\sigma(x) = 0$, and let $y = xz^\sigma$. Let $R[0] = \mathbb{Q}[x,y,z] \leq R$, $R[n] = \mathbb{Q}[x, y, z, z^\sigma, \ldots, z^{\sigma^n}]$. Let $V = \mathbb{A}^3 = \operatorname{Spec} R[0]$, $X[n] = \operatorname{Spec}^\sigma R[n] \subset V^{n+1}$, $a_n = \dim X[n]$. Then $(a_0, a_1, a_2, a_3, \ldots) = (3, 3, 4, 5, \ldots)$.

Remark 4.23 *Assume $X \subset Y$ are two algebraically integral difference subschemes of the algebraic variety V. If X, Y have the same transformal dimension and dimension growth degree, then $X = Y$.*

Proof In fact, for every sufficiently large n, $X[n] \subset Y[n]$, and $\dim(X[n]) = \dim(Y[n])$; so $X[n] = Y[n]$; and so $X = Y$. □

4.5 A chain condition

We present a strong version (4.26) of the Ritt-Raudenbush finite basis theorem for perfect difference ideals.

But first, a simple lemma.

Lemma 4.24 *Let k be a difference field, D an algebraically integral finitely generated k-difference algebra. Then for some n, $\sigma^n(D)$ is a difference domain.*

In fact if R is a finitely generated k-subalgebra of D, generating D as a difference ring, and b is the dimension growth degree of $\mathrm{Spec}^\sigma(D)$ *as a difference subscheme of* $\mathrm{Spec}(R)$*, then one can take $n \leq b$.*

Proof Let a be the transformal dimension of $\mathrm{Spec}^\sigma D$ over k; this equals the transformal dimension of $\mathrm{Spec}^\sigma \sigma^n(D)$ over $\sigma^n(k)$, for any n. Let b_n be the dimension growth degree of $\mathrm{Spec}^\sigma \sigma^n(D)$ as a difference subscheme of $\mathrm{Spec}\, \sigma^n(R)$, a variety over $\sigma^n(k)$. In other words, let $R_{n,m}$ be the subring of D generated by $\sigma^n(R) \cup \ldots \cup \sigma^{n+m-1}(R)$, $K_{n,m}$ the field of fractions of $R_{n,m}$, $c(n,m)$ the transcendence degree of $K_{n,m}$ over $\sigma^n(k)$. Then for any n, for large enough m, $c(n,m) = am + b_n$. We have $b = b_0 \geq b_1 \geq \ldots$, so $b_i = b_{i+1}$ for some $i \leq b$. The lemma now follows by applying the following claim to $\sigma^i(D)$.

Claim If $b = b_1$, then D is a difference domain.

Proof Pick any $c \in D, c \neq 0$, in order to show that $\sigma(c) \neq 0$. Let m be large enough so that $c \in R_{0,m}$, that $c(0,m) = am + b$ and $c(1,m) = am + b_1$. As $b = b_1$, $c(0,m) = c(1,m)$, i.e. $tr.deg._k K_{0,m} = tr.deg._{\sigma(k)} K_{1,m}$. Thus by Krull's theorem, the Krull dimensions of $R_{0,m}$ and of $R_{1,m}$ are equal. Now the surjective homomorphism $\sigma : R_{0,m} \to R_{1,m}$ has a prime ideal p for kernel (as $R_{1,m}$ is an integral domain.) . Thus $R_{0,m}/p \cong R_{1,m}$ has the same Krull dimension as $R_{0,m}$. As 0 is also a a prime ideal of $R_{0,m}$, this forces $0 = p$. So $c \notin p$, and thus $\sigma(c) \neq 0$. □

Lemma 4.25 *There is no strictly descending infinite chain of algebraically integral difference subschemes of an algebraic variety V*

Proof In such a descending chain, the transformal dimension must be non-increasing, and eventually stabilize; after that point, the dimension growth degree cannot increase, and eventually stabilizes too. Indeed by 4.23, if $X \subset Y \subset V$ are algebraically integral, and have the same transformal dimension and dimension growth degree, then $X = Y$. □

Corollary 4.26 *Let k be a difference field, R a finitely generated difference k-algebra. There is no infinite ascending chain of radical, well-mixed difference ideals.*

Proof By 2.10, any radical well-mixed difference ideal is an intersection of algebraically prime difference ideals.

Any algebraically prime difference ideal q contains, by 4.25, a finite subset X, such that any algebraically prime difference ideal q' containing X must also contain q. By the remark above, any radical well-mixed ideal containing X also contains q. Thus at least algebraically prime difference ideals are finitely generated as radical well-mixed ideals.

Suppose not every radical, well-mixed difference is finitely generated as such; then there exists a maximal ideal p with this property. We will get a contradiction once we show that p is algebraically prime. If $ab \in p$, let $p_1 = Ann(a/p)$; then p_1 is well-mixed and radical. Let $p_2 = Ann(p_1/p)$; then p_2 is well-mixed and radical, and $p^2 \subseteq p_1 p_2 \subseteq p$. If $p \neq p_1$ and $p \neq p_2$ then p_1, p_2 are finitely generated as well-mixed radical ideals, say by X_1, X_2. Any algebraically prime difference ideal containing $X_1 X_2$ must contain p_1 or p_2; hence also p. Thus any radical well-mixed difference ideal containing $X_1 X_2$ must contain p. So p is finitely generated, a contradiction. Thus $p = p_1$ or $p = p_2$; so $b \in p$ or $a \in p$. Thus p is prime. □

Corollary 4.27 *Let k be a difference field, R a finitely generated difference k-algebra. If p is a well-mixed ideal and is radical, then there exists a finite number of algebraically prime difference ideals whose intersection is p.*

Proof Using Lemma 4.26, we can prove the statement by Noetherian induction on well-mixed, radical ideals. Let p be such an ideal. If p is algebraically prime, we are done; otherwise, as in the Claim of Lemma 4.26, there exist well-mixed, radical p_1, p_2 with $p_1 p_2 \subseteq p \subseteq p_1 \cap p_2$. Let $q = p_1 \cap p_2$; then $q^2 \subseteq p_1 p_2 \subseteq p$, so as p is radical, $q \subseteq p$; and thus $p = p_1 \cap p_2$. Now by induction, each p_i is the intersection of finitely many algebraically prime difference ideals; hence so is p. □

The corollary can be restated as follows: $S = R/p$ has a finite number of minimal algebraically prime difference ideals; their intersection equals 0 in this ring. (Indeed if $p_1, \ldots, p_m$ are algebraically prime difference ideal of the difference ring S, no one contained in another, and with $p_1 \cap \ldots \cap p_m = (0)$ then each is minimal.)

Let $P_i = \{a : \sigma^m(a) \in p_i, \text{ some } m\}$. By 4.24 applied to R/p_i, for some fixed m, $P_i = \{a : \sigma^m(a) \in p_i\}$. Thus P_i is a transformally prime ideal. Every transformally prime ideal contains some p_i and hence some P_i.

Remark 4.28

Every minimal transformally prime ideal must therefore equal some P_i; but not every P_i must be minimal. (Consider $\mathbb{Q}(x,y)_\sigma/(xx^\sigma, xy^\sigma)$.)

The prime ideals of $R[a^{-1}]_\sigma$ correspond to primes p of R with $a^{-\sigma^n} \notin p$ for all n. This correspondence preserves inclusion, and also preserves the set of difference ideals. The minimal algebraically prime difference ideals of $R[a^{-1}]_\sigma$ are thus precisely the proper $p_i' = p_i R[a^{-1}]_\sigma$, and $p_i' \cap R = p_i$. Thus for any irreducible difference scheme X, one obtains canonically a finite set of algebraically prime difference ideal sheaves; the corresponding closed subschemes are called the *weak components* of X. (They are not necessarily transformally reduced.)

Definition 4.29 *Let X be a difference scheme, Z a component. The* weak components *of X along Z are the weak components of the well-mixed subscheme of X supported on Z. (cf. 3.20).*

Lemma and Definition 4.30 *Let X be a well-mixed closed subscheme of an algebraic variety V over k. There exist integers a, b such that for all sufficiently large $n \in \mathbb{N}$,*

$$\dim (X[n]) = an + b$$

a is the transformal dimension of X, while b is called transformal degree *or the* dimension growth degree. *If $a = 0$ then b is the total dimension of X.*

Proof We may assume $V = \operatorname{Spec}^\sigma A$ is affine; so $X = \operatorname{Spec}^\sigma ([\sigma]_k A)/I$. Moreover, passing to the radical using Lemma 2.3 (an operation that does not change any of the dimensions,) we may assume that I is radical as well as well-mixed. By 4.27, $I = \cap_{i=1}^{m} p_i$, with p_i an algebraically prime difference ideal. Let $X(i) = \mathrm{Spec}^{\sigma}\left([\sigma]_k A)/p_i\right.$. For any n, the dim $(X[n])$ equals the maximum of the corresponding quantity for the $X(i)$; similarly for the transformal and total dimensions. For each i, the result is known by 4.21. It follows for X by looking at the behavior of the integral linear functions involved.

5 Transformal multiplicity

5.1 Transformally separable extensions

Later we will study the locus where the maps $X \to X_n$ fail to be finite. Here we will take a brief look at their behavior on the generic point of X. For this purpose we may pass to the fields of fractions. We define an invariant ι dual to the limit degree. The invariant ι' of 5.2 below figure in the asymptotics in q of the multiplicity of the points of $M_q(X)$, the reduction of X to a difference scheme with structure endomorphism the q-Frobenius. Most likely these lemmas appear in [Cohn].

Lemma 5.1 *Let F be an inversive difference field, K a difference field extension of F with $tr.deg._F(K) < \infty$. Then K^a is inversive.*

Proof $F \subset \sigma(K) \subset K$, and $tr.deg._F(\sigma(K)) = tr.deg._F(K)$. □

Lemma 5.2 *Let $K \subset L \subset M$ be difference fields, with M of finite transcendence degree over K. Consider L, K, M as subfields of M^{inv}.*

1. *L^{inv} is an algebraic extension of $K^{inv}L$. Thus if L is finitely generated over K as a difference field, then LK^{inv} is a finite extension of $\sigma(L)K^{inv}$.*

   Write $\iota(L/K) = [LK^{inv} : \sigma(L)K^{inv}]$, $\iota'(L/K) = [LK^{inv} : \sigma(L)K^{inv}]_{insep}$

2. *Assume M/L is a regular field extension. Then M/K is transformally separable iff M/L and L/K are transformally separable.*

3. *Assume M is finitely generated and transformally separable over K. Then $\iota(M/K) = \iota(M/L)\iota(L/K)$, and similarly for ι'.*

Proof

By Lemma 5.1, we have:

Claim Let K be an inversive difference field, L a difference field extension of K of finite transcendence degree. Then L^{inv} is an algebraic extension of L.

(1) follows: since the transcendence degree of L over K is finite, so is that of LK^{inv} over K^{inv}, and the claim applies.

(2) The "if" direction follows from the transitivity properties of linear disjointness. Assume M/K is transformally separable. Then L/K is a fortiori transformally separable. We need to prove that M/L is transformally separable. M is linearly disjoint from the compositum $K^{inv}L$ over L. We must show that $K^{inv}M$ is linearly disjoint from L^{inv} over $K^{inv}L$. Since M/L is a regular field extension , and M is linearly disjoint from LK^{inv} over L,

MK^{inv}/LK^{inv} is a regular extension. In other words MK^{inv} is linearly disjoint over LK^{inv} from the algebraic closure of LK^{inv}. By (1), $L^{inv} \subset (LK^{inv})^{alg}$.

(3) We may assume here that K is inversive. We have

$$\iota(M/K) = [M : M^\sigma] = [M : M^\sigma L][M^\sigma L : M^\sigma]$$

Yet M is linearly disjoint from $L^{\sigma^{-1}}$ over L, so M^σ is linearly disjoint over L^σ from L, and

$$[M^\sigma L : M^\sigma] = [L : L^\sigma] = \iota(L/K)$$

Similarly

$$[M : M^\sigma L] = [ML^{inv} : M^\sigma L^{inv}] = \iota(M/L)$$

□

5.2 Transformally radicial morphisms

Definition 5.3 *Let $f : R \to S$ be a morphism of difference rings. S is* transformally radicial *over R if for any $s \in S$, for some $m \geq 0$, $s^{\sigma^m} \in f(R)$. A difference scheme X over a difference scheme Y is* transformally radicial *if there exists an open affine covering $\{U_j\}$ of Y and $\{V_{ij}\}$ of X, such that $\mathcal{O}_X(V_{ij})$ is transformally radicial over $\mathcal{O}_Y(U_j)$.*

Remark 5.4 *If S is a finitely generated R-difference algebra and is transformally radicial over R, then S is a finitely generated R-algebra.*

In particular, if a difference scheme X of finite type over Y is transformally radicial over Y, then X has finite total dimension over Y.

Proof Let H be a finite set of generators of R as a difference k-algebra. For $a \in H$, $\sigma^m(a) \in k \cdot 1$ for some m; let $T(a) = \{a, \sigma(a), \ldots, \sigma^{m-1}(a)\}$. Then $H' = \cup_{a \in H} T(a)$ is a finite set of generators of R over as a k- algebra. □

Remark 5.5 *Let k be an inversive difference field, let $R \subset S$ be difference k-algebras, and assume S is transformally radicial over R. Then R,S have the same reduced total dimension over k.*

Proof Let $I = \{r \in R : (\exists m \geq 1)\sigma^m(r) = 0\}$, $J = \{s \in S : (\exists m \geq 1)\sigma^m(s) = 0\}$. Any difference ring homomorphism on R into a difference domain must factor through R/I; thus R,R/I have the same reduced total dimension, and similarly so do $S, S/J$. But $I = J \cap R$, and $R/I \simeq_k S/J$. □

5.3 The reduction sequence of a difference scheme

Let X be a well-mixed scheme. We define a sequence of functors B_n, and maps giving a sequence $X \to B_1X \to B_2X \mapsto \ldots$. Denote B_nX by X_n in this section. We will also define functorially maps $r_n : X \to X_n$ as well as $i_n : X_n \to X$. The maps i_n will allow us to identify X_n with a subscheme of X, but it is the sequence of maps r_n that will really interest us.

When $X = \operatorname{Spec}^\sigma R$ is affine, we let $X_n = \operatorname{Spec}^\sigma \sigma^n(R)$. Let r_n be induced by the inclusion $\sigma^n(R) \subset R$. Let i_n be induced by the surjective homomorphism $\sigma^n : R \to \sigma^n(R)$.

If $f : R \to S$ is a ring homomorphism,, we let $B_n(f) : \sigma^n(R) \to \sigma^n(S)$ be the restriction of f; and with $f^* : \operatorname{Spec}^\sigma S \to \operatorname{Spec}^\sigma R$ the corresponding map of difference schemes, $B_n(f^*) = B_n(f)^*$.

If X is a multiplicatively and transformally closed subset of R, $0 \notin X$, let $S = R[X^{-1}]$. We compare $\sigma^n(S)$ to the localization $R'[\sigma^n(X)^{-1}]$, where $R' = \sigma^n(R)$. The natural map $R'[\sigma^n(X)^{-1}] \to \sigma^n(S)$ is clearly surjective, and since R is assumed to be well-mixed, injective too. Thus

$B_n(j) : B_n(\operatorname{Spec}^\sigma S) \to B_n(\operatorname{Spec}^\sigma R)$

is compatible with localizations. By gluing we obtain a functor B_n on well-mixed schemes.

Observe that $R[X^{-1}]$ and $R[\sigma^n(X)^{-1}]$ may not be the same ring, but they have the same affinization. Indeed an element $a \in X$ may not be invertible in $R[\sigma^n(X)^{-1}]$, but $\sigma^n(a)$ will be invertible, and therefore a will be a σ-unit in this ring.

The maps $i_n : \operatorname{Spec}^\sigma R \to B_n(\operatorname{Spec}^\sigma R)$ and $r_n : B_n(\operatorname{Spec}^\sigma R) \to \operatorname{Spec}^\sigma R$, also globalize.

Remark 5.6

If k is inversive, $B_n(k) = k$, and B_n induces a functor on difference schemes over k. The maps r_n are maps of k-difference schemes, but the maps i_n are not (so that the identification of X_n with a subscheme of X involves a twisting vis a vis k.)

Let X be a well-mixed scheme of finite type (or, of finite type over an inversive difference field K.) It follows from the Remark 5.4 that each ring $\mathcal{O}_X(U)$ is finitely generated over $\mathcal{O}_{B_n(X)}(U)$. In particular it has finite relative total dimension τ_n. If X has finite total dimension, then all the τ_n are bounded by this dimension.

Definition 5.7 *The* transformal multiplicity *of a difference scheme X is the supremum of the total dimensions of the morphisms $X \to B_n(X)$.*

If X is a well-mixed scheme of finite type over a difference field k, we define the transformal multiplicity of X over k *to be the transformal multiplicity of the difference scheme $X \otimes_k K$, where K is the inversive closure of k.*

If X is a well-mixed scheme of finite type over a difference scheme Y, define the relative transformal multiplicity of X over Y *to be the supremum of the transformal multiplicity of X_y over L, where L is a difference field and y is an L-valued point of Y.*

Observe that as a map of points, $\operatorname{Spec}^\sigma X \to \operatorname{Spec}^\sigma B_n(X)$ is bijective. (Every transformally prime ideal p of $\sigma^n(R)$ extends uniquely to a transformally prime ideal of R, namely $\sigma^{-n}(p)$.) While $B_n(X)$ may not be Noetherian as a scheme, $X \to B_n(X)$ is a map of finite type of schemes, and by Lemma 4.5, the transformal multiplicity is the maximum dimension of a fiber of this map (an algebraic scheme) over a point $p \in \operatorname{Spec}^\sigma B_n(X)$.

Example 5.8

Consider the subscheme X of $\mathbb{A}^1$ defined over an algebraically closed difference field K by $f(X^\sigma, \ldots, X^{\sigma^{m+1}}) = 0$, f an irreducible polynomial. If K is inversive, then we may write

$f = g^{\sigma}$, and the corresponding reduced scheme will be given by $g(X, \ldots, X^{si^m}) = 0$; the transformal multiplicity is 1. In general X may be transformally reduced; but it becomes reducible after base change, and the transformal multiplicity is at all events equal to 1.

Proposition 5.9 *Let K be an inversive difference field, X/K a well-mixed difference scheme of finite type and of total dimension d. Let $k \geq 1$. There exist canonically defined closed subschemes $Mlt_k X$ of X such that:*

1. *$X \setminus Mlt_k X$ has transformal multiplicity $< k$ over K.*
2. *$Mlt_k X$ has reduced total dimension $\leq d - k$ over K.*
3. *If $Mlt_k X$ has reduced total dimension $d - k$, then it contains a weak component of X. In this case, X has transformal multiplicity $\geq k$.*

Proof

Let $Y = B_{k+1}X$, and let $r = r_{k+1} : X \to Y$ be the reduction sequence map.

Let $Y_k = D_k(r)$; cf. 4.18.

Let $Mlt_k X = r^{-1} Y_k$. (Or the underlying reduced scheme).

(1) By Lemma 5.13 below, we have to show that if L' is an inversive difference ring and y is an L'-valued point of $B_{k+1}(X \setminus Mlt_k X)$, then X_y has total dimension $< k$ over L'. Now $y \notin Y_k$; by definition of $Mlt_k(Y)$, and 4.18, the total dimension of X_y is $< k$.

(2),(3) come from 4.18 (3), (4).

□

Notation 5.10 $Z_0 X = X \setminus Mlt_1 X$. *For $k \geq$, $Z_k X = Mlt_k X \setminus Mlt_{k+1} X$.*

Remark 5.11 *The assumption in 5.9 that K is inversive is not necessary.*

Proof Let $X' = X \times_K K'$, and let $X'_k = Mlt_k(X')$, satisfying the conclusion of Proposition 5.9. Let $j : X' \to X$ be the natural map of difference schemes. Then the radicial map j induces a bijection between the points of X' and of X, or between the perfectly reduced subschemes of X' and of X, preserving reduced total dimension. Let $Mlt_k X$ be the perfectly reduced subscheme of X corresponding to $Mlt_k X'$. Then (1),(2),(3) are clear. (cf. 4.11). □

The next lemma, 5.12, falls a little short of concluding, when X itself has transformal multiplicity $\leq n$, that $B_n(X)$ must have transformal multiplicity 0. The obstacle can be explained in terms of sheaves of difference algebras; $(B_n)_*(\mathcal{O}_X)$ need not coincide with $\mathcal{O}_{B_n(X)}$.

Lemma 5.12 *Let $X = \mathrm{Spec}^{\sigma} R$ be a well-mixed difference scheme, and assume $X \to B_{n'}(X)$ has total dimension $\leq n < n'$. Then for any difference field L and any L-valued point $y_{n'}$ of $B_{n'}(X)$, and $y \in X(L)$ lifting $y_{n'}$, $\sigma^n(\mathcal{O}_{X,y} \times_{B_{n'}(X), y_{n'}} L)$ has total dimension* 0.

Proof By 4.11, we may take L to be inversive.

We may assume $X = \mathrm{Spec}^{\sigma} R$, R a well-mixed difference ring. Let $R_n = \sigma^n(R)$. Let y be the unique extension of $y_{n'}$ to a difference ring morphism $y : R \to L$; let $y_m = y | R_m$. Let $\otimes'$ denote the tensor product in the well-mixed category; $A \otimes'_B C$ is the quotient of $A \otimes_B C$

by the smallest well-mixed ideal of that ring. Let $S = R\otimes'_{R_{n'},y_{n'}} L$, $z : S \to L$ the induced map, $S_n = \sigma^n(S)$, $z_n = z|S_n$.

We have to show that S_n has total dimension 0. Now S is a well-mixed L-algebra, and $\sigma^{n'}(S) \subset L$. By assumption, S has total dimension $\leq n$. We are reduced to showing:

Claim Let L be an inversive difference field, S a finitely generated well-mixed difference L-algebra, with $\sigma^{n'}(S) \subset L$. Assume S has total dimension $n < n'$. Then $\sigma^n(S)$ has total dimension 0.

There is no harm factoring out the nil ideal of S (as a k-algebra), as this is a difference ideal, and the total dimension of S_n will not be effected. As S is well-mixed, and Noetherian, 0 is the intersection of finitely many prime ideals $p_1, \ldots, p_r$, and they are difference ideals. We have $\cap_{i=1}^r (p_i \cap S_n) = 0$, so it suffices to show that $S_n/(p_i \cap S_n)$ has Krull dimension 0 for each i; for this we may work with S/p_i. Thus we may assume S is an integral domain. At this point we will show that $S = L$. Otherwise let $a \in S \setminus L$. As L is inversive, $\sigma^{n'}(a) = \sigma^{n'}(b)$ for some $b \in L$, so $a - b \notin L$, and $\sigma^{n'}(a - b) = 0$. Thus it suffices to show that for $n \leq m < n'$, if $a \in S$ and $\sigma^{m+1}(a) = 0$ then $\sigma^m(a) = 0$. S is a f.g. domain of Krull dimension $\leq n$. If $a \in S$, then $a, \sigma(a), \ldots, \sigma^n(a)$ are algebraically dependent over L in the field of fractions of S; so there exists an $F \neq 0 \in L[X, X_1, \ldots, X_n]$ with $F(a, \ldots, \sigma^n(a)) = 0$. Take $0 \neq F \in L[X, X_1, \ldots, X_m]$ of smallest number of monomials, and then least degree, such that $F(a, \ldots, \sigma^m(a)) = 0$. Let i_0 be least such that F has a monomial from $L[X, X_1, \ldots, X_{i_0}]$. Apply σ^{m-i_0} to F. The monomial involving X_i for some $i > i_0$ disappear when applied to a (as $\sigma^{m+1}(a) = 0$); deleting them, we obtain a shorter polynomial (though in higher -indexed variables) vanishing on $(a, \ldots, \sigma^m(a))$. This is impossible; so no monomials disappear; i.e. all monomials are from $L[X, X_1, \ldots, X_{i_0}]$; but none are from $L[X, X_1, \ldots, X_{i_0-1}]$. So F has the form $X_{i_0}F'$; and being irreducible, it is just a multiple of X_{i_0}. So $\sigma^{i_0}(a) = 0$. □

Corollary 5.13 *A difference scheme X has transformal multiplicity $\leq n$ iff $X \to B_{n+1}X$ has total dimension $\leq n$.*

Proof It suffices to show that if $n < n'$ and $X \to B_{n'}(X)$ has total dimension $\leq n$, then so does $X \to B_{n'+1}(X)$. Let k be a difference field, y an k-valued point of $B_{n'+1}$. In terms of local rings, we have $R \supset \sigma^{n'}(R) \supset \sigma^{n'+1}(R)$, and $z : \sigma^{n'+1}(R) \to k$. We have to show that $S = R\otimes_{\sigma^{n'+1}(R),z} k$ has total dimension $\leq n$. Let (h, D, L) be as in the definition of total dimension. Let S_m denote the image of $\sigma^m(R)\otimes_{\sigma^{n'+1}(R),z} k$ in S.

Observe in general that if $X \to_f Y$ has relative total dimension $\leq n$, then so does the induced map $BX \to BY$. Indeed, by Remark 4.8, the pullback $f^{-1}BY \to BY$ has total dimension $\leq n$, and BX is a subscheme of $f^{-1}BY$, so 4.8 applies again. In particular, as $X \to B_{n+1}X$ has total dimension $\leq n$, so does $B_{n'-n}X \to B_{n'+1}X$.

By lemma 5.12, $\sigma^n(\sigma^{n'-n}R\otimes_{\sigma^{n'+1}(R),z} k)$ has total dimension 0. Thus $h(S_{n'})$ is finite dimensional over k, so it is a difference subfield of the domain D, call it k'. So $h(\sigma^{n'}(R)) \subset k'$. As $X \to B_{n'}(X)$ has total dimension $\leq n$, $h(R)$ is contained in a field F of k'-transcendence degree $\leq n$. But $[k' : k] < \infty$, so $tr.deg._k(F) = tr.deg_{k'}(F) \leq n$. □

The reduction sequence: two side remarks

Lemma 5.14 *If X is algebraically integral,of finite type over a field K, then the X_n stabilize as subschemes of X: for some n, $X_n = X_{n+1} = \dots$.*

Proof This reduces to the affine case by taking an open affine cover; there it follows from Lemma 4.24. □

(What can be said without the algebraic integrality assumption?

If X is transformally integral, the schemes X_n are all isomorphic to X. The maps $r_n : X \to X_n$ need not be isomorphisms however.

If X is a scheme over an inversive difference field K, then X_n is not necessarily isomorphic to X over K, but is the transform of X under σ^n. In this case we can also define X_n for negative values of n, obtaining a sequence $\dots X_{-2} \to X_{-1} \to X \to X_1 \to \dots$.)

6 Directly presented difference schemes

Definition 6.1

Let D be a difference (well-mixed) ring. Let $F = D[X_1, \dots, X_n]_\sigma$ be the difference polynomial ring over D in variables $X_1, \dots, X_n$. Let I be a difference (well-mixed) ideal of F. I is *directly generated* if I is generated as a difference (resp. well-mixed) ideal by $I \cap F_1$. A surjective difference homomorphism $h : F \to R$ is said to be a *direct presentation of R as a difference (resp. well-mixed) D-algebra* via $R_1 = h(F_1)$ if $Ker(h)$ is directly generated.

In other words, R is generated by R_1, and the relations between the generators can all be deduced from relations between R_1 and $\sigma(R_1)$. Every finitely presented difference D-algebra admits a direct presentation. We will say that the *direct total dimension* of R (for this presentation) is the Krull dimension over D of the ring R_2 generated by $R_1 \cup \sigma(R_1)$, and that R is *directly reduced (irreducible, absolutely irreducible)* if $f = 0$ whenever $f \in R_1$ and $f^n \sigma(f)^m = 0,\ n, m > 0$ (respectively, if R_2 has no zero-divisors, $R_2 \otimes_D L$ has no zero-divisors for some algebraically closed field L containing D.)

A similar definition can be made for difference schemes. Let V be an algebraic scheme over a difference ring D, S a subscheme of $Y \times Y^\sigma$. We let Σ denote the graph of σ on $[\sigma]_D V \times_{\mathrm{Spec}^\sigma D} [\sigma]_D V^\sigma$; more precisely, in any affine open neighborhood where $V = \mathrm{Spec}^\sigma R$, R a difference D-algebra, we let I_σ be the ideal in $R \otimes_D R$ generated by the elements $\sigma(r) \otimes 1 - 1 \otimes r$; Σ is the corresponding closed difference subscheme. We also let $S \star \Sigma$ be the projection to $[\sigma]_D V$ of the difference scheme $S \cap \Sigma$. It is isomorphic to $S \cap \Sigma$, and will sometimes be confounded with it.

A direct presentation of X is an embedding of X into $[\sigma]_D V$, V an algebraic scheme of finite type over D, such that the image of X has the form $S \star \Sigma$ as above.

Lemma 6.2 *Let X be an affine or projective difference scheme (over $\mathbb{Z}$ or over* $\mathrm{Spec}^\sigma k$*, k a difference field.) Then X may be embedded as a closed subscheme of a directly presented (over k) difference scheme $\widetilde{X}$. Moreover, $X, \widetilde{X}$ have the same topological space and the same underlying algebraically reduced well-mixed scheme $X_{wm,red} = \widetilde{X}_{wm,red}$.*

Proof (Affine case.) We may write $X = \mathrm{Spec}^{\sigma}\,([\sigma]_D D[X]/J)$, where $D[X]$ is a polynomial ring over D in finitely many variables, and J a prime ideal of $[\sigma]_D D[X]$. By Proposition 4.26, J is finitely generated as a well-mixed, algebraically reduced ideal. Adding variables, we may assume J is so generated by $J_0 = J \cap D[X, \sigma(X)]$. Let $\widetilde{X} = \mathrm{Spec}^{\sigma}\,([\sigma]_D D[X]/J_0)$. □

.

Definition 6.3 *(The limit degree, cf. [Cohn]). Let Z be an irreducible difference variety over a difference field K, of transformal dimension* 0*. Let a be a generic point of Z over K. Then $K_n = K(a, \sigma(a), \ldots, \sigma^n(a))$ is a field; for large n, K_{n+1} is a finite algebraic extension of K_n; and the degree $[K_{n+1} : K_n]$ is non-increasing with n. Let $\deg_{lim}(Z)$ be the eventual value of this degree.*

Proposition 6.4 *Let X be an algebraic variety over a difference field K. Let $S \subset X \times X^{\sigma}$ be an absolutely reduced and irreducible subvariety,* $\dim(X) = d$, $\dim(S) = d + e$. *Assume S projects dominantly to X and to X^{σ}. Let Z be the difference subscheme of $[\sigma]_K V$ cut out by: $(x, x^{\sigma}) \in S$, W a weak component of Z, and let $W_n \leq X \times \ldots \times X^{\sigma^n}$ be the n'th-order weak Zariski closure of W (cf. 4.21.) Assume $W_1 = S$.*

Then $\dim(W_n) = d + ne$ for all $n \geq 1$.

In particular, W has transformal dimension e and transformal degree $d - e$.

Moreover, if $e = 0$,

$$\deg_{cor} S = \sum_{W} \deg_{cor} Z'$$

Where the sum is taken over all components W of Z that are Zariski dense in X.

Proof

For generic $a \in X$, $S(a) = \{y : (a, y) \in S\}$ has dimension e.

We may freely pass to Zariski open subsets of X. In particular we may assume X is smooth, and that $\dim S(a) \leq e$ for all a.

First assume that for generic $a \in X$, $S(a)$ is absolutely irreducible. Let $(a_0, a_1, \ldots)$ be a sequence of elements in an extension field of K, with $(a_i, a_{i+1}) \in S^{\sigma^i}$, and $tr.deg._K K(a_0, \ldots, a_m) = d + me$. Then $K(a_0, \ldots, a_i)$ is linearly disjoint from $K(a_i, a_{i+1})$ over $K(a_i)$, so the isomorphism type of the field $K(a_1, \ldots, a_n)$ over K is completely determined. Thus $S \times_{X^{\sigma}} S^{\sigma} \times_{X^{\sigma^2}} S^{\sigma^2} \times \ldots \times_{X^{\sigma^{n-1}}} S^{\sigma^{n-1}}$ has a unique component projecting dominantly to each X^{σ^i}. This sequence of components is clearly compatible, forming a difference scheme. Thus there is only one Zariski dense component Z' of Z, and the dimensions and degrees are as predicted.

In general, for generic $a \in X$, $S(a)$ may not be absolutely irreducible; but there always exists a quasi-finite map $\pi : \widetilde{X} \to X$ and a variety $\widetilde{S} \subset (\widetilde{X} \times X^{\sigma})$, such that for generic $\widetilde{a} \in tX$, $\widetilde{S}(\widetilde{a})$ is absolutely irreducible, and $S(a) = \cup_{\pi(\widetilde{a}) = a} \widetilde{S}(\widetilde{a})$. Let

$$S_0(a) = \cup_{\widetilde{a} \neq \widetilde{a}' \in \pi^{-1}(a)} \widetilde{S}(\widetilde{a}) \cap \widetilde{S}(\widetilde{a}')$$

Then $S(a) \setminus S_0(a)$ is the disjoint union of the absolutely irreducible varieties $\widetilde{S}(\widetilde{a})$. We may assume by passing to a Zariski open subset of X that this holds for all $a \in X$, and that π is étale and surjective. Let us write π also for the induced map $\widetilde{S} \to S$. Let $S' = (S \setminus S_0), \widetilde{S}' = \widetilde{S} \setminus \pi^{-1}(S_0)$.

Define inductively $S(m) \subset X \times \ldots X^{\sigma^m}$: $S(0) = X$, $S(1) = S'$,

$$S(m+1) = S(m) \times_{X^{\sigma^m}} \sigma^m(S')$$

By the smoothness of X, all components of $S(m)$ have dimension at least $d + me$. Since $\dim S(a) \leq e$ for all a, they have precisely this dimension.

Moreover, one shows inductively that distinct irreducible components of $S(m)$ are *disjoint*. (An irreducible component U of $S(m+1)$ arises from an irreducible component W of $S^*(m) =_{def} S(m) \times_{X^{\sigma^m}} \widetilde{X}^{\sigma^m}$, as the push-forward by $(Id, \sigma^{m+1}(\pi))$ of $W \times_{\widetilde{X}^{\sigma^m}} \widetilde{S}$. It suffices therefore to show that distinct irreducible components of $S^*(m)$, projecting to the same component of $S(m)$, are disjoint. This is clear since $S^*(m)$ is an étale cover of $S(m)$.)

Let W be a component of Z, weakly Zariski dense in X; then W_m is an irreducible subvariety of S_m, so W_m is contained in a unique component $U(m)$ of S_m. By uniqueness it is clear that the $U(m)$ are compatible with σ and σ^{-1}, so that $U = \cap_m V(U(m), \sigma)$ is a perfectly irreducible difference scheme, containing W. As W is a component, we must have $W = U$, and so $\dim(W_m) = \dim(U_m) = d + me$.

Finally, assume $e = 0$. The additivity of degree may be proved by induction on $\deg_{cor}(S)$. If for all m, $S(m+1)$ is absolutely irreducible, (or at least has a unique component projecting dominantly to $S(m)$), then there is a unique Zariski dense component Z' and $\deg_{lim}(Z') = \deg_{cor}(S)$. Otherwise, consider the minimal m such that $S(m+1)$ is reducible. Let $Y = S(m)$, and let

$$T = \{((a_0, \ldots, a_m), (a_1, \ldots, a_{m+1})) \in (Y \times Y^\sigma) : (a_0, \ldots, a_{m+1}) \in S(m+1)\}$$

Let $\{T_j\}$ be the components of T projecting dominantly to Y. They map finite-to-one to $Y\sigma$, hence dominantly. Then

$$\sum_j \deg_{cor}(T_j) = \deg_{cor}(T)$$

Since $\deg_{cor}(T_j) < \deg_{cor}(Z)$ for each j,

$$\deg_{cor}(T_j) = \sum \deg_{lim}(Z')$$

where the Z' here range over the components of Z , Zariski dense in X, such that for $a \in Z'$

$$(a, \sigma(a), \ldots, \sigma^{m+1}(a)) \in T_j$$

Since each Zariski dense Z' falls into a unique T_j in this sense, the required equation follows.

□

Here is a sketch of a slightly different argument. Suppose Z' is a Zariski dense (in X) component of the wrong growth rate. Let $a \in Z'$ be generic, and let X' be the intersection of X with a generic linear space L of codimension e passing through a, in a suitable projective embedding of X (near a). One obtains an absolutely irreducible $S' \subset X' \times X^\sigma$, with $\dim(S') = \dim(X') = \dim(X) - e$ (Bertini.) Moreover $Z' \cap X'$ is a component of X' whose growth is computed using Bertini, and also seen to be wrong. Since a is generic in X', one can remove the ramification locus of $S' \to X$. Now one obtains a contradiction to the following:

Lemma 6.5 *Let V be a smooth algebraic variety over a difference field K. Let $S \subset V \times V^{\sigma}$ be an absolutely irreducible subvariety,* $\dim(X) = d = \dim(S)$*, and assume either the projection $S \to V$ or $S \to V^{\sigma}$ is étale. Let $Z = \sigma_K S \star \Sigma \subset [\sigma]_K V$. Then Z is a disjoint union of components. Each weak component W is a component, is Zariski dense in V, and has* $\dim(W_n) = d$ *for all n. Z is perfectly reduced.*

Proof Define $S[n]$ Let $V_n = V \times V^{\sigma} \times \ldots \times V^{\sigma^n}$, $f_i : V_n \to (V^{\sigma^i} \times V^{\sigma^{i+1}})$ the projection, $S[n]$ the scheme-theoretic intersection of ${f_i}^{-1}(S^{\sigma^i})$. As in the proof of 6.4, $S[n]$ is a disjoint union of irreducible varieties; hence a (nonempty) weak component of Z is determined by a (compatible) sequence W_n of components of $S[n]$; and $\dim W_n = d$. It follows that Z is perfectly reduced. □

Let $f[1](S) = \{a : \dim(f^{-1}(a) \cap S) \geq 1\}$. Note also, in the converse direction to 6.4:

Remark 6.6

Let $S \subset V \times V^{\sigma}$ be k-varieties, $X = [\sigma]_k S \star \Sigma$, $pr_0[1](S) = \{a : \dim({pr_0}^{-1}(a) \cap S) \geq 1\}$. Suppose $X \cap pr_0[1](S) = \emptyset$. Then every weak component of $[\sigma]_k S \star \Sigma$ of total dimension $d = \dim V$ is a component, and is Zariski dense in V. *Proof*

Since X is contained in the complement of $pr_0[1](S)$, an open subvariety of V, we may pass to this open subvariety; so we may assume $pr_0[1](S) = \emptyset$. In this case, the statement is obvious. But here is a more direct argument: Let $(a_0, a_1, \ldots)$ be a generic point of the pro-algebraic variety corresponding to the weak component C, i.e. $(a_0, \ldots, a_n)$ is a generic point of $C[n]$ for each n. Then a_n cannot be in the closed set $\sigma^n(pr_0[1](S))$ (for instance, because if this were the case for some n, it would be true for all larger values of n. But for sufficiently large n, $a_n \in \sigma^n(X)$, and $\sigma^n(X) \cap \sigma^n(pr_0[1](S)) = \emptyset$.) Now $(a_n, a_{n+1}) \in S^{\sigma^n}$, so $a_{n+1} \in k(a_n)^{alg}$ for each n. So $tr.deg._k k(a_0) = tr.deg._k k(a_0, a_1, \ldots) = d$.

This shows already that C is weakly Zariski dense in V, i.e. $C[0] = V$. Moreover, interpolating $k(a_0, a_1)$ above, we see that $tr.deg._k k(a_0, a_1) = d$, so (a_0, a_1) is a generic point of S, and thus $a_0 \in acl(a_1)$. So the specialization $(a_0, \ldots, a_n) \to (a_1, \ldots, a_{n+1})$ is an isomorphism. Thus $k(a_0, \ldots)$ is a difference field, and hence the prime ideal corresponding to C is a transformally prime ideal, and C is a component. □

7 Transformal valuation rings in transformal dimension one

7.1 Definitions

Notation 7.1 *When K is a valued field, $\mathcal{O}_K, \mathcal{M}_K$,* $\mathrm{val}(K), K_{\mathrm{res}}$ *will denote the valuation ring, maximal ideal; value group, residue field of K; but we will often denote $R = \mathcal{O}_K, M = \mathcal{M}_K, \bar{K} = K_{\mathrm{res}}, \Gamma =$* $\mathrm{val}(K)$.

Definition 7.2 *1. A transformal valuation ring (domain) is a valuation ring R that is also a difference ring, such that $\sigma(M) \subset M$ (and σ injective.)*

2. *The valuation v will be said to be m-increasing (resp. strictly increasing) if for all $a \in R$ with $(v(a) > 0$, $v(\sigma(a)) \geq m \cdot v(a))$ (resp. $v(\sigma(a)) > v(a))$.*

3. *A weakly transformal valued field is an octuple $(K, R, M, \sigma, \Gamma, \bar{K}, \text{ val}, \text{res})$ such that (R, M, σ) is a transformal valuation ring, K is the field of fractions of R, $res : K \to \bar{K}$ is the residue map, $val : K^* \to \Gamma = G_m(K)/G_m(R)$ is the valuation map. We will also apply the term to parts of the data. If σ extends to an endomorphism of K, we denote it too by σ, and we say that K is a transformal valued field.*

4. *Assume given a distinguished $t \in R$, or at least a distinguished $\tau \in \text{ val}(R)$ $(\tau = \text{ val}(t))$. Let*
$$\Delta = \{v \in \text{ val}(R) : -\tau < nv < \tau,\ n = 1, 2, \ldots\}$$
K has transformal ramification dimension *$rk_{ram}(K) = r$ (over τ) if r is the length of a maximal chain $v_0, \ldots, v_r \in \text{ val}(R)$ with $0 = v_0 << v_1 << \ldots << v_r << \tau$. When K is ω- increasing, $rk_{ram}(K) \leq \dim_{\mathbb{Q}}(\mathbb{Q} \otimes \Delta)$. (In transformal dimension one, equality holds; cf 7.6.) When K extends $k(t)_\sigma$, taking $\text{val}(t)$ distinguished, we define the* valuative rank *of K over $k(t)_\sigma$ to be $rk_{ram}(K/k(t)|si) = rk_{\text{val}}(K) = rk_{ram}(K) + tr.deg._k \bar{K}$.*

Lemma 7.3 *Let $(K, R, M, \sigma, \Gamma, \bar{K}, val)$ be a weakly transformal valued field.*

1. *$\sigma^{-1}(M) = M$. $\sigma(M) = M \cap \sigma(R)$. If K is 1-increasing, then every ideal of R is a well-mixed difference ideal.*

2. *There exists a unique difference field structure on $\bar{K}$, such that the residue map $R \to \bar{K}$ is a morphism of difference rings.*

3. *There exits a convex subgroup Γ' of Γ and a homomorphism $\sigma_\Gamma : \Gamma' \to \Gamma$ of ordered Abelian groups, such $val(a) \in \Gamma'$ and $\sigma_\Gamma(val(a)) = val(\sigma(a))$ when $a \in R$, $\sigma(a) \neq 0$.*

   If the valuation is k-increasing, then $kv \leq \sigma_\Gamma(v)$ for $v \in \Gamma, v \geq 0$.

4. *Let L be a subfield of K. Then $\sigma(L^a \cap R) \subset (\sigma(L \cap R))^a$. ($S^a$ denotes the algebraic closure of the field of fractions of S.)*

5. *If L is a subfield of K, $\sigma(L \cap R) \subset L^a$, then $L^a \cap R$ is a difference subring of R.*

6. *Let Δ be a convex subgroup of Γ with ${\sigma_\Gamma}^{-1}(\Delta) \subset \Delta$. (An automatic condition for type 1-increasing valuations.) Let $\hat{v}$ be the valuation $val(a) + \Delta$; with value group Γ/Δ, residue field $\hat{K}$. Then a weakly transformal valued field structure is induced on $\hat{K}$, with residue field $\bar{K}$, value group Δ.*

Proof

1. $\sigma(M) \subset M$ by assumption, while $\sigma^{-1}(M) \subset M$ since M is the unique maximal ideal of R. Thus $\sigma^{-1}(M) = M$. If $\sigma(r) \notin \sigma(M)$, then $r \notin M$, so r is a unit of R, $rr' = 1$, so $\sigma(r)\sigma(r') = 1$, and thus $\sigma(r) \notin M$.

2. σ induces an endomorphism $\sigma : \bar{K} = R/M \to \sigma(R)/\sigma(M) = \sigma(R)/(M \cap \sigma(R)) \subset \bar{K}$ of $\bar{K}$.

3. Let $R' = \{r \in R : \sigma(r) \neq 0\}$, $\Gamma' = \{\pm val(a) : a \in R, a, \sigma(a) \neq 0\}$. If $a, b \in R$ and $ab \in R'$ then $a \in R'$, so Γ' is convex. If $a, b \in R'$ and $val(a) = val(b)$, then $a = cb, b = da$

for some $b, d \in R$, so $\sigma(a) = \sigma(c)\sigma(b), \sigma(b) = \sigma(d)\sigma(a)$, and thus $val\sigma(a) = val\sigma(b)$. Define $\sigma_\Gamma(val(a)) = val(\sigma(a))$, and observe that $\sigma : val(R') \to \Gamma$ is a homomorphism of ordered semi-groups. Extend it to an ordered group homomorphism $\Gamma' \to \Gamma$.

4. Let $b \in L^a \cap R$. There is a nonzero polynomial $P \in L[Y]$ such that $P(b) = 0$. Dividing by the coefficient of lowest value, we can assume the coefficients of P are in R, and at least one of them has value 0. By (1), it follows that $P^\sigma \neq 0 \in \sigma((L \cap R))[X]$, and $P^\sigma(b^\sigma) = 0$.

5. From (4): $\sigma(L^a \cap R) \subset (\sigma(L \cap R))^a \subset L^a$, so $\sigma(L^a \cap R) \subset (L^a \cap R)$.

6. Let $R_\Delta = \{a \in K : (\exists \delta \in \Delta)(\delta \leq val(a))\}$, $M_{>\Delta} = \{r \in K : (\forall \delta \in \Delta)\, \delta < val(r)\}$, $\pi : R_\Delta \to \hat{K}$ the natural map. Let $\hat{R} = \pi(R)$; it is a valuation ring of $\hat{K}$, with maximal ideal $\hat{M} = \pi(M)$.

   If $r \in M_{>\Delta}$, then $\sigma(r) = 0$, or $val(\sigma)(r) = \sigma_\Gamma(val(r)) \notin \Delta$ (otherwise $val(r) \in \Delta$.) So $\sigma(M_{>\Delta}) \subset M_{>\Delta}$, and σ induces an endomorphism $\sigma : \hat{R} \to \hat{R}$. π is a morphism of difference rings, and $\sigma\hat{M} = \sigma(\pi(M)) = \pi(\sigma(M)) \subset \pi(M) = \hat{M}$.

   We have $val(\sigma(r)) \geq val(r)$, implying in particular when $val(r) \in \Delta$ that $val(\sigma\hat{r}) \geq val_{\hat{K}}(\hat{r})$, so that $(\hat{R}, \sigma)$ is 1-increasing if (R, σ) is.

7.2 Transformal discrete valuation rings

We will be interested in ω-increasing transformal valued fields L, finitely generated and of transformal dimension one over trivially valued difference subfield F. We will call these *transformal discrete valuation rings.* Picking any $t \in L$, $\mathrm{val}\,(t) > 0$, we can view L as an extension of $F(t)_\sigma$ of finite transcendence degree.

Let $\mathbb{Z}_\sigma = \mathbb{Z}[\sigma]$, $\mathbb{Q}_\sigma = \mathbb{Q}[\sigma]$, viewed as ordered $\mathbb{Z}[\sigma]$-modules, (with $\mathbb{Q} < \mathbb{Q}\sigma < \ldots$). $\mathbb{Z}_\sigma$,$\mathbb{Q}_\sigma$ are the value groups of $F(t)_\sigma$, $(F(t)_\sigma)^a$.

In this section, all transformal valued fields are assumed to be ω-increasing, with value group contained in $\mathbb{Q}_\sigma$.

Let L be a transformal valued field. .

Lemma 7.4 *Let L^h be the Henselization of L as a valued field. The endomorphism σ of L lifts uniquely to a valued field endomorphism of L^h. If L is ω-increasing, so is L^h.*

Proof $\sigma : L \to \sigma(L)$ is an isomorphism of valued fields; by the universal property of the Henselization, for any Henselian valued field M containing $\sigma(L)$, σ extends uniquely to an embedding $L^h \to M$; in particular, with $M = L^h$, σ extends to $\sigma : L^h \to L^h$. The property of being ω-increasing depends on the value group, which does not change. □

One can also canonically define a *transformal Henselization* of a transformal valued field, using difference polynomials and their derivatives; cf. the remarks following 7.27, as well as 8.4. Since we will only work with the transformal analogue of discrete valuation rings, we will be able to use the somewhat softer notion of topological closure.

Completion and closure Let L be a transformal valued field, with value group contained in $\mathbb{Q}_s i$. We define a topology on L^h, a basic open set being a ball of nonzero radius. This topology is in general incompatible with valued field extensions. However, according to lemma 7.5, all the valued field extensions we will consider here have cofinal value groups; so the inclusions of valued fields are continuous, and the induced topology on a subfield coincides with the intrinsic one.

We can construct the completion $\hat{L}$ of L^h for this topology; an element of $\hat{L}$ is represented by a sequence a_n of elements of L^h, with $\mathrm{val}\,(a_{n+1} - a_n) \geq \sigma^n(v)$ for some $v \in \mathrm{val}\,, v > 0$. This completion carries a natural ω-increasing transformal valued field structure.

If $K \leq L$ are transformal discrete valuation rings over F, the topological closure of K^h within $\hat{L}$ can be identified with $\hat{K}$.

The Henselization and completion processes do not change the value group or residue field. By Lemmas 7.5, 7.13, the residue field is an extension of F of finite transcendence degree, and the value group is a $\mathbb{Z}[\sigma]$-submodule of $\mathbb{Z}_\sigma$.

7.3 The value group

Lemma 7.5 *Let L be an ω-increasing transformal valued field of transformal dimension 1 over F. Then the value group $\mathrm{val}\,(L^a)$ of L^a is isomorphic to $\mathbb{Q}_\sigma$, as an ordered $\mathbb{Z}[\sigma]$-module. Let $K \leq L$ be a difference subfield, nontrivially valued. Then $\mathrm{val}\,(K)$ is cofinal in $\mathrm{val}\,(L)$. $\mathrm{val}\,(L^a)/\,\mathrm{val}\,(K^a)$ is a principal, torsion $\mathbb{Q}[\sigma]$-module. If L is a finitely generated difference field extension of F, then $\mathrm{val}\,(L)$ is isomorphic to a $\mathbb{Z}[\sigma]$-submodule of $\mathbb{Z}_\sigma$.*

Proof Pick $t \in K$, $\mathrm{val}\,(t) > 0$. Since $tr.deg._{F(t)_\sigma} L < \infty$, $\mathrm{val}\,(L^a)/\,\mathrm{val}\,(F(t)_\sigma{}^a)$ is a finite-dimensional $\mathbb{Q}$-space. Thus $\mathrm{val}\,(L^a)$ is a finitely generated $\mathbb{Q}[\sigma]$-module. Since $\mathbb{Q}[\sigma]$ is a principal ideal domain, and $\mathrm{val}\,(L^a)$ is torsion free, $\mathrm{val}\,(L^a)$ is a free $\mathbb{Q}[\sigma]$-module. Now the quotient $\mathrm{val}\,(L^a)/\,\mathrm{val}\,(F(t)_\sigma{}^a)$ is finite-dimensional, so the rank of $\mathrm{val}\,(L^a)$ must equal one. Let $h : \mathbb{Q}_\sigma \to \mathrm{val}\,(L^a)$ be an isomorphism of $\mathbb{Z}[\sigma]$-modules. Replacing h by $-h$ if necessary, we may assume $h(1) > 0$. It follows that h is order preserving. Any nonzero $\mathbb{Z}[\sigma]$ - submodule of $\mathbb{Q}_\sigma$ is cofinal, and co-torsion.

There remains the last statement, concerning finite generation. Let $t \in L$, $\mathrm{val}\,(t) > 0$, $K = F(t)_\sigma$. Then $\mathrm{val}\,K = \mathbb{Z}[\sigma]$, and L is a finitely generated difference field extension of K of finite transcendence degree. Let L_0 be a subfield of L, finitely generated over K as a field, generating L as a difference field, and of the same K-transcendence degree as L. Let L_n be the subfield of L generated by $\cup_{k \leq n} \sigma^k(L_0)$. Then $[L_{n+1} : L_n] \leq [L_1 : L_0] < \infty$. Let $A_n = \mathrm{val}\,L_n$. Then $A_0 \subset A_1 \subset \ldots \subset \mathrm{val}\,L^a \simeq \mathbb{Q}[\sigma]$, A_0 is finitely generated, and A_{n+1}/A_n is bounded. By lemma 7.7 below, $\mathrm{val}\,L$ is a finitely generated $\mathbb{Z}[\sigma]$-module. And by 7.8, every finitely generated $\mathbb{Z}[\sigma]$ submodule of $\mathbb{Q}[\sigma]$ is contained in a free $\mathbb{Z}[\sigma]$-module of rank one. □

In particular, there are no "irrational" values in $\mathrm{val}\,(L)$; for any $0 < u < v \in \mathrm{val}\,(L)$, for some (unique) $0 \leq q \in \mathbb{Q}$, $|qv - u| << v$. Thus:

Corollary 7.6 *Let L be an ω-increasing transformal valued field of transformal dimension 1 over F, $t \in L$, $\operatorname{val}(t) > 0$. Then $rk_{ram}(L/F(t)_\sigma) = rk_{\mathbb{Q}} \operatorname{val}(L)/\operatorname{val} F(t)_\sigma$.*

Lemma 7.7 *Let $A_0 \subseteq A_1 \subseteq \ldots$ be finitely generated subgroups of $\mathbb{Q}[T]$ containing $\mathbb{Z}[T]$, $TA_n \subseteq A_{n+1}$, with A_{n+1}/A_n finite and bounded. Then $A = \cup_n A_n$ is a finitely generated $\mathbb{Z}[T]$-module.*

Proof If $A/\mathbb{Z}[T]$ is finitely generated, so is A, thus we may replace A_i by $A_i/\mathbb{Z}[T] \leq \mathbb{Q}[T]/\mathbb{Z}[T]$.

We use induction on $m = \limsup_n |A_{n+1}/A_n|$. If $A \neq 0$, there exists $c \in A$, $c \neq 0$, $lc = 0$, l prime. Say $c \in A_{n_0}$.

Claim Let $E = E_l = \{x : lx = 0\}$. Then for $n \geq n_0$, $E \cap A_{n+1} \not\subseteq A_n$.

Suppose the claim is false; then any $x \in E \cap A_n$ satisfies $Tx \in E \cap A_{n+1} = E \cap A_n$. Thus $\mathbb{Z}[\sigma]c \subset A_n$. But A_n is a finitely generated group, while $\mathbb{Z}[\sigma]c$ is not: $c, c^\sigma, \ldots$ are $\mathbb{Z}/l\mathbb{Z}$ -linearly independent. A contradiction.

Thus the natural surjective map $e_l : A_{n+1}/A_n \to lA_{n+1}/lA_n$, $e_l(x) = lx$, is not injective, for $n \geq n_0$. So $\limsup_n |lA_{n+1}/lA_n| < m$. By induction, lA is a finitely generated $\mathbb{Z}[\sigma]$-module.

Now $E_l/(\mathbb{Z}[\sigma]c)$ is finite; indeed $E_l \simeq (\mathbb{Z}/l\mathbb{Z})[T]$, so $E_l/(\mathbb{Z}[\sigma]c) \simeq (\mathbb{Z}/l\mathbb{Z})[T]/f$ for some nonzero polynomial $f(T)$. Thus $(A \cap E_l)/(\mathbb{Z}[\sigma]c)$ is finite. Since $\mathbb{Z}[\sigma]c$ is a finitely generated $\mathbb{Z}[\sigma]$-module, so is $A \cap E_l = \ker(x \mapsto lx)$. We saw lA was finitely generated; hence so is A. □

Lemma 7.8 *Let M be a finitely generated $\mathbb{Z}[\sigma]$-submodule of $\mathbb{Q}_\sigma$. Let $\widetilde{M}$ be the union of all submodules N of $\mathbb{Q}_\sigma$ with N/M finite. Then $\widetilde{M}$ is 1-generated.*

Proof If M is generated by $a_1, \ldots, a_n$, $m_i a_i \in \mathbb{Z}_\sigma$ for some $0 < m_i \in \mathbb{N}$; so each $a_i \in \mathbb{Z}_\sigma[1/m]$, where $m = \Pi_i m_i$. Thus $M \subset N$ for some principal $N \leq \mathbb{Q}_\sigma$.

Claim For any ideal J of $\mathbb{Z}[\sigma]$, there exists a principal ideal J' containing J with J'/J finite.

Proof J is k-generated for some finite k. The claim reduces inductively to the case $k = 2$. So say $J = \mathbb{Z}[\sigma](a, b)$. Since $\mathbb{Z}[\sigma]$ is a unique factorization domain, we may write $a = a'c, b = b'c$ with no prime element of $\mathbb{Z}[\sigma]$ dividing both a' and b'. As $\mathbb{Z}[\sigma]$ has Krull dimension 2, $\mathbb{Z}[\sigma]/(a', b')$ must be finite.

Let $J' = \mathbb{Z}[\sigma]c$. Then a', b' annihilate J'/J, so as J'/J is a quotient of $\mathbb{Z}[\sigma]/(a', b')$, and hence is also finite.

Actually, J' is unique:

Claim If $N \leq \mathbb{Q}_\sigma$ is principal, then $\mathbb{Q}_\sigma/N$ has no finite nonzero $\mathbb{Z}[\sigma]$-submodules. For otherwise it would have a finitely generated one, so again one contained in a principal $\mathbb{Z}[\sigma]$-module. Thus it suffices to show that if J is a submodule of $\mathbb{Z}[\sigma]$ containing a principal module K, then J/K is zero or infinite. If J/K is nonzero and finite, so is J'/K, where J'/J is finite and J' is principal. So $\mathbb{Z}[\sigma]/f\mathbb{Z}[\sigma]$ is finite nonzero; but this is clearly absurd.

Thus $\widetilde{M} \subset N$. Let $I = \{r \in \mathbb{Z}[\sigma] : rN \subset \widetilde{M}\}$. So $\widetilde{M} = IN$. By definition of $\widetilde{M}$, if I'/I is finite then $I' = I$. Thus by the Claim I is principal. So $\widetilde{M}$ is 1-generated.

□

Example 7.9 $\mathrm{val}\,(L)$ *need not itself be free.*

Take $F(t^\sigma, t^2)_\sigma \leq F(t)_\sigma$; the value group is $M = \mathbb{Z}[2v, \sigma(v)] \subset \mathbb{Z}[v, \sigma(v), \ldots]$. We have $M/\sigma(M) \simeq \mathbb{Z} \oplus (\mathbb{Z}/2\mathbb{Z})[\sigma]$. □

At all events, when L is finitely generated, of limit degree d, $\mathbb{Z}[1/d!]\,\mathrm{val}\,(L)$ is already free on one generator.

This suggests that L may have an algebraic extension whose value group is a free $\mathbb{Z}[\sigma]$-module of rank one. We prove this statement in characteristic 0; in positive characteristic $p > 0$, the proof works up to localization at p, and we did not investigate it further.

Proposition 7.10 *Let F be a difference field of characteristic* 0*, L a transformal discrete valuation ring over F. Then L has a finite σ-invariant extension M, whose value group is isomorphic to $\mathbb{Z}[\sigma]$.*

Proof

By 7.8, $\mathrm{val}\,L \subset \widetilde{M}$ for some free 1-generated $\mathbb{Z}[\sigma]$-module $\widetilde{M}$ with $\widetilde{M}/\,\mathrm{val}\,L$ finite; we have $\mathrm{val}\,L = I\widetilde{M}$ for some ideal I, and necessarily $n \in I$ for some $n > 0$.

We may assume F has n'th roots of unity; adding them will not change the value groups.

Consider the surjection

$$val : L^* \to> \mathrm{val}\,L$$

It induces a surjection

$$(L^*)^n \to> n\,\mathrm{val}\,L$$

Let H be the pullback of $n\widetilde{M} \subset \mathrm{val}\,L$. We have

$$H \to> n\widetilde{M}$$

and $\sigma(H) \subset H$. Also $H/(L^*)^n \simeq n\widetilde{M}/n\,\mathrm{val}\,L \simeq \widetilde{M}/\,\mathrm{val}\,(L)$. So $H/(L^*)^n$ is finite. By Kummer theory, there exists a (unique) Galois extension M of L, such that $(M^*)^n \cap L = H$, and $[M : L] = [H : (L^*)^n] = \widetilde{M}/\,\mathrm{val}\,L$. Tracing back the isomorphisms we see that $\widetilde{M} \subset \mathrm{val}\,(M)$, so $\widetilde{M} = \mathrm{val}\,(M)$. □

Notation 7.11 $\mathcal{H}(X) = \{y : (\exists x \in X)(y \leq x)\}$

Remark 7.12 *Let $M \subset \mathbb{Z}_\sigma$ be a $\mathbb{Z}[\sigma]$-module, and let $Y = \mathcal{H}(Y) \subset M$, $0 \in Y$. Then for some nonzero convex subgroup S of M, and $0 \leq a \in M$,*

$$Y = \mathcal{H}(a + S) = \{y : (\exists c \in S) y \leq a + c\}$$

The nonzero convex subgroups form a single σ-orbit $\{E_n : n = 1, 2, \ldots\}$ (in the sense that E_{n+1} is the convex hull of $\sigma(E_n)$; $E_n = \sigma^{-1}(E_{n+1})$.)

Proof Let $(0) = C_0 \subset C_1 \subset \ldots$ be the proper convex subgroups of $\mathbb{Z}_\sigma$ (with $\sigma^{-1}(C_{n+1}) = C_n$). Then for some m, $M \cap C_{m+1} \neq (0)$, $M \cap C_m = (0)$. So for $n \geq 0$, $E_n = C_{n+m} \cap M$ are convex subgroups of M. If E is any proper convex subgroup of M, then the convex hull of E in $\mathbb{Z}_\sigma$ must be some C_{n+m}, and it follows that $E = E_n$.

If $c \in E_1 \setminus E_0$, then $\sigma^n(c) \in E_{n+1} \setminus E_n$. Now $C_{n+1}/C_n \simeq \mathbb{Z}$, so E_{n+1}/E_n is isomorphic to a nonzero subgroup of $\mathbb{Z}$, hence also isomorphic to $\mathbb{Z}$.

Let $S = \{a : a + Y \subset Y\}$. This is a convex subgroup of M; so $S = E_n$ for some $n \geq 0$.

There exists an E_{n+1}-class Z with $Z \cap Y \neq \emptyset$ and $Z \not\subseteq Y$. Now Z/E_n is order-isomorphic to $\mathbb{Z}$. If $(Z \cap Y)/E_n$ has no greatest element, then $Z \cap Y$ is cofinal in Z so $Z \subset Y$, a contradiction. Similarly, Y can have no element above $Z \cap Y$. Thus $(Z \cap Y)/E_n$ has a greatest element a/E_n, and this is also a greatest element of Y/E_n. It follows that $Y = \mathcal{H}(a + E_n)$, □

7.4 The residue field

Lemma 7.13 *Let L be an ω-increasing transformal valued field of transformal dimension 1 over F. The residue field $\operatorname{res}(L)$ of L has finite transcendence degree over F.*

Proof The residue field of $F(t)_\sigma$ is F itself, and L has finite transcendence degree over $F(t)_\sigma$. □

The analogy with 7.5 raises the question, that I did not look into: If in 7.13 $L/F(t)_\sigma$ is finitely generated, is $\operatorname{res} L$ finitely generated over $\operatorname{res}(K)$ as a difference field (up to purely inseparable extensions)?

7.5 Valued field lemmas

We will need an observation from the theory of valued fields. (Compare [Haskell-Hrushovski-Macpherson] Part I, §2.5 ("Independence and orthogonality for unary types".))

Let $K \leq L$ be an inclusion of valued fields, $c \in K$. We let

$$T(c/K) = \{val(c - b) : b \in K\}$$

Also let $E(c/K)$ be the stabilizer of $T = T(c/K)$ in $\operatorname{val}(K)$, i.e.

$$E = \{e \in \operatorname{val} K : e + T = T\}$$

If $T(c/K)$ has no greatest element, then (7.14) $\operatorname{val} K(c) = \operatorname{val} K$. In this case $T(c/K)$ is a downwards-closed subset of $\operatorname{val} K$. Hence E is a convex subgroup of $\operatorname{val} K$.

Lemma 7.14 *Let K be an algebraically closed valued field, L an extension of transcendence degree 1. Let $c \in L \setminus K$. Then the following conditions are equivalent:*

1. $\operatorname{res}(K) \neq \operatorname{res}(L)$ *or* $\operatorname{val}(K) \neq \operatorname{val}(L)$
2. $\operatorname{res}(K(c)) \neq \operatorname{res}(L)$ *or* $\operatorname{val}(K(c)) \neq \operatorname{val}(L)$
3. $T(c/K) = \{val(c - b) : b \in K\}$ *has a maximal element.*

In particular, the third condition depends on K, L alone and not on the choice of c.

$T(c/K) = \operatorname{val} K$ *iff L embeds into the completion $\widehat{K}$ as a valued field.*

Proof : (1) implies (2) since if $\operatorname{res}(K(c)) = \operatorname{res}(K)$, and $\operatorname{val}(K(c)) = \operatorname{val}(K)$, then $\operatorname{res}(K(c))$ is algebraically closed and $\operatorname{val}(K(c))$ is divisible; so they cannot change under algebraic field extensions. Now assume $\mathrm{val}\,(K(c)) \neq \mathrm{val}\,(L)$. Then $\mathrm{val}\, f(c) \notin \mathrm{val}\,(K)$ for some $f \in K[X]$. Splitting f into linear factors, wee see that we can take it to be linear. So $\mathrm{val}\,(c-b) \notin \mathrm{val}\,(K)$ for some $b \in K$. It follows that $\mathrm{val}\,(c-b') \leq \mathrm{val}\,(c-b)$ for all $b' \in K$. (Otherwise $\mathrm{val}\,(c-b) = \mathrm{val}\,(b-b')$.) Next suppose $\mathrm{val}\,(K(c)) = \mathrm{val}\,(L)$, but $\mathrm{res}\, f(c) \notin \mathrm{res}\,(K)$ for some $f \in K[X]$ with $f(c) \in \mathcal{O}_K$. Taking into account $\mathrm{val}\,(K(c)) = \mathrm{val}\,(K)$, we can split f into linear factors f_i such that each $f_i(c) \in \mathcal{O}_K$. Thus $ac - b \in \mathcal{O}_K$, $\mathrm{res}\,(ac-b) \notin \mathrm{res}\,(K)$ for some $a, b \in K$. It follows that $\mathrm{val}\,(ac-b') \leq 0$ for all $b' \in K$. So $\mathrm{val}\,(c-b/a) \geq \mathrm{val}\,(c-b'')$ for all $b'' \in K$. This proves that (2) implies (3).

Now assume (3): $\mathrm{val}\,(c-b) \geq \mathrm{val}\,(c-b')$ for all $b' \in K$. If $\mathrm{val}\,(c-b) \notin \mathrm{val}\,(K)$, (2) holds. If $\mathrm{val}\,(c-b) = \mathrm{val}\,(d)$, $d \in K$, then $(c-b)/d \in \mathcal{O}_K$, and $\mathrm{val}\,((c-b)/d - d') \leq 0$ for all d'; so $\mathrm{res}\,((c-b)/d) \notin \mathrm{res}\,(K)$. Thus (1). □

Example 7.15 *Let F be a field, $K = F(t_0, t_1, \ldots)^a$, valued over F with $0 < \mathrm{val}\,(t_0) << \mathrm{val}\,(t_1) << \cdots$. Let $K' = F(t_e, t_{e+1}, \ldots)^a$, and let $c \in \widehat{K} \setminus \widehat{K'}$. Then $T(c/K') = \{\, \mathrm{val}\,(c-b) : b \in K'\}$ has a maximal element.*

Proof If $T(c/K')$ is unbounded in $\mathrm{val}\,(K')$, then $c \in \widehat{K'}$. Otherwise, for all $b \in K'$, $\mathrm{val}\,(c-b) \leq \alpha$. Find $c' \in K$ with $\mathrm{val}\,(c-c') > \alpha$. Then $T(c'/K') = T(c/K')$. Now $K^a/(K')^a$ is generated (as an algebraically closed field) by e elements, whose values are $\mathbb{Q}$-linearly independent over $\mathrm{val}\,(K')$. Thus $\mathrm{val}\, K'(c') \neq \mathrm{val}\, K'$; so 7.14 applies, and shows $T(c'/K')$ has a maximal element. □

Lemma 7.16 *Let K be an algebraically closed valued field, L an extension of transcendence degree 1, $c, d \in L \setminus K$. Then $E(c/K) = E(d/K)$.*

Proof

If one of $T(c/K), T(d/K)$ has a last element, then by 7.14 so does the other, and $E(c/K) = (0) = E(d/K)$. Thus we may assume $T(c/K), T(d/K)$ have no last element, so that $\mathrm{val}\, K = \mathrm{val}\, L =: \Gamma$; and that $E(c/K) \subseteq E := E(d/K)$.

Special Case $E = \Gamma$.

In this case, by the last statement of 7.14, L embeds into $\widehat{K}$, and hence $E(c/K) = \Gamma$ too.

In general, let $\Gamma' = \Gamma/E$, $r : \Gamma \to \Gamma'$ the quotient map, and let $\mathrm{val}\,' : L \to \Gamma'$ be the induced valuation. Note that $\mathrm{val}\,'(L) = \mathrm{val}\,'(K) = \Gamma'$. Call the residue fields K', L'. Let $T'(y/K) = \{val'(y-b) : b \in K\}$.

If $E(c/K) \neq E(d/K)$, then $T'(c/K)$ has a last element. (Indeed let $\gamma \in E(d/K)$ with $\gamma > 0$ and $\gamma \notin E(c/K)$. Then there exists $\alpha \in T(c/K)$ with $\alpha + \gamma \notin T(c/K)$. Let α' be the common image of $\alpha, \alpha + \gamma$ in Γ'. Then clearly α' is the greatest element of $T'(c/K)$.) By 7.14, $T'(d/K)$ has a greatest element too. Effecting additive and multiplicative translations, we may assume these greatest elements are $\mathrm{val}\,'(c) = 0$, $\mathrm{val}\,'(d) = 0$. Let c', d' be the residues of c, d.

Since $\mathrm{val}\,'(L) = \mathrm{val}\,'(K) = \Gamma'$, by 7.14 we must have $K' \neq L'$. We have an induced valuation $\mathrm{val}\,'' : L' \to E$, $\mathrm{val}\,''(\mathrm{res}\,'(a)) = \mathrm{val}\,(a)$ for a with $\mathrm{val}\,(a) \in E$. Clearly $T''(c'/K') = \{\, \mathrm{val}\,''(c' - b' : b' \in K')\} = \mathcal{H}T(c/K)$, and similarly for d; thus $E''(c'/K') =$

$E(c/K), E''(d'/K') = E(d/K)$. But now $T''(d/K)$ is the entire value group $E(d/K)$, so by the special case, $E''(c'/K') = E''(d'/K')$. □

Corollary 7.17 *Let $K' = K(c)^a, K'' = K(d)^a$ be two valued field extensions of an algebraically closed valued field K, with $E(c/K) \neq E(d/K)$. Then K', K'' have a unique valued field amalgam.*

We have $T(c/K'') = \mathcal{H}T(c/K)$.

In particular if $\operatorname{val} K'' = \operatorname{val} K$ then $T(c/K'') = T(c/K)$, and $E(c/K'') = E(c/K)$.

Proof Assume K', K'' are embedded in $L = K'K''$. Let $a \in K', b \in K''$. Then $E(a/K) = E(c/K) \neq E(d/K) = E(b/K)$, so $T(a/K) \neq T(b/K)$. One of $T(a/K), T(b/K)$ is bigger, say $T(a/K) \subset T(b/K)$. Find $e \in K$ with $\operatorname{val}(b-e) > T(a/K)$. Then $\operatorname{val}(a-b) = \operatorname{val}(a-e)$. This shows that the values of elements $a - b$ are all determined. In particular the values $\operatorname{val}(c-b)$ for $b \in K(d)$ are determined; this determines $tp(c/K'')$ and hence $tp(K'/K'')$.

Moreover we see that every element $\operatorname{val}(c-a)$ of $T(c/K'')$ either equals $\operatorname{val}(c-e)$ for some $e \in K$, so that it lies in $T(c/K)$, or else equals $\operatorname{val}(a-e)$ where $\operatorname{val}(c-e) > \operatorname{val}(a-e)$; so in either case it lies in $\mathcal{H}T(c/K)$. □

This proof uses the stationarity theorem of [Haskell-Hrushovski-Macpherson] (cf. Theorem 2.11); alternatively one can prove the corollary directly, along the lines of the proof of 7.16.

Proposition 7.18 *Let L be a valued field extension of an algebraically closed valued field K. Let $a_i \in L \setminus K$,*

$$T_i = T(a_i/K) = \{val(a_i - b) : b \in K\}$$

Assume that the convex subgroups

$$E_i = E(a_i/K) = \{e \in \operatorname{val} K : e + T_i = T_i\}$$

are distinct. Then $a_1, \ldots, a_n$ are algebraically independent over K. The valued field structure of $K(a_1, \ldots, a_n)$ is uniquely determined by the T_i.

Proof By assumption, $a_i \notin K^{alg}$. Use induction on n. Some E_i is nonzero; say $E_1 \neq (0)$. Let $K' = K(a_1)^a$. By 7.14, $\operatorname{val} K' = \operatorname{val} K$. By 7.16, a_i, a_1 are algebraically independent over K for $i \neq 1$. By 7.17, the type of $(a_i, a_1)/K$ is determined, and $E(a_i/K') = E(a_i/K)$. By induction, $a_2, \ldots, a_n$ are independent over K', and their type over K' is determined. The conclusion follows. □

The assumption $K = K^a$ in 7.18 can be weakened to: K is perfect Henselian. Indeed:

Remark 7.19 *Let L be an immediate valued field extension of a perfect, Henselian valued field K. If $c \in L \setminus K$ then $T(c/K^a) = \mathcal{H}T(c/K)$; hence $E(c/K^a) = \mathbb{Q}E(c/K)$.*

Proof Let P be the intersection of all K-definable balls containing c. It suffices to show that there is no $b \in K^a$ with $\operatorname{val}(b-c) > T(c/K)$. Otherwise some nonzero separable polynomial f over K has a root in P. Some derivative of f has a single simple root in P. By the Hensel property, this root must lie in K; but $P(K) = \emptyset$. □

Let L be a model, and let Let $F \leq K_1, \dots, K_n \leq L$. We say that K_1, K_2 are almost-orthogonal over F if for any $a = (a_1, \dots, a_n)$, a_i a tuple of elements of K_i, $\cup_i tp(a_i/F)$ implies $tp(a/F)$.

In the case of an algebraically closed valued field L, K_1, K_2 are almost-orthogonal over F if whenever $f_i : K_i \to M$ are embeddings of valued fields, with $f_1|F = F_2|F$, there exists a valued field embedding $f : K \to M$, with $f|K_i = f_i$.

In this language, 7.18 asserts the almost-orthogonality of certain unary types.

Let F, K, L, K', L denote algebraically closed valued fields. $\widehat{F}$ denotes the completion of F.

When $K \leq L$ is an extension of valued fields, $K = K^a$, and $tr.deg._K L = tr.deg._{K_{\mathrm{res}}} L_{\mathrm{res}}$, we will say that L/K is *purely inertial.* If instead $tr.deg._K L = rk_{\mathbb{Q}}(\mathrm{val}(L)/\mathrm{val}(K)$, we will say that L/K is *purely ramified.*

We will use Theorem 2.11 of [Haskell-Hrushovski-Macpherson], or rather the following corollary: if $F(a), K$ are almost-orthogonal over $F = F^a$, then so are $F(a)^a, K$. (Note that this is immediate for perfect Henselian closure in place of algebraic closure.)

Lemma 7.20 *Let $F \leq K_1, K_2 \leq \widehat{F}$ be algebraically closed valued fields. Then K_1, K_2 are almost-orthogonal over F if (and only if) they are linearly disjoint over F.*

Proof We may assume here that $tr.deg._F K_1 < \infty$. We have $\widehat{F} = \widehat{L}$ for any intermediate field $F \leq L \leq \widehat{F}$; so by the transitivity of the two notions, the lemma reduces inductively to the case $tr.deg._F K_1 = 1$. Let $a \in K_1 \setminus F$. Then $a \notin K_2$. Let B be a K_2-definable ball with $a \in B$. So $B = B_\rho(b) = \{y : \mathrm{val}(y - b) \geq \rho\}$, $\rho \in \mathrm{val}(K_2), b \in K_2$. As $a \notin K_2$, $\rho < \infty$. Since $b \in \widehat{F}$ and $\rho \in \mathrm{val}(\widehat{F}) = \mathrm{val}(F)$, there exists $b' \in F$ with $\mathrm{val}(b - b') > \rho$. So B is F-definable. Thus $tp(a/F)$ implies $a \in B$. Since B was arbitrary, and K_2 is algebraically closed, the formulas $x \in B$ of this kind generate $tp(a/K_2)$. Thus $tp(a/F)$ implies $tp(a/K_2)$. So $F(a), K_2$ are almost orthogonal over F. By Theorem 2.11 of [Haskell-Hrushovski-Macpherson], K_1, K_2 are almost-orthogonal over F. □

Lemma 7.21 *Let $K = K_0 \subset K_1 \subset \cdots \subset K_n = L$ be algebraically closed valued fields. Assume: $\mathrm{val}(K_0)$ is cofinal in $\mathrm{val}(K_n)$; and for each i, K_{i+1}/K_i is purely ramified, or purely inertial, or $K_{i+1} \subset \widehat{K_i}$. Assume also $\bar{K}$ is an immediate extension of K, generated by elements c with $T(c;K)$ bounded. Then $\hat{L}, \bar{K}$ are almost orthogonal over K.*

Hence $\hat{L}$ is dominated by L_{res} over $\hat{K} \cup \mathrm{val}(L)$.

Proof The domination follows from the almost orthogonality by [Haskell-Hrushovski-Macpherson], Theorem 6.13, taking $\bar{K}$ to be a maximal immediate extension of $\hat{K}$.

To prove the almost-orthogonality, consider first the case $n = 1$; we will also show that $M = L\bar{K}$ is an immediate extension of L. If L/K is purely ramified, or purely inertial, the almost orthogonality is clear. So is the fact that $M = L\bar{K}$ is an immediate extension of L:

$$tr.deg._{K_{\mathrm{res}}} M_{\mathrm{res}} + rk_{\mathbb{Q}} \mathrm{val}(M)/\mathrm{val}(K) \leq tr.deg.\bar{K}M \leq tr.deg._K L = tr.deg._{K_{\mathrm{res}}} L_{\mathrm{res}} + rk_{\mathbb{Q}} \mathrm{val}(L)/\mathrm{val}(K)$$

so the numbers are equal, and thus $M_{\text{res}} = L_{\text{res}}$, $\text{val}(L) = \text{val}(M)$. The remaining case is that $L \subset \hat{K}$. This follows from 7.18.

For $n > 1$, we can use induction: assume $\bar{K}$ is almost-orthogonal to K_i over K, and $\bar{K}K_i$ is an immediate extension of K_i; show the same is true for $i+1$. This follows from what has been proved, for K_i, K_{i+1} in place of K, L. □

Lemma 7.22 *Let k be an algebraically closed field, $K = k(t)^a$. Let v be a valuation of K/k, R the valuation ring. Let X be a quasi-projective scheme over R, with $X(K)$ and $X(k)$ finite. Consider the residue homomorphism $R \to k$, and the induced map $r : X(R) \to X(k)$. If $p \in X(k)$ and $r^{-1}(p)$ has n distinct points, then p has geometric multiplicity $\geq n$ on the scheme $X_0 = X \otimes_R k$. More generally, this holds true if $r^{-1}(p)$ has n points counted with multiplicities on $(X \otimes_R K)$.*

Proof Find an affine open subscheme containing the said n points $q_1, \ldots, q_n$, and thus reduce to $X = \text{Spec}^\sigma A$, A a finitely generated R-algebra. Let B be the image of A in $A \otimes_R K$. As X is 0-dimensional, $A \otimes_R K$ is a finite dimensional K-space, and so the finitely generated R-submodule B is a free R-module of finite rank. It follows that $\dim_k(A \otimes_R k) \geq \dim_k(B \otimes_R k) = \dim_K(B \otimes_R K)$. □

7.6 Structure theory

Example 7.23 *Let K be a transformal discrete valuation ring, with value group $\mathbb{Z}[\sigma]$, and residue field $F(a)_\sigma$, where $g(a) = 0$ for some difference polynomial g while $g'(a) \neq 0$. Then there exists a 1-generated transformal discrete valuation rings over F, with the same value group and residue field.*

To see this, we may at first assume K is Henselian. pick t such that $\text{val}(t)$ generates the value group. Lift g to $\mathcal{O}_K(x)$, and lift a to $\mathcal{O}_K$. Then $\text{val}(g(a)) > 0$, so $\text{val}((g-t)(a) > 0)$. By one step of Hensel's lemma, we can perturb a so that $\text{val}((g-t)(a)) > \text{val}(t)$. Then $\text{val}\, g(a) = \text{val}(t)$. So $F(a)_\sigma$ has the same residue field and value group as K.

Example 7.24

1. Let F be a valued field. Form $F(t)_\sigma$, and let L be the completion. Let $a = \sum_{n=0}^{\infty} t^{\sigma^n} \in L$, $K = F(t,a)_\sigma$. We have $\sigma(a) - a - t = 0$.
2. Let $y^\sigma - y = t^{-1}$. Then $(F(t)_\sigma)^{inv}(y)$ is an immediate extension of $(F(t)_\sigma)^{inv}$:
$$y = t^{-1/\sigma} + t^{-1/\sigma^2} + \ldots$$
$F(t,y)_\sigma$ is a "ramified" extension of $F(t)_\sigma$, in the sense that $rk\, \text{val}(F(t,y)_\sigma)/\text{val}(F(t,y)_\sigma) = 1$.

Proposition 7.25 *Let K be a perfect Henselian transformal valued field with value group $\text{val}(K) \subset \mathbb{Q}_\sigma$. Assume: (*) for any $c \in K$, $\{\text{val}(c-b) : b \in \sigma(K)\}$ has a greatest element.*

Let L be a valued difference field extension of K, with tr.deg.$_K(L) < \infty$, $\text{val}(L) = \text{val}(K)$, and $L_{\text{res}} = K_{\text{res}}$. Then $\widehat{L} = \widehat{K}$.

Proof Let $a \in L$, $a_n = \sigma^n(a)$, $C_n = \{ \operatorname{val}(a_n - b) : b \in K\}$, $C = C_0$. Then C_n is closed downwards in $\operatorname{val}(K)$. C_n cannot have a maximal element, since L is an immediate extension of K.

If C is unbounded, then $a \in \widehat{K}$. Suppose for contradiction that C is bounded. We have $\mathbb{Q}_\sigma = \mathbb{Q}v \oplus \mathbb{Q}v^\sigma + \ldots$; let $E_m = \mathbb{Q}v \oplus \ldots \oplus \mathbb{Q}v^{\sigma^m}$. Let m be least such that $C \subset E_m + c_0$ for some $c_0 \in C$. Since $\operatorname{val}(L) = \operatorname{val}(K)$, dividing a by an element of K with value c_0, we may assume $C \subset E_m$. Then $\sigma(C) \subset E_{m+1}$. Let H_1 be the conex hull of $\sigma(C)$.

Claim $E_m + H_1 = H_1$. *Proof* Let $e \in E_m$, $0 < c \in C$. Since $C \not\subseteq c + E_{m-1}$, there exists $c' \in C$, $l \in \mathbb{N}$ with $c' - c > e/l$. Then $\sigma(c' - c) > l(c' - c) > e$. So $e + \sigma(c) \leq \sigma(c')$.

Claim $H_1 = C_1$ *Proof* Let $c_1 < c_2 < \ldots$ be cofinal in C. Let $b_n \in K$, $\operatorname{val}(a - b_n) = c_n$; let $B_n = \{x : \operatorname{val}(x - b_n) \geq c_n\}$; $P = \cap_{n \in \mathbb{N}} B_n$; ${B_n}^\sigma = \{x : \operatorname{val}(x - \sigma(b_n)) \geq \sigma(c_n)\}$ $P^\sigma = \cap_{n \in \mathbb{N}} {B_n}^\sigma$;

Then $P(K) = \emptyset$. So $P^\sigma(K^\sigma) = \emptyset$. By the assumption (*) on K/K^σ, $P^\sigma(K) = \emptyset$. Since $c_1 \in P^\sigma$, it follows that there can be no $b \in K$ with $\operatorname{val}(b - c_1) > \sigma(C)$. Thus $\sigma(C) = H_1$.

Now σ^n induces an isomorphism $(K, \sigma(K)) \to (\sigma^n(K), \sigma^{n+1}(K))$. So the convex hull of $\sigma^n(C)$ is C_n for all n; $E_{m+n} + C_{n+1} = C_{n+1}$; and $C_n \subset \mathcal{H}(E_{m+n})$. Thus the hypotheses of Lemma 7.18 are valid of the C_m.

But now by 7.18, 7.19, the elements a_n must be algebraically independent over K. This contradicts the assumption of finite transcendence degree. □

The condition (*) is easily seen to apply in the fundamental case: $K = F(t)^a_\sigma$. (See 7.15.)

Example 7.26 *The composition of difference polynomials extends to $F(t)_\sigma$, and by continuity to the completion L of $F(t)_\sigma$. By 7.25, difference polynomials with nonzero linear term have a left compositional inverse in L.*

Difference polynomials of the form $t + t^\sigma h$ are compositionally invertible in L. For instance, the compositional inverse of $t + t^{\sigma+1}$ is the repeated fraction:

$$\frac{t}{1 + \frac{t^\sigma}{1 + \frac{t^{\sigma^2}}{1 + \cdots}}}$$

Transformally Henselian fields

Lemma 7.27 *Let K be a complete, algebraically Henselian transformal valued field over F, with value group $\operatorname{val}(K) \subset \mathbb{Q}_\sigma$. Let $f \in \mathcal{O}_K[X]_\sigma$ be a difference polynomial, $f_\nu = \partial_\nu f$. Suppose $a \in \mathcal{O}_K$, $v = \operatorname{val}(f(a)) > 2 \operatorname{val}(f_1(a)) = 2v'$. Then there exists $b \in K$, $f(b) = 0$, $\operatorname{val}(a - b) \geq v - v'$.*

Proof Let $e_0 = f(a)$. We will find $a_1 \in \mathcal{O}_K$, $\operatorname{val}(a - a_1) \geq v - v'$, $\operatorname{val} f(a_1) \geq \sigma(v)$. Since $\operatorname{val}(f_1(a) - f_1(a_1)) \geq v - v' > v' = \operatorname{val} f_1(a)$, it follows that $\operatorname{val} f_1(a_1) = v'$. Iterating this we get a sequence a_n, with $v_n =_{def} \operatorname{val} f(a_n) \geq \sigma(v_{n-1})$, and $\operatorname{val}(a_{n-1} - a_n) \geq v_n - v'$. As K is complete, there exists a unique $b \in K$ such that the sequence a_n converges to b; by continuity, $f(b) = 0$.

To obtain a_1, write $a_1 = a + re$, with $e = e_0/f_1(a)$, and $r \in \mathcal{O}_K$ to be found. Then

$$f(a_1) = f(a+re) = e_0 + \sum_{i=1}^{m} f_i(a)(re)^i + \sum_{\nu} f_\nu(a)(re)^\nu$$

Here the f_i are transformal derivatives, cf. 2.12 ν runs over indices $\geq \sigma$, while i ranges over the nonzero finite indices. Note that $\operatorname{val} f_\nu(a)(re)^\nu \geq \operatorname{val}(\sigma(e)) = \sigma(v)$. Thus it suffices to find r with $e_0 + \sum_{i=1}^{m} f_i(a)(re)^i = 0$, or with r a root of

$$g(x) = 1 + x + f_2(a)(e^2/e_0)x^2 + \ldots + f_m(a)(e^m/e_0)x^m$$

(Divide through by e_0, noting $f_1(a)(e/e_0) = 1$.) This (ordinary) polynomial has derivative

$$g'(x) = 1 + 2f_2(a)(e^2/e_0)x + \ldots$$

Now for $k \geq 2$, $\operatorname{val}(e^k/e_0) = k(v-v') - v \geq 2(v-v') - v = v - 2v' > 0$. Thus $\operatorname{res} g'$ is a nonzero constant polynomial. By the ordinary Hensel lemma, there exists $r \in \mathcal{O}_K$ with $g(r) = 0$, near any $c \in \mathcal{O}_K$ with $\operatorname{val} g(c) > 0$. Letting $c = -1$ we see that there exists $r \in \mathcal{O}_K$ with $g(r) = 0$. □

We will tentatively call a transformal valued field satisfying the conclusion of 7.27 of characteristic 0 *transformally henselian.* In characteristic $p > 0$, we demand also that the field L be perfect, and that all the Frobenius twists $(L, \sigma \circ (\phi_p)^n)$ also satisfy the conclusion of 7.27. ($\phi_p(x) = x^p$).

Lifting the residue field

Corollary 7.28 *Let $K = K^a \leq L$ be transformal valued fields, with L transformally henselian. Let $f \in K_{\text{res}}[X]_\sigma$ be a difference polynomial of order r, $f' = \partial_1 f$. Let $\bar{a} \in L_{\text{res}}$, $f(\bar{a}) = 0$, $f'(\bar{a}) \neq 0$, tr.deg.$_{K_{\text{res}}} K_{\text{res}}(\bar{a}) = r$. Then there exists a purely inertial extension $K' = K(b)_\sigma$ of K, $K' \leq L$, $\bar{a} \in \operatorname{res}(K')$.*

Proof Pick any $a \in \mathcal{O}_L$ with $\operatorname{res}(a) = \bar{a}$, and also lift f to $\mathcal{O}_K[X]_\sigma$. Then $\operatorname{val} f(a) > 0$, $\operatorname{val} f'(a) = 0$, so by the transformal Hensel property (cf. 7.27) there exists $b \in \mathcal{O}_K$ with $f(b) = 0$, $\operatorname{res}(b) = \bar{a}$. Clearly $K' = K(b)_\sigma$ is purely inertial over K. □

Let K be an ω-increasing transformal valued field over F. The set of elements of K satisfying nontrivial difference polynomials over F forms a difference subfield F' of K; it is the union of all subfields of transformal transcendence degree 0 over F. The residue map res is injective on F', since if $a \in F'$ and $\operatorname{val}(a) > 0$, then $0 < val(a) << \operatorname{val}(\sigma(a)) << \ldots$, so $a, \sigma(a), \ldots$ are algebraically independent over F.

Corollary 7.29 *Let K be an ω-increasing transformally henselian valued field over F. Let $\bar{a} \in K_{\text{res}}$, $f(\bar{a}) = 0$, $f'(\bar{a}) \neq 0$, $f \in F[X]_\sigma$ a difference polynomial, $f' = \partial_1 f$. Then $F(\bar{a})_\sigma$ lifts to K: there exists a subfield $F(b)_\sigma$ of K such that the residue map restricts to an isomorphism $F(b)_\sigma \to F(\bar{a})_\sigma$ (of difference F-algebras.)*

Proof Pick any $a \in \mathcal{O}_K$ with $\operatorname{res}(a) = \bar{a}$, and also lift f to $\mathcal{O}_K[X]_\sigma$. Then $\operatorname{val} f(a) > 0$, $\operatorname{val} f'(a) = 0$, so by 7.27, there exists $b \in \mathcal{O}_K$ with $f(b) = 0$, $\operatorname{res}(b) = \bar{a}$. The difference field $F(b)_\sigma$ has finite transcendence degree m over F, and hence contained in F'; thus it is trivially valued, so the residue map is injective on $F(b)_\sigma$. since $\mathrm{res}\,(b) = \bar{a}$, it carries $F(b)_\sigma$ into and onto $F(\bar{a})_\sigma$. □

Proposition 7.30 *Let K be a perfect ω-increasing transformally henselian valued field over a difference field F.*

Assume $F^{inv} \subset F^a$, and K_{res} has transformal transcendence degree 0 over F. Then $\mathrm{res} : F' \to K_{res}$ *is surjective.*

Proof By Lemma 5.1, $(F')^{inv} \subset (F')^a$; and clearly F' is perfect. We may assume $F' = F$, and show that $\mathrm{res}\,(K) = F$. Suppose $K_{\mathrm{res}} \neq F$. In characteristic 0, choose a difference polynomial $f \in F[X]_\sigma$ of smallest possible transformal order and degree, and $\bar{a} \in K_{\mathrm{res}}$ with $f(\bar{a}) = 0$. If $f \in F[X^\sigma]_\sigma$, then replacing $\bar{a}$ by $\sigma(\bar{a})$ we may lower the degree of f; unless $\sigma(\bar{a}) \in F$. But in this case, $\bar{a} \in F^a$, so the order of f must be 0, so $f \in F[X^\sigma]$ is impossible. Thus $f \notin F[X^\sigma]$. It follows that $f' \neq 0$ and f' has smaller order or degree. So the hypothesis of 7.29 is met, and we can increase F'.

In positive characteristic, this may not be the case, because of polynomials such as $x^\sigma - x^p$. However, after replacing σ by $\tau = \sigma(x^{1/p^m})$ for large enough m, it is possible to find such an $\bar{a}$ for τ, lowering the degree of f further. Assuming the residue field of finite total dimension, this permits lifting the residue field; we may then return from τ to σ. □

Remark 7.31

In positive characteristic, when K is a transformal discrete valuation ring over F, the Proposition applies to the perfect closure of K. But we can lift K_{res} to a subfield of the completion of a smaller extension L of $\hat{K}$ within the perfect closure, such that L is still the completion of a transformal discrete valuation ring, and in particular $\mathrm{val}\,(L) \subset \mathbb{Z}_\sigma$; the details are left to the reader.

Example 7.32

The assumption in 7.30 that F^a is inversive is necessary: take $F = \mathbb{Q}(b, b^\sigma, \ldots), K = F(t, c), \sigma(c) = b + t$.

Characterizing the completion Here is one version of a theorem characterizing completions of algebraically closed valued fields of transformal dimension one over a difference field F.

Lemma 7.33 *Let $L = L^a$ be a transformal valued field over $F = F^{inv}$, with value group $\subset \mathbb{Q}_\sigma$. Assume L is complete. Then the following conditions are equivalent:*

1. *$L = \widehat{L_0}$ for some transformal valued field L_0 of transformal dimension 1 over F.*
2. *Whenever $L' = (L')^a \leq L$ is complete with the same value group and residue field, $L' = L$.*
3. *$L \simeq \widehat{K}$, where $K = (L_{\mathrm{res}}\,(t)_\sigma)^a$.*

Proof Assume(1), and let L' be as in (2). By 7.13, L has value group $\mathbb{Q}_\sigma$, and L_{res} has transformal transcendence degree 0 over F. So L_{res} is inversive.

Let $t \in L'$ be such that $\text{val}(t)$ generates $\text{val}(L)$ as a $\mathbb{Q}[\sigma]$-module. By 7.30, L_{res} lifts to a trivially valued difference subfield F' of L'. Let $K = (F'(t)_\sigma)^a$. Then $K_{\text{res}} = L_{\text{res}}$, and $\text{val}(K) = \text{val}(L)$.

We have $\text{val}(K) \neq \text{val}(K^\sigma)$, while $tr.deg.\sigma(K)K = 1$. By 7.14, the hypothesis (*) of 7.25 holds for K. Thus by 7.25, $L = \widehat{K}$. Since $\widehat{K} \leq L' \leq L$, we have also $L' = L$. This proves (2).

The same proof shows that K embeds into L; so assuming (2), we conclude (3), using the assumption of (2) in place of 7.25.

(3) implies (1) trivially. □

Lifting the value group Let $K = K^a$ be a valued field. If K'/K, K''/K' are purely ramified, then clearly K''/K' is purely ramified. Thus K has a maximal purely ramified extension K' within a given extension L (not necessarily unique.)

Let L be a transformally valued field. Let

$$H_e = H_e(L) = \{v \in \text{val}(L) : (\forall u > 0)(|v| < \sigma^e(u))\}$$

H_e is a convex subgroup of $\text{val}(L)$; in case $\text{val}(L) = \mathbb{Q}_\sigma$, H_e is the e'th nonzero convex subgroup. When $v > c$ for every $c \in H_e$, we write: $v > H_e$. Let $h_e : \text{val}(L) \to \text{val}(L)/H$ be the quotient map, and let

$$VAL_e = h_e \circ \text{val} : L^* \to \text{val}(L)/H$$

We obtain auxiliary valuations, with residue map denoted RES_e. These are not transformal valuations.

When $L = F(t)_\sigma{}^a$, RES_e induces an isomorphism $F(t, t^\sigma, \ldots, t^{\sigma^{e-1}})^a \to \text{RES}_e(L)$ naturally. When $L = \widehat{F(t)_\sigma}^a$ the same is true, since any element of L is close to an element of $F(t)_\sigma$ to within a value $> H$.

Proposition 7.34 *Let $L = L^a$ be a transformal valued field of transformal dimension 1 over an inversive difference field $F = L_{res}$; with value group $\mathbb{Q}_\sigma$. Let $F < K = K^a \leq \widehat{L}$. Then there exists a difference field $K' \leq \widehat{L}$, with*

1. *$K \leq K'$, and K'/K is purely ramified.*
2. *With $e = rk_{\mathbb{Q}} \text{val}(L)/\text{val}(K')$, $\sigma^e \widehat{L} = \widehat{K'}$.*

Proof

Let K' be a maximal purely ramified difference field extension of K within $\widehat{L}$. Then $K' = (K')^a$.

By 7.33, $\widehat{L} = \widehat{F(t)^a_\sigma}$ for some t. Let $e = rk_{\mathbb{Q}} \text{val}(L)/\text{val}(K')$. Then there exists $s \in K'$, $\text{val}(s) > 0$, $s \in H_{e+1}$. Fix $N \geq e+1$ for a moment. Let $t_n = \sigma^n(t), s_m = \sigma^m(s)$, $\bar{t_n} = \text{RES}_N(t_n)$, $\bar{s_n} = \text{RES}_N(s)$.

We saw that $tr.deg._F \,\mathrm{RES}_N(L) = N$. Thus $\bar{t_0}, \dots, \bar{t_e}, \bar{s_0}, \dots, \bar{s}_{N-e-1}$ are algebraically dependent over F. Since $s \in H_{e+1}$, $\bar{s_0}, \dots, \bar{s}_{N-e-1}$ are algebraically independent over F. Let m be maximal such that $\bar{t_m}, \dots, \bar{t_e}, \bar{s_0}, \dots, \bar{s}_{N-e-1}$ are algebraically dependent over F.

Write $f(\bar{t_m}, \dots, \bar{t_e}, \bar{s_0}, \dots, \bar{s}_{N-e-1}) = 0$, $f \in F[X_m, \dots, X_e, \bar{s_0}, \dots, \bar{s}_{N-e-1})]$ of minimal X_m- degree.

We now argue that $m = e$. Lift f to $f(X_m, \dots, X_e, s_0, \dots, s_{N-e-1}) \in \mathcal{O}_L[X_m]_\sigma$, viewed as a difference polynomial in X_m (with coefficients involving the s_i, and $X_{m+j} = (X_m)^{\sigma^j}$.) In characteristic 0, let f' be the (first) derivative of f with respect to X_m; then $f' \neq 0$, and so by minimality of f, $RES_N f'(\bar{t_m}, \dots, \bar{t_e}) \neq 0$, i.e. $\mathrm{val}\, f'(t_m, \dots, t_e) \in H_N$. By 7.27, applied to the element t_m, there exists $b \in \widehat{L}$, $\mathrm{val}\,(t_m - b) > H_N$, $f(b, \dots, \sigma^{e-m}(b), s_0, \dots, s_{N-e-1}) = 0$. Let $L' = K'(b)_\sigma$. Clearly $tr.deg._K L' \leq e - m$, while $rk_{\mathbb{Q}} \,\mathrm{val}\,(L')/\,\mathrm{val}\,(K') \geq e - m$ (since $\mathrm{val}\,(\sigma^i(b)) = \mathrm{val}\,(t_{m+i})$, and $\mathrm{val}\,(t_m) << \mathrm{val}\,(t_{m+1}) << \dots << \mathrm{val}\,(t_{e-1}) << \mathrm{val}\,(c)$ for any $c \in K'$, $\mathrm{val}\,(c) > 0$.) However $rk_{\mathbb{Q}} \,\mathrm{val}\,(L')/\,\mathrm{val}\,(K') \leq tr.deg._{K'} L'$, so these numbers are equal, and L'/K' is purely ramified, of transcendence degree $m - e$. But K' is a maximal purely unramified extension; so $K' = L'$, and $m = e$.

In characteristic p, we first replace σ by $\sigma \circ \phi^{-l}$ where $\phi(x) = x^p$, reduce the equation if it is now a p'th power, till t_m occurs with an exponent that is not a power of p. We then use the above argument. So $m = e$ in any characteristic.

Thus $\bar{t_e} \in \mathrm{RES}_N(K')$ (since this residue field is algebraically closed). So for some $b_N \in K'$, $\mathrm{val}\,(t_e - b_N) > H_N(L)$. Now letting $N \to \infty$, we see that $t_e \in \widehat{K'}$. Let $K'' = F(t_e)^a_\sigma$; then $\sigma^e \widehat{L} = \widehat{K''} \subset \widehat{K'}$, and it remains to show equality. Note that $\mathrm{val}\,(K'') = \mathbb{Q}[\sigma] \,\mathrm{val}\,(t_e) = \mathrm{val}\,(K')$, and $\mathrm{res}\,(K'') = F = \mathrm{res}\,(K')$. Thus K'/K'' is immediate, so for any $c \in K'$, $T(c) = \{\,\mathrm{val}\,(c - b) : b \in K''\}$ can have no maximum value (7.14). But by 7.15, for any $c \in \widehat{L} \setminus \widehat{K''}$, $T(c)$ does have a maximal element. Thus $K' = K''$. □

Corollary 7.35 *If L satisfies 7.33 (1)-(3), and K is a closed difference subfield of L, $K_{\mathrm{res}} = L_{\mathrm{res}}$, then either K is trivially valued or K satisfies the same conditions.*

Proof By 7.30, we may assume L, K are transformally valued fields over F, $F = K_{\mathrm{res}} = L_{\mathrm{res}}$. Let K' be as in 7.34. Then K'/K is purely ramified, of finite transcendence degree r; there exist therefore $K_0 \leq K, K_0' \leq K'$, $tr.deg._F(K_0)$ finite, $tr.deg._{K_0}(K_0') = r = rk_{\mathbb{Q}} \,\mathrm{val}\,(K_0')/\,\mathrm{val}\,(K_0)$, $K' = KK_0'$. We can choose K_0 algebraically closed, and with value group equal to $\mathrm{val}\,(K)$. So $\widehat{K_0} \leq \widehat{K_0'} \leq \widehat{K'}$. $\widehat{K'}$ is isomorphic (by 7.34 (3)) to L, so it satisfies 7.33 (1)-(3); by (2), since $\widehat{K_0'}$ is a complete difference subfield with the same value group and residue field, $\widehat{K'} = \widehat{K_0'}$. Thus $\widehat{K_0} \leq K \leq \widehat{K'} = \widehat{K_0'}$. If $c \in \widehat{K'} \setminus \widehat{K_0}$, as K_0'/K_0 is purely ramified, and as in the proof of 7.15, $K_0(c)/K_0$ is also purely ramified. In particular if $c \in K \setminus \widehat{K_0}$ then $\mathrm{val}\,(c) \notin K_0$; but we chose $\mathrm{val}\,(K_0) = \mathrm{val}\,(K)$. Thus $K = \widehat{K_0}$. Now $\widehat{K_0}$ satisfies 7.33 (1). □

When E/F is an extension of valued fields, let

$$rk_{\,\mathrm{val}}\,(E/F) = \dim_{\mathbb{Q}}(\mathbb{Q} \otimes \mathrm{val}\,(E)/\,\mathrm{val}\,(F)) + tr.deg._F \mathrm{res}\, E$$

When E, F are subfields of a valued field L, we let $rk_{\mathrm{val}}(E/F) = rk_{\mathrm{val}}(EF/F)$, EF being the compositum of E, F in L. A similar convention holds for transcendence degree.

Corollary 7.36 *Let $K \leq L = L^a$ be an ω-increasing transformal valued field of transformal dimension 1 over an inversive difference field F, with value group $\mathbb{Q}_\sigma$. Then the following are equivalent:*

1. *$rk_{\mathrm{val}}(L/K) \leq e$*
2. *There exist difference fields $K \leq K_r \leq \widehat{K}$, and $K_r \leq K_1 \leq K_2 \leq \widehat{L}$, with:*
   - *(a) $tr.deg._F K_r \leq tr.deg._F K_{\mathrm{res}}$*
   - *(b) $tr.deg._{K_r} K_1 = tr.deg._{K_{\mathrm{res}}} L_{\mathrm{res}} = e_{\mathrm{res}}$*
   - *(c) $tr.deg._{K_1} K_2 = rk_{\mathbb{Q}} \operatorname{val}(K_2)/\operatorname{val}(K_1) = e_1$*
   - *(d) $\widehat{K_2} = \widehat{L}^{\sigma^{e_2}}$*
   - *(e) $e_{\mathrm{res}} + e_1 + e_2 \leq e$*
3. *There exists a difference field K', $\widehat{K} \leq K' \leq \widehat{L}$, $tr.deg._{\widehat{K}} K' + e_2 \leq e$, $\widehat{K'} = \widehat{L}^{\sigma^{e_2}}$.*
4. *There exist difference fields K', M, $\widehat{K} \leq K', \widehat{K'} \leq M, \widehat{M} = \widehat{L}$, with $tr.deg._{\widehat{K}} K' + tr.deg._{\widehat{K'}}(M) \leq e$.*

Thus, if $rk_{\mathrm{val}}(L/K) < \infty$, there exists a chain of difference subfields $K = M_0 \leq M_1 \leq \cdots \leq M_5 = \hat{L}$ such that $M_{i+1} = \widehat{M_i}$ for even $i < 5$, while for odd $i < 5$ $tr.deg._{M_i} M_{i+1} = rk_{\mathrm{val}}(M_{i+1}/M_i)$.

Proof Since F is inversive, and L_{res} is algebraically closed, by 7.13, L_{res} is also inversive. Thus for any $d \geq 1$, $\operatorname{res}(L^{\sigma^d}) = \operatorname{res}(L)$ By 7.5, $rk_{\mathrm{val}}(\widehat{L}/\widehat{L}^{\sigma^d}) = rk_{\mathbb{Q}}(\mathbb{Q}_\sigma/\sigma^d(\mathbb{Q}_\sigma)) = d$.

To show that (1) implies (2), we may assume $K = K^a$. By 7.30, there exists a field of representatives $F_r \leq \widehat{K}$ for the residue map of $\widehat{K}$. We have $tr.deg_F F_r = tr.deg._F K_{\mathrm{res}}$. Let $K_r = F_r K$.

By 7.30 again, applied to $\widehat{L}$ over K_r, there exists a difference field $F_r \leq L_1$ such that $\operatorname{res} : L_1 \rightarrow L_{\mathrm{res}}$ is an isomorphism. Let $K_1 = K_r L_1$. Then $tr.deg._{K_r} K_1 \leq e_{\mathrm{res}}$; equality holds by comparing the residue fields. Thus (b).

Now over K_1, 7.34 applies, and gives K_2 with (c,d). We have $K \leq \widehat{L}^{\sigma^{e_2}}$, and $e_1 \leq rk_{\mathbb{Q}} \operatorname{val}(\widehat{L}^{\sigma^{e_2}}/K)$; so $e_1 + e_2 \leq rk_{\mathbb{Q}}(\operatorname{val}(\widehat{L})/\operatorname{val} K) = rk_{\mathbb{Q}}(\operatorname{val} L/\operatorname{val} K)$. Thus (e).

(2) obviously implies (3), with $K' = K_2\widehat{K}$. To go from (3) to (4), let K' be as in (3). By 7.33, $\widehat{L} \simeq (L_{\mathrm{res}}\widehat{(t)}_\sigma)^a$, $\widehat{K'} = \widehat{L}^{\sigma^{e_2}} = (L_{\mathrm{res}}\widehat{(t^{e_2})}_\sigma)^a$. Let $M = (L_{\mathrm{res}}\widehat{(t^{e_2})}_\sigma)^a(t)$. Then $\widehat{M} = \widehat{L}$ and $tr.deg._{\widehat{K'}}(M) = e_2$.

(4) implies (1) since rk_{val} is additive in towers, bounded by transcendence degree, and vanishes for completions ($rk_{\mathrm{val}\,K}\widehat{K} = 0$.)

The final conclusion follows from (4) by taking $e = rk_{\mathrm{val}}(L/K)$. We obtain $M_0, \ldots, M_5$ with $M_{i+1} = \widehat{M_i}$ for even $i < 5$, and $tr.deg._{M_1} M_2 + tr.deg._{M_3} M_4 \leq e$. But $e = rk_{\mathrm{val}}(M_2/M_1) + rk_{\mathrm{val}}(M_4/M_3)$, and $rk_{\mathrm{val}}(M_{i+1}/M_i) \leq tr.deg._{M_{i+1}} M_i$; so all inequalities must be equalities

□

Corollary 7.37 *Let L be a transformal valued field of transformal dimension 1 over an inversive difference field F, with value group $\mathbb{Q}_\sigma$. Let K be a nontrivially valued subfield of*

L. Let K' be an extension of K within some ω-increasing transformal valued field containing L; assume $tr.deg._K K' < \infty$. Then $rk_{\mathrm{val}}(L/K') \leq rk_{\mathrm{val}}(L/K)$.

Proof Note that taking algebraic closure or completion does not change rk_{val}. We will thus assume all these fields are algebraically closed; and we will prove a more general statement, allowing L to be the completion of a transformal valued field of transformal dimension 1.

All residue fields are therefore also algebraically closed and, being extensions of K_{res} of transformal transcendence degree 0, are inversive. Note (say by 7.33) that $L' = \widehat{K'L}$ is also the completion of a transformal valued field of transformal dimension 1.

If $K \leq M \leq \widehat{L}$, then $rk_{\mathrm{val}}(L/K) = rk_{\mathrm{val}}(L/M) + rk_{\mathrm{val}}(M/K)$, and $rk_{\mathrm{val}}(L/K') = rk_{\mathrm{val}}(L/MK') + rk_{\mathrm{val}}(M/K')$; so the lemma for $M/K; K'$ and for $L/M; MK'$, implies the statement for $L/K; K'$. By 7.36, we may thus assume one of the following cases holds:

1. $L \leq \widehat{K}$
2. $tr.deg._K(L) = rk_{\mathrm{val}}(L/K)$

In case (1), $LK' \leq \widehat{K'}$, so $rk_{\mathrm{val}}(L/K') = rk_{\mathrm{val}}(\widehat{K'L}/\widehat{K'}) = 0$.

In case (2), $rk_{\mathrm{val}}(L/K') \leq tr.deg._{K'}(K'L) \leq tr.deg._K(L) = rk_{\mathrm{val}}(L/K)$.

□

Example 7.38 *In 7.37, the hypothesis that the value group of L is $\mathbb{Q}_\sigma$ is necessary.*

Let F be an inversive difference field, $a \in F$ K the inversive hull of $F(t)_\sigma$. Consider also the field of generalized power series $F((t^G))$, where the coefficient group G is $\cup_n \sigma^{-n}\mathbb{Q}_\sigma \subset \mathbb{Q}[\sigma, \sigma^{-1}]$ (ordered naturally.)

Let $L = K(b)$, where b solves the equation:

$$\sigma(x) - tx = a$$

We can take b to be a generic over $F((t^G))$.

Let

$$c = a^{\sigma^{-1}} + a^{\sigma^{-2}} t^{\sigma^{-1}} + a^{\sigma^{-3}} t^{\sigma^{-1}+\sigma^{-2}} + a^{\sigma^{-4}} t^{\sigma^{-1}+\sigma^{-2}+\sigma^{-3}} + \cdots \in F((t^G))$$

Let $K' = K(c)$.

Then $rk_{\mathrm{val}}(L/K) = rk_{\mathrm{val}}(K'/K) = 0$. But $rk_{\mathrm{val}}(L/K') = 1$, since $d = b - c$ is a nonzero solution of $\sigma(x) - tx = 0$, and hence $\mathrm{val}(d) = (\sigma - 1)^{-1}\mathrm{val}(t) \notin \mathbb{Q}_\sigma$.

Proposition 7.39 *Let L be a transformal valued field of transformal dimension 1 over F. Let $K = K^a \leq \widehat{L}$ be a complete subfield. Let $\bar{K}$ be a maximally complete valued field containing K. Then $L, \bar{K}$ (and even $\widehat{L}, \bar{K}$) are almost orthogonal over K in ACVF.*

Hence L is dominated by L_{res} over $K \cup \mathrm{val}(L)$.

Proof By 7.36, there exists a tower $K = K_0 \leq \cdots \leq K_6 = \widehat{L}$ of difference field extensions of K, with K_{i+1} purely ramified or purely inertial over K_i, or contained in $\widehat{K_i}$. (Namely $K \leq K_r \leq K_1 \leq K_2 \leq \widehat{K_2} \leq M \leq \widehat{L}$; with M as in 7.36 (4).) The proposition follows from 7.21.

Discussion The notation *tp* refers to the quantifier-free valued field type; so does the term "dominated". However this implies domination in the sense of transformal value fields too: if K' is a (transformal) valued field extension of K, with $\operatorname{res}(K'), \operatorname{res}(L)$ algebraically independent over $\operatorname{res}(K)$, then the quantifier-free valued field type of $K'/K \cup \operatorname{val}(L)$ implies the type of K'/L.

This implies the existence of canonical base change of L/K to extension fields of K, relative to a base change in $\operatorname{res}(L)$ and in $\operatorname{val}(L)$; the canonicity is such that the result goes through for transformal valued fields. We will make no use of this tool in the present paper.

8 Transformal valuation rings and analyzability

8.1 The residue map on difference varieties

We begin with some criteria for the residue map to be total, injective and surjective. We will not use them much, but they clarify the picture.

Lemma 8.1 *Let $D \subset D' = D[a]$ be difference domains; $a = (a_1, \ldots, a_n)$ a tuple of generators of D' as a difference D-algebra. Assume $\sigma(a) \in D[a]^{int}$ (the integral closure of the domain $D[a]$.). Then for large enough m, for any morphism $h : D' \to K$, with (K, R, v) an m-increasing transformal valuation field, if $h(D) \subset R$ then $h(D') \in R$.*

Proof For each i, $\sigma(a_i)$ is a root of a monic polynomial $f_i = \sum_{l=0}^{d_i} c_{i,l} X^l \in D[a][X]$. The coefficients $c_{i,l}$ are themselves polynomials in a over D; take m bigger than the total degree of all these polynomials. Then if v is a valuation, with $v(d) \geq 0$ for $d \in D$ and $v(a_i) \geq \delta$ for each i, (where $\delta < 0$), then for each i, l, $v(c_{i,l}) > m\delta$.

Assume $(K, R, \operatorname{val})$ is m-increasing, and $h(D) \subset R$. We may replace D, D' by their images under h, and assume h is the inclusion. Say $\operatorname{val}(ha_1) \leq \operatorname{val}(ha_2) \leq \ldots \leq val(ha_n)$. We have to show that $\delta = \operatorname{val}(a_1) \geq 0$. Otherwise, as the (monic) leading monomial of $f_i(\sigma(a_i))$ cannot have valuation less than all other monomials, we have $d_1 val(\sigma(a_1)) \geq \operatorname{val}(c_{1,l}) + l val(\sigma(a_1))$ for some $l < d_1$. Thus $(d_1 - l) val(\sigma(a_1)) \geq \operatorname{val}(c_{1,l}) > m\delta$. So $\operatorname{val}\sigma({a_1}^{-1}) \leq m \operatorname{val} {a_1}^{-1}$. This contradicts the assumption that val is m-increasing, and proves the lemma. □

Definition 8.2

For the sake of the lemma 8.3, define a *standard unramified map* to be a scheme morphism $\operatorname{Spec}^\sigma S \to \operatorname{Spec}^\sigma R$, where R is a (not necessarily Noetherian) commutative ring, $S = R[a_1, \ldots, a_n]$, and we have $f_i(a_1, \ldots, a_n) = 0$ for some polynomials $f_1, \ldots, f_n$ over R, with invertible Jacobian matrix. If the (f_i) give a presentation of S, i.e. $S = R[X_1, \ldots, X_n]/(f_1, \ldots, f_n)$, the map is said to be *étale*. Let us say that a map $f : X \to Y$ of schemes is *unramified* if for any two points of Y over a field, there exists an open subscheme Y' of Y containing the two points, $X' = f^{-1}(Y')$, with $X' \to Y'$ (isomorphic to) a standard unramified map. When Y is a difference scheme, we view the map $Y \to B_1 Y$ as a map of schemes, and apply the same terminology.

Lemma 8.3 *Let $(K, R, \bar{K})$ be a strictly increasing transformal valued field. Let X be a finitely generated quasi-projective difference scheme over R, $X_0 = X \otimes_R \bar{K}$. Assume the reduction sequence map $X_0 \to B_1 X_0$ is unramified (8.2). Then the residue map $X(R) \to X(\bar{K})$ is injective.*

Proof Consider first the special case $X = \mathrm{Spec}^\sigma R[x]/f(x)$, $f(x) = x + g(x) + h(x) \in R[x]_\sigma$; where g has coefficients in the maximal ideal M of R, and h is a sum of monomials x^ν, $\nu \in \mathbb{N}[\sigma]$, $nu > 1$. Let $a, b \in X(R)$ with $\mathrm{res}\,(a) = \mathrm{res}\,(b)$; we will show $a = b$. We may assume $b = 0$. Then $\mathrm{res}\,(a) = 0$ so $\mathrm{val}\,(a) > 0$; but then $\mathrm{val}\,(h(a)) > \mathrm{val}\,(a)$, and also $\mathrm{val}\,(g(a)) > \mathrm{val}\,(a)$, so $f(a) = 0$ implies $a = 0$.

In general, we may assume $X = \mathrm{Spec}^\sigma A$; $X_0 = \mathrm{Spec}^\sigma A_0$, $B_1 X_0 = \mathrm{Spec}^\sigma {A_0}'$, with $A_0 = A \otimes_R \bar{K}$, ${A_0}' = \sigma(A_0)$, with $X_0 \to B_1 X_0$ a standard unramified map (8.2): A_0 has generators $y_1, \ldots, y_n$ over ${A_0}'$, admitting relations $F_1, \ldots, F_n$, where $(F_i) \in {A_0}'[y_1, \ldots, y_n]$ has invertible Jacobian matrix $\det(\partial_{y_i} F_j) \in G_m({A_0}')$.

Complete $y_1, \ldots, y_n$ to a system of generators of A_0 over $\bar{K}$; since the y_i already generate A over ${A_0}'$, the additional generators y_j may be chosen from ${A_0}'$. Each new y_j solves the inhomogeneous linear polynomial $Y_j - y_j \in {A_0}'[Y_j]$; adding these generators and relations to the system clearly leaves the Jacobian invertible.

Now the coefficients of the F_i lie in $\sigma(A_0)$, so they may themselves be expressed as difference polynomials in the y_i, indeed in the $\sigma(y_i)$. Replacing the coefficients by these polynomials, we may assume that $F_i \in \bar{K}[Y_1, \ldots, Y_n]$ has coefficients in $\bar{K}$, rather than in $\sigma(A_0)$. If we convene that $\partial \sigma(Y_j)/\partial Y_k = 0$, the Jacobian matrix remains invertible.

Lift the generators y_i to $x_i \in A$. Then any element of A can be written as a difference polynomial over $\sigma(A)$ in the x_i, up to an element of $\mathcal{M}A = \ker A \to A \otimes_R \bar{K}$. (Here $\mathcal{M}$ is the maximal ideal of R.) So A has generators $x_1, \ldots, x_n, x_{n+1}, \ldots, x_N$, where for $i > n$ the image of x_i in A_0 is $y_i = 0$. We still have $A_0 = {A_0}'[y_1, \ldots, y_N]/(F_1, \ldots, F_N)$, where $F_i = y_i$ for $i > n$. Lift the F_i to $f_i \in \sigma(A)[x_1, \ldots, x_N]$. Then in A we have a relation $f_i(x_1, \ldots, x_N) = g_i(x_1, \ldots, x_N)$, where g_i has coefficients in $\mathcal{M}$.

Consider two elements $e = (e_1, \ldots, e_N), e' = (e'_1, \ldots, e'_N)$ of $X(R)$ specializing to the same point of X_0. Replacing x_i by $x_i - e'_i$, we may assume $e' = 0$. Suppose $e \neq 0$; let $\min_i \mathrm{val}\,(e_i) = \beta > 0$. In these coordinates, let L_i be the linear (monomials of order 1) part of f_i. The L_i are then the rows of an invertible matrix L over R, and $L_i x - g_i(x) + h_i(x) = 0$, where h_i is a sum of σ-monomials of degree > 1. Since $e \in X(R)$, we have $Le - g(e) + h(e) = 0$. (Where g, h are the matrices of σ-polynomials whose rows are g_i, h_i.)

Now $\min_i \mathrm{val}\,(\sigma(e_i)) = \sigma(\beta) > \beta$, so $\mathrm{val}\, h_i(e) > \beta$. Also $\mathrm{val}\, g_i(e) > \beta$. Writing $e = -G^{-1}(g(e) + h(e))$, we see that $\min_i \mathrm{val}\,(e_i) > \beta$, a contradiction. □

Remark 8.4 *Assume in 8.3 that $X \to B_1 X$ is étale, and that K is maximally complete as a valued field. Then $X(R) \to X(\bar{K})$ is surjective.*

Proof We can put X in the form: $F(X) = 0$, $X = (x_1, \ldots, x_n)$, $F = (f_1, \ldots, f_n)$, where the f_i are difference polynomials over R, and $J = \partial f_i / \partial x_j \in GL_n(R)$. (With the convention that $\partial({x_j}^\sigma)/\partial x_i = 0$.) Replacing F by $J^{-1}F$, we can write: $F(x) = c + x +$ (higher terms) .

Assuming $F(e_0) \in M^n$, we seek $e \in R^n$, $e - e_0 \in M^n$, and with $F(e) = 0$. The usual proof of Hensel's lemma provides such an e, by a transfinite sequence of successive approximations. (If $F(a) = \epsilon b$, $\epsilon \in R$, $v(\epsilon) > 0$, $b \in R^n$, then $F(a - \epsilon b) = \epsilon^2(\ldots)$; maximal completeness permits jumping over limit steps.)

Remark 8.5

In this paper, we will only use 8.4 when σ is the standard Frobenius, $\sigma(x) = x^q$ on K, K a complete discrete valuation ring. Here the convention that $\partial\sigma(x_j)/\partial x_i = 0$ follows from Leibniz's rule: $\partial({x_j}^q)/\partial x_i = 0$ in characteristic p.

8.2 Sets of finite total dimension are residually analyzable

Definition 8.6 *Let K be a valued field. A subset $X \subset K$ is* scattered *if $\{|x - y| : x, y \in X\}$ is finite. $X \subset K^n$ is scattered if $pr_i X \subset K$ is scattered for each i.*

The definition can be made more generally an ultrametric spaces. We will use it in contexts where X is definable in some expansion of K, and one can envisage X in arbitrary elementary extensions. Then the notion appears sufficiently close to the usual topological one to permit our choice of the term "scattered".

A scattered set X may consist of clusters of points of distance $\alpha < 1$. Each cluster can be enlarged, to reveal a new set of clusters separated by the residue map. After finitely many iterations, this process separates points of X.

Lemma 8.7 *Assume X is scattered. Then there are a finite number of equivalence relations $E_0 \subset E_1 \subset \ldots E_n$, such that $E_0 = Id$, and for each E_{i+1}-class Y of X, there exists a map f_i^Y embedding Y/E_i into the residue field. The E_i and f_i^Y are quantifier-free definable in the language of valued fields; f_i^Y is defined by a formula with a parameter depending on Y, while E_i requires no parameters beyond those needed to define X.*

Proof First take $X \subset K$. Let $\rho_0 < \rho_1 < \ldots \rho_n$ be the possible values of $|x - y|$ for $x, y \in X$. Say $\rho_i = |c_i|$. Let $E_i(x, y) \equiv |x - y| \le \rho_i$. Given Y, pick $b \in Y$, and let $f_i(x) = \operatorname{res} {c_i}^{-1}(x - b)$.

When $X \subset K^2$, we let $E_i(x, y) \equiv (|pr_0(x) - pr_0(y)| \le \rho_i \& pr_1(x) = pr_1(y))$ for $i \le n$, $E_j(x, y) \equiv |pr_1(x) - pr_1(y)| \le \rho_{j-n}$ for $j > n$; etc. □

Proposition 8.8 *Let K be a transformal valued field, with Γ a torsion-free $\mathbb{N}[\sigma]$-module. Let $F \in K[x]_\sigma$ be a nonzero transformal polynomial. Then $X = \{x : F(x) = 0\}$ is scattered. More generally, any $X \subset K^n$ of finite total dimension is scattered.*

Proof This reduces to $X \subset K$, using projections. When $X \subset K$ has finite total dimension, there exists a nonzero difference polynomial F over K with $X \subset \{x : F(x) = 0\}$. Recall the transformal derivatives F_ν. We have $F(x + y) = \sum_\nu F_\nu(x) y^\nu$. So if $a, a + b \in X$, then b is a root of $\sum_{\nu > 0} F_\nu(a) Y^\nu = 0$. If $F_\nu(a) = 0$ for all $\nu > 0$, then F is constant, so $X = \emptyset$. Otherwise, either $b = 0$ or $\operatorname{val} F_\nu(a) b^\nu = \operatorname{val} F_\mu(a) b^\mu$ for some $\nu < \mu$; so $(\mu - \nu) \operatorname{val}(b) = \operatorname{val} F_\mu(a) - \operatorname{val} F_\nu(a)$. By the Claim below, $\{\operatorname{val} F_\nu(a) : a \in X\}$ is finite for each ν; so $\{\operatorname{val}(b) : b \ne 0, a, a + b \in X\}$ is also finite.

Claim Let $G \in K[x]_\sigma$. Then $\{\operatorname{val} G(a) : \ F(a) = 0\}$ is finite.

It suffices to see that $\{\operatorname{val} G(a) : \ F(a) = 0\}$ remains bounded in any elementary extension L of K. In fact if $F(a) = 0$ then $\lambda \operatorname{val}(a) \in \operatorname{val}(K)$ for some $0 \neq \lambda \in \mathbb{N}[\sigma]$. Let $a \in L$, $F(a) = 0$, $c = G(a)$. Then $K(a)_\sigma$ has transcendence degree $\leq \deg_\sigma F$ over K; hence so does the subfield $K(c)_\sigma$. So $H(c) = 0$ for some nonzero $H \in K[x]_\sigma$, $H = \sum_\nu d_\nu x^\nu$. Arguing as above, we see that $(\mu - \nu) \operatorname{val}(c) = \operatorname{val} d_\nu - \operatorname{val} d_\mu$ for some $\nu < \mu$. □

Remarks on the liaison groups In most run of the mill cases, if S is the set of zeroes of a difference polynomial F within some ball B, some transformal derivative F' of F will have a unique root in B, and applying F' will map S into a ball of known radius around 0; so a direct coordinatization will be possible (internalization without additional parameters.) This process will fail in those some cases where a transformal derivative is constant, but non zero. The polynomial $F(X) = X^\sigma + X - a$ is an example. It seems likely that the examples can all be shown to have an Artin-Schreier aspect, and to become generalized Artin-Schreier extensions upon application of M_q.

Another approach, that does not explicitly look at the form of the polynomial, is in [Hrushovski02]. The associated groups of automorphisms of clusters over the residue field are all subgroups of the additive group; this has to do with "higher ramification groups" (cf. [Serre]).

We will not go in these directions here. However, in the next subsection, we will explain how to assign a dimension to scattered sets, given a notion of dimension over the residue field. (We will apply this to the total dimension over the residue field, obtaining something finer than total dimension of the closure over the valued field.) It will be convenient to explain the way that this dimension is induced in a more abstract setting. Our only application however will be the above mentioned one, and the reader is welcome to use 8.9,8.19 as a dictionary.

8.3 Analyzability and residual dimension

Let L be a language with a distinguished sort V. x, y, v will denote tuples of variables; and we will write $a \in M$ to mean: a is a tuple of elements of M.

Variables $v = (v_1, \ldots, v_m)$ will be reserved for elements of V. We use the notation $\phi(x; y)$ when we have in mind the formulas $\phi(x; b)$ with $b \in M \models T$. (This corresponds to the use of relative language, for schemes over a given scheme, in algebraic geometry.)

Data

A theory T (not necessarily complete). A set Φ of quantifier-free L-formulas and a set Φ_{fn} of basic functions of L. We assume Φ is closed under conjunctions and substitutions of functions from Φ_{fn}. The set of formulas of Φ in the variables x is denoted $\Phi(x)$; similarly $\Phi_{\mathrm{fn}}(x, t)$ refers to domain variables x and range variable t. We allow partial functions, whose domain is given by some $P \in \Phi(x)$. $\Phi(x; y)$ denotes formulas in $\Phi(xy)$, together with a partition of the variables, as indicated.

A map $d_V : \Phi(v; y) \to \mathbb{N}$.

Example 8.9 $L =$ the language of transformal valued rings, over a field F, with a distinguished sort $V = V_{\mathrm{res}}$ for the residue field.

$T =$ the theory of ω-increasing transformal valued fields.

$\Phi(x, v) =$ all quantifier-free formulas $\phi(x, v)$, implying that x has transformal dimension ≤ 1 over F, and v has transformal dimension 0 over F.

Φ_{fn} includes polynomials, and maps of the form $x \to \text{res}\,((x-y_1)/(y_1-y_2))$; with domain $\{(x, y_1, y_2) : y_1 \neq y_2, \text{ val}\,(x - y_1) = \text{ val}\,(y_1 - y_2)\}$.

$d_V(\phi(v; y)) \leq n$ iff for any $b \leq M \models T$, $\phi(v; b)$ implies $v \in X$ for some difference scheme X over M_{res} of total dimension $\leq n$.

Later, we will use a rank $Rk(L/K)$ on substructures; it corresponds to $rk_{\text{val}}\,(L/K)$.

In this case, the V-dimension $\dim_V$ defined below will be referred to as the *inertial dimension.*

We define V-co-analyzability (relative to T, Φ) and the V-dimension of $P \in \Phi$ (relative also to d_V), using recursion on $h \in (1/2)\mathbb{N}$. (Compare [Herwig-Hrushovski-Macpherson].) The second half-step between any two integers does not directly relate to x, but serves to provide additional parameters for future steps. Recall that the variables v refer to n-tuples from V. Since V is fixed, we will simply refer to co-analyzability (but will continue to talk of V-dimension.)

Definition 8.10 *1. $P(x; u)$ is 0-step coanalyzable over V (with V-dimension 0, i.e. $\leq n$ for all n) if $T \models P(x, u) \wedge P(x', u). \Rightarrow x = x'$.*

2. $P(x; u)$ is co-analyzable in $h + 1/2$ steps (with V-dimension $\leq n$) if there exist $Q \in \Phi(x; u, v)$, co-analyzable in h steps (with V-dimension $\leq n_1$), $R \in \Phi(v; u)$ (with $d_V(R) \leq n_2$), and $g \in \Phi_{\text{fn}}(x, u; v)$, such that

$$T \models P(x; u) \Rightarrow Q(x, u, g(x, u)) \wedge R(g(x, u); u))$$

(and $n_1 + n_2 \leq n$).

3. $P(x; u)$ is co-analyzable in $h+1$ steps (with V-dimension $\leq n$) if there are finitely many $Q_j \in \Phi(y; u)$ such that $T \models P(x; u) \Rightarrow \bigvee_j (\exists y) Q_j(y, u)$, and for each j,

$$(P(x; u) \wedge Q_j(y, u)) \in \Phi(x; uy)$$

is coanalyzable in $h + 1/2$ steps (with V-dimension $\leq n$).

We say P is (V-)co-analyzable if it is so in some finite number of steps. We will write

$$\dim_V(P) \leq r$$

for: " P is (V-)co-analyzable, of V-dimension $\leq r$. ".

In applications, it is convenient to apply this terminology to ∞-definable $P = \wedge_{i \in I} P_i$. An ∞-definable $P(x, u)$ is given, by definition, by a collection $\{P_i(x; u_i) : i \in I\}$, $P_i \in \Phi(x; u_i)$. u may be an infinite list, containing all the finite lists u_i. . We have in mind $P(M) =_{def} \cap_{i \in I} P_i(M)$. We simply define: $\dim_V(P) = \inf\{\dim_V(\wedge_{i \in I'} P_i) : I' \subset I, I' \text{ finite.}\}$.

When $K \leq L \leq M \models T$, and $a \in L$, we will write $\dim_V(a/K) \leq r$ if there exists $P \in \Phi(x; y)$, $\dim_V(P) \leq r$, and $b \in K$, such that $M \models P(a, b)$.

Example 8.11 *$P(x;y)$ is called V-*internal *if it is V-co-analyzable in* 1 *step.* In this case, if $b \in M \models T$, there exists an M-definable injective map $f : P(x,b) \to V^{eq}$. All such injective maps have the same image, up to an M-definable bijection; if V is stably embedded, and the given dimension on V is an invariant of definable bijections, then the V-dimension of P equals the dimension of the image of any of these maps f.

See [Hrushovski02] (appendix B) for a general treatment of internality, and associated definable groups.

Example 8.12 *Let $P(x) \in \Phi(x)$, $E_0 \subset \ldots \subset E_h \in \Phi(x,x')$ equivalence relations on $P(x)$, such that for each i, the set of E_{i+1}-classes is internal relative to the E_i-classes.* (I.e. for every class Y of E_i, there exists a function $f_Y = f(y,b) : Y \to V^{eq}$, $f \in \Phi_{\mathrm{fn}}$, parameters b depending on Y, such that f_Y is injective modulo E_{i+1}.) In this situation P is said to be *V-analyzable.* This is a stronger notion than co-analyzability. If the V-dimension of Y/E_{i+1} is n_i for each E_i-class Y, then $\sum n_i$ is an obvious upper bound for the V-dimension of P.

Remark 8.13 *In the case of valued fields, 8.9, if P is scattered, then it is in fact analyzable.*

Consider structures $K \models T_\forall$; (more precisely, subsets K of models M of T, closed under Φ_{fn}, with the formulas in Φ interpreted according to M.) We will consider pairs $K \leq L$ with L generated over K by a tuple c; write $L = K(c)$ in this case. Let

$$tp_\Phi(c/K) = \{\phi(x;b) : \phi(x;y) \in \Phi, b \in K, L \models \phi(c,b)\}$$

Assume given an $\mathbb{N} \cup \infty$-valued function Rk on such Φ-types. We will assume that the Rk does not depend on the choice of a generator c of L/K, and write $Rk(L/K)$ for $Rk(c/K)$ when $L = K(c)$.

Lemma 8.14 *Assume:*

1. *If $Rk(L/K) \leq n$, $a \in V(L)$, then for some $\phi(v;u) \in \Phi(v;u)$ with $d_V(\phi) \leq n$, and $b \in K$, $M \models \phi(a,b)$.*
2. *If $K \leq K' \leq L \leq M$, $Rk(L/K') + Rk(K'/K) = Rk(L/K)$.*
3. *If $K \leq K' \leq M$, $c \in M$, $Rk(K(c)/K) < \infty$, then $Rk(K'(c)/K') \leq Rk(K(c)/K)$.*

Let $K \leq L \models T_\forall$, with $Rk(L/K) \leq n$. Let $a \in L$, a/K co-analyzable. Then $\dim_V(a/K) \leq n$.

Proof We use induction on the number of steps of co-analyzability (the case of 0 steps being clear.) Suppose a/K is co-analyzable in $h + 1/2$ steps, i.e. $L \models P(a,b)$, $b \in K$, P co-analyzable in $h + 1/2$ steps. Then there exists $Q \in \Phi(x;u,v)$, co-analyzable in h steps, $g \in \Phi_{\mathrm{fn}}(x,u;v)$, and $d = g(a,b)$, such that $M \models Q(a;b,d)$. By induction, there exists $Q' \in \Phi(x;y',v)$ and $b' \in K$ with $M \models Q(a,b',d)$, $\dim_V(Q') \leq Rk(L/K(d))$. (Any parameters from $K(d)$ can be written as terms in elements b' of K, and d.)

By (1), there exists $R'(v;u) \in \Phi(v;u)$, $d_V(R') \leq Rk(K(d)/K)$, and $b'' \in K$, with $R(d,b'')$.

Let $P'(x;y,y',u) = Q'(x,y',g(x,y)) \wedge R(g(x,y);y''))$. Then $P'(a;b,b',b'')$ shows that the V-dimension of a/K is $\leq Rk(L/K(d)) + Rk(K(d)/K) = Rk(L/K)$.

Now suppose a/K is co-analyzable in $h+1$ steps. Then $L \models P(a,b)$, $b \in K$, $T \models P(x;u) \Rightarrow \bigvee_j (\exists y) Q_j(y,u)$, and for each j, (*) $P(x;u) \wedge Q_j(y,u)$ is coanalyzable in $h+1/2$ steps. Let $r = Rk(L/K)$,

$$\Xi = \{S \in \Phi(x;u,u',y) : \dim_V(S) \leq r\}$$

and let $\Xi' = \{\neg S(x,b,b',y) : S \in \Xi, b' \in K\}$.

Claim $T_\forall \cup tp_\Phi(a/K) \cup \Xi' \cup \{Q_j(y,b)\}$ is inconsistent.

Proof Suppose otherwise. Then in some $M \models T_\forall$, $K \leq M$, we can find $a' \models tp_\Phi(a/K)$, and d with $Q_j(d,b)$, such that (**) for any $S \in \Xi$, $M \models \neg S(a,b,b',d)$. Let $K' = K(d)$, $L' = K(a',d)$. by (3), $Rk(L'/K') \leq r$. By (*), a'/K' is coanalyzable in $h+1/2$ steps. Thus by induction, $\dim_V(a'/K') \leq r$. So there exists $S \in tp_\Phi(a'/K')$ with $\dim_V(S) \leq r$. We can take $S = S(x,b,b',d)$, $S \in \Phi(x,u,u',y)$. So $S \in \Xi$. But this contradicts (**). □

By compactness, for some finite disjunction $\bigvee_{j'} S_{jj'}(x,u,u',y)$ (with $S_{jj'} \in \Xi$) and some $P'(x,b,b') \in tp_\Phi(a/K)$,

$$T_\forall \models P'(x,u,u') \wedge Q_j(y,u). \Rightarrow \bigvee_{j'} S_{jj'}(x,u,u',y)$$

Since this is the case for each index j,

$$T \models P'(x,u,u') \Rightarrow \bigvee_{jj'} (\exists y) S_{jj'}(y,u,u',y)$$

So by 8.10 (3), $\dim_V(P') \leq r$. □

Dually, we have:

Proposition 8.15 *Let Rk satisfy the hypotheses of 8.14. Let $K \models T_\forall$, and let $P \in \Phi(x;y)$. (Or more generally, $P = \wedge_{i \in I} P_i$, $P_i \in \Phi(x;y)$.) Assume: whenever $K \leq M \models T$, $a \in M$, $b \in K$, if $P(a,b)$ then $Rk(K(a)/K) \leq r$. Then* $\dim_V(P) \leq r$

Proof For any $M \models T$, if $a,b \in M$, $M \models P(a,b)$, letting K be the substructure of M generated by b, we see that there exists $Q \in \Phi(x;y)$, $M \models Q(a,b)$, $\dim_V(Q) \leq r$. By compactness, $T \models P \Rightarrow \bigvee_j Q_j$, with $\dim_V(Q_j) \leq r$. It follows that $\dim_V(P) \leq r$. □

Specializing this to our example 8.9, we have:

Proposition 8.16 *Let F be an algebraically closed ω-increasing transformal valued fields of transformal dimension 1 over an inversive difference field, with value group $\mathbb{Q}_\sigma$. Let $\phi(x) \in \Phi(x)$ be a quantifier-free formula in the language of transformal valued fields over F, cf. 8.9. Assume ϕ- is V_{res}-analyzable, and: for any ω-increasing transformal valued field extension $L = F(c)$ of F, with $\phi(c)$, $rk_{\text{val}}(L/F) \leq n$. Then ϕ has inertial dimension $\leq n$.*

Proof This is a special case of 8.15. Note that the proof of 8.14 uses $Rk(K'/K'')$ only for finitely generated extensions K', K'' of K of transformal dimension 0 over K. For such extensions, take $Rk(K'/K'') = rk_{\text{val}}(K'/K'')$. Then hypothesis (1) of 8.14 is clear since if $rk_{\text{val}}(K'/K'') \leq n$ then $tr.deg.K'_{\text{res}}/K''_{\text{res}} \leq n$. (2) follows from additivity of transcendence degree and vector space dimension. And the truth of (3) is the content of 7.37 (with L the completion of $K(c)_\sigma$.) □

Remark 8.17 *There is a canonical function Rk satisfying 8.14 (1-3).* To define it, let

- $Rk_0(L/K) = \sup\{d_V(a/K) : a \in V(L)\}$
- $Rk_{n+1/2}(L/K) = \sup\{Rk_n(L/K') + Rk_n(K'/K) : K \leq K' \leq L\}$
- $Rk_{n+1}(L/K) = \sup\{Rk_n(LK'/K') : K \leq K', L \leq M\}$
- $Rk_\infty(L/K) = \sup_n Rk_n(L/K)$

Remark 8.18 *Let $P(x; y) \in \Phi(x; y)$ be V co-analyzable in h steps, with V-dimension $\leq n$, with respect to T, Φ.*

1. *For some finite $T_0 \subset T$, $\Phi_0 \subset \Phi$, $\Phi_{\mathrm{fn}0} \subset \Phi_{\mathrm{fn}}$, P has V-dimension $\leq N$ with respect to $T_0, \Phi_0, \Phi_{\mathrm{fn}0}$.*
2. *Let $M_q \models T_0$ be a family of models of T_0, indexed by an infinite set of integers $\{q\}$, and suppose that for any $P(v, y) \in \Phi_0$, for some β, for all q and all $b \in M_q$, $|P(M, b)| \leq \beta q^{d(P)}$.*

   Then for any $P(x, y) \in \Phi_0$ of V-dimension $\leq n$, for some β, for all q and all $b \in M_q$, $|P(M, b)| \leq \beta q^n$.

Proof (1) is obvious from the definition; (2) also follows immediately from the definition, by induction on the number of steps. (For each n, the half-step from n to $n + 1/2$ increases the V-dimension and the exponent; the half-step from $n + 1/2$ to $n + 1$ increases only the constant β.)

Example 8.19 (continuing 8.9). Here T_0 will be, for some k, the theory of k-increasing transformal valued fields. The family of models M_q can be taken to be the fields $K_q(t)^{alg}$, endowed with a nontrivial valuation over K_q, and with the q-Frobenius automorphism. The validity of the assumption will be seen in 10.8.

8.4 Direct generation

Let F be an inversive difference field, K a transformal discrete valuation ring over F. Then K is generated as a difference field by a subfield $K(0)$ with $tr.deg._F K(0) < \infty$, and (letting $K(n)$ be the field generated by $K(0) \cup \ldots \cup \sigma^n K(0)$), $tr.deg._{K(0)} K(1) = 1$. We wish to show that this automatically implies: $K_{\mathrm{res}} \subset (K(0)_{\mathrm{res}})^a$. We phrase it a little more generally.

Lemma 8.20 *Let F be an inversive difference field, and let $(K, R, M, \bar{K})$ be a weakly transformal valued field extending $(F(t)_\sigma, F[\breve{t}]_\sigma, tF[\breve{t}]_\sigma, F)$. Let $K(0) \subset K(1) \subset \ldots$ be subfields of K with $K = \cup_n K(n)$, $t \in K(1)$; let $R(n) = K(n) \cap R$, $\bar{K}(0) = res(R(0))$.*

Assume $tr.deg._F K(0) < \infty$, $tr.deg._{K(n)} K(n + 1) \leq 1$, and $\sigma(R(n)) \subset R(n + 1)$.

Then $\bar{K} \subset \bar{K}(0)^a$.

Proof Let us first reduce to the case that K is a transformal valued field, and the $val(t_n)$ are cofinal in Γ. Let Γ_∞ be the convex subgroup of Γ generated by the elements $val(t_n)$. (If

$r \in R = \cup_n R(n)$, $val(r) \leq val(t_n)$ then $\sigma(r) \neq 0$, so Γ' of 7.3(3) contains $Gamma_\infty$.) Let $\hat{K}, \hat{R}, \hat{M}, \pi, R_{\Gamma_\infty}$ be as in 7.3(6), and let $\hat{K}(n) = \pi(K(n) \cap \hat{R})$.

Note that σ extends to $\hat{R}$: an element of $\hat{R}$ has the form ba^{-1}, with $a, b \in R$, $0 \leq val(a) \leq val(t_n)$ for some n. Then $val(\sigma(a)) \leq val(t_{n+1})$ (in particular, $\sigma(a) \neq 0$.) Let $\sigma(ba^{-1}) = \sigma(b)\sigma(a)^{-1}$.

It follows that $\hat{K}$ is a transformal valued field. Since $\hat{M} \cap F[\check{t}]_\sigma = (0)$, $(\hat{K}, \hat{R}, \hat{M}, \bar{K})$ extend $(F(t)_\sigma, F[\check{t}]_\sigma, tF[\check{t}]_\sigma, F)$. This effects the reduction.

Thus σ extends to an endomorphism of K; and we have $\sigma(K(n)) \subset K(n+1)$.

Let Γ_n denote the divisible hull within $\Gamma(= \Gamma_\infty)$ of $\{val(a) : a \in K(n)\}$. So Γ_n is a group of finite rank.

Claim 2 Let $\Gamma_I = \{u \in \Gamma : u << \sigma(u)\}$. If one of $u, \sigma(u) \in \Gamma_I$ and $0 < u \leq v \leq \sigma(u)$ then $v \in \Gamma_I$. Each $t_m \in \Gamma_I$.

Proof If $\sigma(u) \leq mu$, then $\sigma^2(u) \leq m\sigma(u)$, contradicting $\sigma(u) \in \Gamma_I$. Thus $u << \sigma(u)$. So we can assume $u \in \Gamma_I$. If $mu \leq v$ for all m, applying σ we obtain $mv < m\sigma(u) \leq \sigma(v)$ for all m, so $v << \sigma(v)$. If $v < mu$ for some m, then $v \leq mu << \sigma(u) \leq \sigma(v)$.

Claim 3 If $u \in \Gamma_n \cap \Gamma_I$, then $\Gamma_n \neq \Gamma_{n+1}$, and in fact Γ_{n+1} has an element greater than any element of Γ_n.

Proof In an ordered Abelian group, if $0 < u_1 << u_2 << u_3 << \dots$ then $u_1, u_2, u_3, \dots$ are linearly independent. Thus there is no infinite $<<$-chain of elements of Γ_n. So we may take u to be $<<$-maximal in $\Gamma_n \cap \Gamma_I$. But then $u << \sigma(u)$, $\sigma(u) \in \Gamma_I$ so we cannot have $\sigma(u) \in \Gamma_n$. At the same time $\sigma(u) \in \Gamma_{n+1}$. So $\Gamma_n \neq \Gamma_{n+1}$.

Claim 4 For all $n \geq 0$, $\Gamma_n \neq \Gamma_{n+1}$. Moreover, Γ_{n+1} has an element greater than any element of Γ_n.

Proof As $val(t) \in \Gamma_1 \cap \Gamma_I$, the previous claim applies for $n \geq 1$. It applies to $n = 0$ too if Γ_0 has an element $v \geq val(t)$. (Since $v \leq val(t^{\sigma^m})$ for some m, so $v \in \Gamma_I$.) Otherwise, $\Gamma_0 < \mathrm{val}\,(t) \in \gamma_1$. This covers all the cases.

Claim 5 $K(n+1)_{\mathrm{res}} \subset (K(n)_{\mathrm{res}})^a$.

Proof This follows from valuation theory, the previous claim, and the assumption $tr.deg._{K(n)}K(n+1) \leq 1$.

This finishes the proof of the lemma. □

As in definition 7.2,let

$$\Delta = \{v \in val(R) : -val(t) < nv < val(t),\ n = 1, 2, \dots\}$$

Lemma 8.21 *In 8.20, we can also conclude:* $\Delta \subset \Gamma_0$.

Proof Note $\Delta \subset \Gamma_\infty = \cup_n \Gamma_n$. By Claim 4 of 8.20, Γ_{n+1} has an element greater than any element of Γ_n. But $tr.deg._{K(n)}K(n+1) \leq 1$ implies $rk_{\mathbb{Q}}(\Gamma_{n+1}/\Gamma_n) \leq 1$, and it it follows that $\Delta \cap \Gamma_n = \Delta \cap \Gamma_{n+1}$. By induction, $\Delta \cap \Gamma_n \subset \Gamma_0$, and the lemma follows. □

Corollary 8.22 *In 8.20, let $d = tr.deg._F K(0)$. Then the difference ring R/tR has a unique minimal prime ideal P. The total dimension of R/tR is $\leq d$; if equality holds, then $tR \cap R(0) = P \cap R(0) = (0)$. Ditto the total valuative dimension $rk_{\mathrm{val}}(K)$.*

If K is ω-increasing, then ${}^{\sigma}\sqrt{}(P) = M$.

Proof Let Δ be as in 8.21, and let $P = \{a \in R : val(a) \notin \Delta\}$. Then P (viewed as P/tR) is the ideal of nilpotents of R/tR, and P is prime.

Let $R' = \{a/b : a, b \in R, val(b) \in \Delta\}$. R' is a valuation ring, with maximal ideal PR', and value group Γ/Δ. R' is not in general invariant under σ, but $R' \cap K_n$ has value group Γ_n/Δ, and by 8.20, 8.21 the rank of this value group grows with n. Thus as in 8.20, if K' is the field of fractions K' of R/P, then $K' \subset K'(0)^a$, where $K'(0) = resR'(0)$. So $tr.deg._F K' = tr.deg._F K'(0) \leq tr.deg.K(0)$; the second inequality is strict unless the R'-valuation is trivial on $K(0)$, and in this case $R(0) \cap P = (0)$.

The total dimension of R/tR equals that of R/PR since $P/tR = \sqrt{}(0)$ in R/tR, and so is bounded by the transcendence degree of the field of fractions of this domain. The total dimension of R/tR equals the reduced total dimension, $r.dim(R/tR) = total.dim(R/M) = tr.deg._F \bar{K}$, plus the transformal multiplicity, or total dimension of $\operatorname{Spec} R/t$ over $\operatorname{Spec} K$. As there is a chain of prime difference ideals of length $rk_{ram}(K)$ between M and P, this total dimension is $\geq rk_{ram}(K)$. Thus

$$total.dim(R/tR) \geq tr.deg._F \bar{K} + rk_{ram}(K) = rk_{\mathrm{val}}(K)$$

So $rk_{\mathrm{val}}(K) \leq d$, and if equality holds then $P \cap R(0) = (0)$.

The last statement is clear from the definition of P: if K is ω-increasing, and $a \in M$, then $\mathrm{val}(a) > 0$; we have $\mathrm{val}(a) << \mathrm{val}(\sigma(a)) << \ldots$, so they are $\mathbb{Q}$- linearly independent, and hence cannot all be in Δ; thus $\sigma^m(a) \in P$ for some $m \leq rk_{\mathbb{Q}}\Delta < \infty$. □

9 Transformal specialization

9.1 Flatness

Intersection theory leads us to study difference schemes "moving over a line"; the behavior of a difference scheme X_t depending on a parameter t, as $t \to 0$. In his Foundations of Algebraic Geometry, Weil could say that (a, t) *specializes* to $(a', 0)$ (written $(a, t) \to (a', 0)$) if $(a', 0)$ lies in every Zariski closed set that (a, t) lies in; i.e. the point of $\operatorname{Spec}^{\sigma} X$ corresponding to $(a', 0)$ lies in the closure of the (generic) point (a, t). By Chevalley's lemma, this automatically implies the existence of a valuation ring, and the theory that comes with that. We think of that as indicating the existence of a "path" from (a, t) to $(a', 0)$. In difference schemes, closure and pathwise closure do not coincide. We will use transformal valuations to define the latter.

We can view the valuations (or open blowing up) as revealing new components, that cannot be separated with difference polynomials, but can be separated using functions involving for instance $x^{\sigma-1}$. (This can also be given responsibility for the failure of the dimension theorem in its original formulation, cf. [Cohn].)

The above phenomenon can lead to a special fiber with higher total dimension than the generic fiber, or even with infinite total dimension; in this case it cannot be seen as a smooth movement of a single object. To prevent this, we need to blow up the base. It is convenient to replace the affine line once and for all with the spectrum of a transformal valuation scheme (a posteriori, a finite blowing-up will suffice to separate off occult components in any particular instance. In general, we will use the language of valuation theory in the present treatment, so that blowing ups occur only in the implicit background.) Over a valuative base, we will show 9.3 that total dimension does not increase.

The lack of jumps in total dimension is analogous to the preservation of dimension (”flatness”) of the classical dynamic theory; but it is not the Frobenius transpose of the classical statement. The latter corresponds rather to the preservation of *transformal* dimension, a fact that holds true already over the usual transformal affine line, without removing hidden components. The good behavior of total dimension is related under Frobenius to another principle of classical algebraic geometry; what Weil called the ”preservation of number”.

We actually require something more than the flatness of total dimension when measured globally over a fiber. Consider a difference subscheme X of an algebraic variety V, or more generally a morphism $j : X \to [\sigma]_k V$. Say X is *evenly spread* (along V, via j) if for any proper subvariety U of V, $j^{-1}([\sigma]_k U)$ has total dimension $< \dim(V)$. We need to know that if the the generic fiber X_t is evenly spread, then the same is true of the special fiber. To achieve this, we need to replace the naive closure X_0 of with a pathwise closure $X_{\to 0} ='' \mathit{lim}_{t \to 0} X_t''$, and the total dimension by a valuative dimension. (cf. 9.10).

As a matter of convenience, since we are interested in the generic point and in one special point at a time, we localize away from the others.

Weakly transformal valuation rings will not be used in the proof of the main results.

Notation 9.1

A $\mathbb{Z}$-polynomial in one variable $F(X)$ is said to be *positive at* ∞ if $F(t) > 0$ for sufficiently large real t.

Given a difference ring k, let $k[t, t^{-1}]_\sigma$ be the transformal localization by t of the transformal polynomial ring $k[t]_\sigma$, and let $k[\breve{t}]_\sigma{}'$ be the sub-difference ring

$$k[\breve{t}]_\sigma{}' = k[t^{F(\sigma)} : F \in \mathbb{Z}[X], F(\infty) > 0] \leq k[t, t^{-1}]_\sigma$$

Write $t_n = \sigma^n(t)$. Note the homomorphism $k[\breve{t}]_\sigma \to k$, with kernel generated by the (t_n).

Alternative description: Let k be a difference field. Let $k(t)_\sigma = k(t_0, t_1, \ldots)$, with $\sigma(t_n) = t_{n+1}$, $t = t_0$. Then $k(t)_\sigma$ admits unique a k-valuation, with $0 < val(t_i) << val(t_j)$ whenever $i < j \in \mathbb{N}$. (Here $\alpha << \beta$ means: $m\alpha < \beta$ for all $m \in \mathbb{Z}$.) $k[\breve{t}]_\sigma$ is the associated valuation ring, and $\breve{A}_k = \mathrm{Spec}^\sigma\, k[\breve{t}]_\sigma$. Note that $k[\breve{t}]_\sigma$ is the localization of $k[\breve{t}]_\sigma{}'$ at the prime $t = 0$.

$\breve{A}_k$ will be used as a base for moving difference varieties, analogously to the affine line in rational equivalence theory of algebraic varieties. Let X be a difference scheme over $\breve{A}_k$; we will write X_t for the generic fiber $X \times_{\breve{A}_k} \mathrm{Spec}^\sigma\, k(t)_\sigma$, and X_0 for the special fiber $X \times_{\breve{A}_k} \mathrm{Spec}^\sigma\, k$ (referring respectively to the inclusion $k[\breve{t}]_\sigma \to k(t)_\sigma$ and the map $k[\breve{t}]_\sigma \to k$, $t \mapsto 0$.)

Let X be a difference subscheme of $\underline{V} = [\sigma]_k V \times_k \breve{A}_k$, V an algebraic variety over k. We denote by $X[n] \subset V \times V^\sigma \times \ldots \times V^{\sigma^n}$ the n'th weak Zariski closure of X, and similarly $X_t[n]$, $X_0[n]$.

Definition 9.2 *(cf. 9.5.) We will say that a difference scheme X of finite type over $\breve{A}_k$ is* flat *over $\breve{A}_k$ if in every local ring, $y \mapsto t_n y$ is injective.*

When X is algebraically reduced, this is equivalent to: X has no component contained in the special fiber X_0.

Lemma 9.3 *Let X be flat over $\breve{A}_k$. If X_t has total dimension d over $k(t)_\sigma$, then X_0 has total dimension $\leq d$ over k.*

Proof View $X[n]$ as a scheme over $\operatorname{Spec} k[\breve{t}]_\sigma$. Since $X_0[n] \subset X[n]_0$, it suffices to show that $\dim X[n]_0 \leq d$ (where $X[n]_0$ is the fiber of $X[n]$ above $t_0 = t_1 = \ldots = 0$). $X[n]$ arises by base extension from a scheme Y of finite type over $\operatorname{Spec} S'$, where S' is a finitely generated k-subalgebra of $S = k[\breve{t}]_\sigma$. Let $K = k(t)_\sigma$, $K[m] = k(t_0, \ldots, t_m)$, $S[m] = k[\breve{t}]_\sigma[m] = k[\breve{t}]_\sigma \cap K[m]$; then for some m we can take $S' = k[\breve{t}]_\sigma[m]$; and $\dim_{K[m]}(Y \otimes_{S[m]} K[m]) = d$.
Claim Let A be a finitely generated $S[m]$-algebra, such that $x \mapsto t_m x$ is injective on A. Suppose $\dim_{K[m]}(A \otimes_{S[m]} K[m]) = d$. (This refers to Krull dimension.) Let $A' = A/JA$, where J is the ideal generated by $(t_m, t_m t_{m-1}^{-1}, t_m t_{m-1}^{-2}, \ldots)$. Then A' is an $S[m]/J = S[m-1]$-algebra, $x \mapsto t_{m-1} x$ is injective on A' , and $\dim_{K[m-1]}(A' \otimes_{S[m-1]} K[m-1]) \leq d$
Proof Suppose $t_{m-1} c \in JA$; so $t_{m-1} c = t_m t_{m-1}^{-r} a$ for some $r \in \mathbb{N}$ and $a \in A$. As $t_{m-1} | t_m$, $x \to t_{m-1} x$ is injective on A. So $c = t_m t_{m-1}^{-r-1} a$. But then $c \in JA$. This shows that $x \mapsto t_{m-1} x$ is injective on A'.

To see that the dimension remains $\leq d$, let $B = A \otimes_{K[m]} K[m-1][t_m]'$, where $K[t_m]'$ is the localization of $K[m-1][t_m]$ at $t_m = 0$. It is easy to see, by looking at the numerator, that t_m is not a 0-divisor in B. We are given that $B \otimes_{K[m-1][t_m]'} K[m-1](t_m)$ has Krull dimension $\leq d$. Thus $\operatorname{Spec} B \to \operatorname{Spec} K[m-1][t_m]$ has generic fiber of dimension $\leq d$, and has no component sitting over $t_m = 0$, so it has special fiber of dimension $\leq d$ as well (over $t_m = 0$); thus $B/t_m B$ has Krull dimension $\leq d$. But $B/t_m B = A' \otimes_{S[m-1]} K[m-1]$, proving the claim.

The lemma follows upon $m+1$ successive applications of the Claim. □

Remark 9.4

Using the main theorems of this paper concerning Frobenius reduction, and the translation this affords, one can conclude:

1. A statement similar to 9.3 holds for transformal dimension; this corresponds to [Hartshorne] III 9.6 (or, using the dimension theorem in $\mathbb{P}^m \times \mathbb{P}^1$, where $X \subset \mathbb{P}^m$, note that no component of the special fiber $t = 0$ can have smaller codimension than a component of a generic fiber of an irreducible variety projecting dominantly to $\mathbb{P}^1$.)
2. When X is a closed subscheme of $V \times_k \breve{A}_k$, V a proper algebraic variety over k, one can conclude that X_0 has total dimension equal to d (though not necessarily reduced total dimension d).

   (The Frobenius specializations give systems of curves over the affine line, having about q^d points over a generic point of the affine line, hence the fiber over 0 cannot be of size $O(q^{d-1})$.)

3. One cannot expect every component of X_0 to have total dimension d (even with completeness and irreducibility assumptions). Nor can one expect the reduced total dimension to equal d. It suffices to note that the diagonal Δ of $\mathbb{P}^1$ can be moved to $(\{pt\} \times \mathbb{P}^1) \cup (\mathbb{P}^1 \times \{pt\})$ (and pull back with the graph of Σ.)

The conclusion of 9.3 below for ordinary schemes is usually obtained from flatness hypotheses; cf. e.g. [Hartshorne] III 9.6. This suggested the terminology, as well as the following remark. We have not been able to use it, though, since it is not obvious how to reduce to schemes of finite type while retaining flatness. The lemma refers to the k-algebra structure of $k[\check{t}]_\sigma$, disregarding the difference structure.

Remark 9.5 *Let k be a difference field, M be a $k[\check{t}]_\sigma$ module, such that $t_m y = 0$ implies $y = 0$. Then M is a flat $k[\check{t}]_\sigma$- module.*

Proof By [Hartshorne] III 9.1A, it suffices to show that for any finitely generated ideal I of A_k, the map $I \otimes_R M \to M$ is injective. Now $k[\check{t}]_\sigma$ is a valuation ring, so any finitely generated ideal is principal, $I = k[\check{t}]_\sigma c$. The map in question is injective if $cy = 0$ implies $y = 0$. But for some $m \in \mathbb{N}$, $c | t_m$ in $k[\check{t}]_\sigma$, so $t_m y = 0$, and thus $y = 0$.

9.2 Boolean-valued valued difference fields

We wish to consider finite products of transformal valuation domains (corresponding to a finite disjoint unions of difference schemes.) As we wish to include them in an axiomatizable class (for purposes of compactness and decidability), we will permit arbitrary products, and so will discuss briefly Boolean-valued transformal valuation domains. Compare [Lipshitz-Saracino73], [Macintyre1973]; but here we will need them only at a definitional level.

Definition 9.6 *(cf. [shoenfield].) Let T be a theory in a language L. The language $L_{boolean}$ has the same constant and function symbols as L; and for each l-place formula ϕ of L, a new l-place function symbol $[\phi]$, taking values in a new sort B; also, functions $\cup, \cap, \neg, 0, 1$ on B. $T_{boolean}$ is the theory of all pairs (M, B), where $(B, \cup, \cap, \neg, 0, 1)$ is a Boolean algebra, and M is a Boolean-valued model of T. Thus $[R](a_1, \ldots, a_l)$ is viewed as the truth value of $R(a_1, \ldots, a_l)$.*

Axioms: universal closures of $[\phi \& \psi] = [\phi] \cap [\psi]$, $[\neg \phi] = \neg [\phi]$; if $\phi = (\exists x)\psi$, $[\phi] \geq [\psi]$, and $(\exists x)([\phi] = [\psi])$.

And: $[\phi] = 1$, where $T \models \phi$.

When $B = [\mathbf{2}]$ is the 2-element Boolean algebra, we obtain an ordinary model of T.

If $(M, B) \models T_{boolean}$, and $h : B \to B'$ is a homomorphism of Boolean algebras, we obtain another model (M, B') of $T_{boolean}$, by letting $[\phi]' = h([\phi])$.

Thus when $(M, B) \models T_{boolean}$, $X = Hom(B, \mathbf{2})$, for each $x \in X$ we obtain a model M_x of T.

When T is a theory of fields, it is not necessary to have B as a separate sort. A model M of $T_{boolean}$ is a ring, and B can be identified with the idempotents of M, via the map $[x = 1] : M \to B$. (The sentence $(\forall y)(\exists x)(y = 0 \to x = 0 \,\&\, y \neq 0 \to x = 1)$, true in T, must have value 1 in $T_{boolean}$, and shows that for any $m \in M$ there exists an idempotent b with $[m = 0] = [b = 0]$.)

When L has an additional unary function symbol σ and T states that $\sigma(0) = 0, \sigma(1) = 1$, $T_{boolean}$ implies that σ fixes the idempotents.

We will just need the case when T is the theory of transformal valued fields (K, R, M). A model of $T_{boolean}$ is then a difference ring, called a *multiple transformal valued field.* $\{r : [val(r) \geq 0] = 1\}$ is a subring S, called a *boolean-valued transformal valuation domain.* Thus a boolean-valued transformal valuation domain is a certain kind of difference ring R, without nilpotents, such that for any ultrafilter U on the Boolean algebra B of idempotents of R, R/U is a transformal valuation domain. This class of rings is closed under Cartesian products.

When the Boolean algebra is finite, we say T is *finitely valued.* A finitely valued transformal valuation domain is just a finite product of transformal valuation domains.

9.3 Pathwise specialization

We will formulate a notion of a "specialization" $\lim_{t \to 0} X_t$ of a difference variety X_t over $F_t = F(t)_\sigma$, as $t \to 0$. This will be a difference subscheme of X_0; we will denote it more briefly as $X_{\to 0}$. In addition, there will be certain data over each point of $X_{\to 0}$ determining the "multiplicity" of that point as a point of specialization.

A *valuative* difference scheme over $\breve{A}_k$ is a difference scheme $T = \mathrm{Spec}^\sigma R$, R a transformal valuation domain extending $k[\breve{t}]_\sigma$, the valuation ring of $k(t)_\sigma$.

Note that any family of closed difference subschemes of a difference scheme X has a "union", a smallest closed subscheme containing each element of the family. This correspond to the fact that the intersection of any family of well-mixed difference ideals is again one. This makes possible the following definition.

Definition 9.7 *Let X be a difference scheme of finite type over $\breve{A}_k$. We define the pathwise specialization $X_{\to 0}$ to be the smallest well-mixed difference subscheme Y of X_0 such that for any valuative difference scheme T over $\breve{A}_k$, and any morphism $f : T \to X$ over $X \to \breve{A}_k$, $f(T_0) \subset Y$ (i.e. the restriction of f to the fiber over 0 factors through Y.)*

Intuitively, T is a kind of transformal smooth curve, and $f(T)$ marks a path from points on X_t to points on X_0.

Remark 9.8

(i) Since T is perfectly reduced, $X_{\to 0}$ depends only on the perfectly reduced difference scheme underlying X. (Though of course $X_{\to 0}$ need not itself be perfectly or even transformally reduced.)

(ii) One could think of allowing more singular paths, using weakly transformal valuation rings, with nilpotent elements permitted. Presumably, when X_t is reduced, and irreducible even in the valuative sense, the two definitions yield the same object.

(iii) A simple point of X_t may specialize to a singular point of X_0; or several simple points may specialize to one. This multiplicity of specialization is not faithfully reflected in geometric multiplicity on $X_{\to 0}$. (For instance, when X is the union of of several "transformal curves" C_i meeting at a point p of the special fiber, each C_i will map into X, but separately; so that p is registered just once on $X_{\to 0}$.) Thus the local algebra of points on $X_{\to 0}$ is insufficient to capture the data needed over each point.

One can modify the definition of $X_{\to 0}$ so that the local algebra will capture the multiplicities; cf. $X_{\to 0}{}'$ in §16.4. We will deal with the issue differently, replacing the local algebra information by scattered sets (definable in the language of transformal valued fields, and residually analyzable.) This approach appears to us much more transparent.

Lemma 9.9 *Let Z be a weak component of $X_{\to 0}$. Then there exists a valuative scheme T over $\breve{A}_k$, and a morphism $T \to X$ over $\breve{A}_k$, such that Z is contained in the image of T_0 in X_0.*

Proof Moving to an open affine difference subscheme of X, we can assume $X = \mathrm{Spec}^\sigma R$; Z corresponds to an algebraically prime difference ideal p; there are homomorphisms $h_j : R \to A_j \subset L_j$, L_j a transformally valued field extending $k[\breve{t}]_\sigma$, with corresponding transformal valuation domain A_j, such that p contains $\cap_{j \in J} {h_j}^{-1}(tA_j)$. Let $h : R \to A_* = \Pi_j A_j \subset L_* = \Pi_j L_j$ be the product homomorphism; then $h^{-1}(tA_*) \subset p$, and A_* is a Boolean-valued transformal valuation domain, with Boolean algebra E of idempotents.

Let U_0 be the filter in E generated by all $e \in E$ such that for some $r \in R \setminus p$, $h(r) \in tA_* + eA_*$. If $r_1, \ldots, r_n \in R \setminus p$, and $h(r_i) \in tA_* + e_i A_*$, then (since p is prime) $r = r_1 \cdot \ldots \cdot r_n \in R \setminus p$; and with $e = e_1 \cap \ldots \cap e_n$, $h(r) \in tA_* + eA_*$; so $e \neq 0$. Thus I_0 is a proper filter, and so extends to an ultrafilter U, with complementary maximal ideal I. Composing h with the map $A_* \to A_*/I =: \bar{A} \subset L_*/I =: \bar{L}$, we obtain $h^* : R \to \bar{A} \subset \bar{L}$, $\bar{A}$ a transformal valuation domain extending $\breve{A}_k$, such that for $r \in R \setminus p$, $h(r) \notin t\bar{A}$. Thus p contains $h^{-1}(t\bar{A})$. □

We come to the main estimate for the "equivalence of the infinite components". Given a difference scheme X, we have the residue map $X_t \to X_0$ (on the points in a transformal valued field.) If X has finite total dimension over $\breve{A}_k$, the pullback of a difference subvariety $W \subset X_0$ is a certain scattered set ${}^*W \subset X_t$. Under certain conditions of direct presentation, we find an upper bound on the inertial dimension of ${}^*W$; it can be smaller than the total dimension of the closure of ${}^*W$ within X_t.

Let V be an algebraic varieties over an inversive difference field k, $\dim(V) = d$. Let $V_{\breve{A}_k} = [\sigma]_k V \times_k \breve{A}_k$, $V_t = V \otimes_k k(t)$.

Proposition 9.10 *Let X be a closed difference subscheme of $V_{\breve{A}_k}$. Then (1) implies (3), and (1 & 2) implies (3 & 4).*

1. *Consider transformally valued fields L generated over $k(t)_\sigma$ by $c \in X_t(L) \subset V(L)$. Whenever $tr.deg._{k(t)_\sigma} L \geq d$, equality holds, and $\sigma(c) \in k(t,c)^a$.*

2. *Every weakly Zariski dense (in V) weak component of X_0 is transformally reduced.*

3. *$X_{\to 0}$ is evenly spread out along V.*

4. *for any subvariety W of V with* $\dim(W) < d = \dim(V)$,

$$^*W = \{x \in X_t(R) : \mathrm{res}_X(x) \in W\}$$

has inertial dimension $< \dim(V)$.

Explanations:

(1e) A point c of $X_t(L)$ corresponds to a morphism from a local difference ring of X_t, into L; let $k(c)_\sigma$ denote the field of fractions of the image. The inclusion $X_t \subset [\sigma]_{k(t)_\sigma} V_t$ allows us to restrict c to a point of $V(L)$, with a morphism from a local ring of V into L; let $k(c)$ denote the field of fractions of the image. Then (1) states that if $tr.deg._{k(t)_\sigma} k(c)_\sigma = d$, then $k(c)^a = (k(c)_\sigma)^a$. It follows also that $k(c)^a = k(V)$.

(2) says that a weakly Zariski dense weak component is actually a component of X_0. Note (2e) that any weakly Zariski dense (in V) difference scheme has total dimension $\geq \dim(V)$.

(3) means: every weak component Y of $X_{\to 0}$ of total dimension $\geq d$ is weakly Zariski dense in V. Note that (1) is a weak form of the statement that X_t is evenly spread out along V.

(4): Here R denotes the valuation ring, a quantifier-free formula in the language of transformal valued rings. res_X denotes the map $X_t(R^L) \to X_0(L_{\mathrm{res}})$ induced by res on a transformally valued field L.

Proof (1) $\Rightarrow$ (3) :

We may assume $X = \mathrm{Spec}^\sigma A$. Let Y be a weak component of $X_{\to 0}$, and let p be the corresponding algebraically prime difference ideal of A. By 9.9, there exists a transformally valued field (K, R) extending $(k(t)_\sigma, \breve{A}_k)$ and a difference ring morphism $h : A \to R$ over $k[\breve{t}]_\sigma$, such that p contains $h^{-1}(tR)$; so A/p embeds into a quotient of R/tR.

Note that $h^{-1}(tR)$ may be bigger than $(kerh, t)$, and $h(A)$ cannot be assumed to generate R (cf. §16.4). But we may assume $h(A)$ generates K as a field over $k(t)_\sigma$. By (1), $tr.deg._k K \leq d$.

Now t is not a 0-divisor in $A/\ker(h)$. By 9.3, $A/(\ker(h), t)$ has total dimension $\leq d$, hence so does $A/h^{-1}(tR)$. Moreover if $tr.deg._k K < d$, or if p is not a minimal prime of $A/h^{-1}(tR)$, then then $Y = \mathrm{Spec}^\sigma A/p$ has total dimension $< d$. So we may assume $tr.deg._k K = d$ (hence by (1), $k(V)$ embeds into K via h), and p is a minimal prime of $A/h^{-1}(tR)$.

Let $K[0] = k(V)^a \cap K$. Let $K[n] = K[0](t_0, \ldots, t_{n-1})^a \cap K$. Then $\sigma(K[0]) \subset K[1]^a$; so $\sigma(K[n]) \subset K[n+1]$. Now 8.20-8.22 apply to $\cup_n K[n]$. If R/tR has total dimension $< d$, then so does A/p, hence Y. In this case (3) holds trivially. Otherwise by 8.22, the nilradical of R/tR lifts to an algebraically prime difference ideal P of R, with $P \cap K[0] = (0)$. If $h(x) \in P$ then $h(x^n) \in tR$ for some n, so $x^n \in p$, and thus $x \in p$. So $h^{-1}(P) \subset p$. But p is a minimal prime, so $h^{-1}(P) = p$. Thus $p \cap k(V) = (0)$, i.e. Y is weakly Zariski dense in V.

(1&2) $\Rightarrow$ (4) : We will use 8.16. Let M be an ω-increasing transformal valued field extending $K = k(t)_\sigma$, with valuation ring $R = R^M$; let $c \in X_t(R^M)$, $\mathrm{res}(c) \in W$. Let $K = k(t, c)_\sigma$. We must show that $rk_{\mathrm{val}} K/k(t)_\sigma < \dim(V)$.

If $tr.deg._{k(t)_\sigma} K < \dim(V)$, this is clear. Otherwise, let $K(0) = k(c)$, $K(n) = k(c, \ldots, \sigma^n(c), t, \ldots, \sigma^{n-1}(t))$. By (1), $\sigma(c) \in k(c,t)^a$ so $\sigma^{n+1}(c) \in k(\sigma^n(c), \sigma^n(t))^a$. So the hypotheses of 8.20 hold. By

8.22, R/tR has a unique minimal prime ideal P; and $rk_{\text{val}}(K) < d$, unless $rk_{\text{val}}(K) = total.dim(R/tR) = d$, and $P \cap R \cap K(0) = (0)$.

In the latter case, we will obtain a contradiction. Let A be the image in R of the local ring of X, corresponding to the point c. Let $A(0)$ be the image of the corresponding restriction to a point of V, as in (1e). So $A(0) \subset R^M \cap K(0)$.

Thus $A(0) \cap P = (0)$. So $Z = \text{Spec}^\sigma A/(P \cap A)$ is a closed difference subscheme of X_0, weakly Zariski dense in V.

By 9.3, X_0 has total dimension $\leq d$; by (2e), Z has total dimension $\geq d$; so Z must be a weak component of X_0, of total dimension d; hence by (2), it is a component, i.e. $P \cap A$ is transformally prime in A.

Let K_A (resp. K') be the field of fractions of A/P (resp. R/P). Then $tr.deg._k K_A = d = tr.deg._k K'$. So $K' \subset (K_A)^a$. Since P' is transformally prime in A, by lemma 2.1, P is transformally prime. But if $a \in R$, $\text{val}(a) > 0$, then $\sigma^k(\text{val}(a)) > \text{val}(t)$ for some k; so $\sigma^k(a) \in P$; and thus $a \in P$. Hence P is the maximal ideal of R. Since $P \cap K(0) = 0$, the residue map is an isomorphism on $K(0)$. But $c \notin W$, $\text{res}(c) \in W$; a contradiction □

10 Frobenius reduction

10.1 The functors M_q

These functors may be viewed as difference-theoretic analogs of "reduction mod p"; we reduce mod p, and also "reduce" σ to a Frobenius map. In the case of subrings of number fields, the situation is what one ordinarily describes using the Artin symbol.

Let R be a difference ring. Let q be a power of a prime number p. We let $J_q(R)$ be the ideal generated by $p = p \cdot 1_R$ together with all elements $r^q - \sigma(r)$. Let $M_q(R) = (R/J_q(R))$. This is a difference ring, on which σ coincides with the Frobenius map $x \mapsto x^q$. In particular, J_q is a difference ideal.

Lemma 10.1 *Let T be a multiplicative subset of R with $\sigma(T) \subset T$, $\bar{T}$ the image of T under the quotient map in $M_q(R)$. Let $R[T^{-1}]$ be the localized ring. Then $M_q(R[T^{-1}]) = M_q(R)[\bar{T}^{-1}]$.*

Proof In whatever order it is obtained, the ring is characterized by the universal property for difference rings W with maps $R \to W$, $p = 0$ in W, such that the image of T is invertible, and $\sigma(x) = x^q$. □

Lemma 10.2 *Let $f : S_i \to R$ be surjective maps of difference rings. Then*

$M_q(S_1) \times_{M_q(R)} M_q(S_2) = M_q(S_1 \times_R S_2)$

Proof Similar to that of 10.1

Lemma 10.3 *Let $h : R \to S$ be a surjective difference ring homomorphism, with kernel I. Then the kernel of $M_q(h) : M_q(R) \to M_q(S)$ is $(I + J_q(R))/J_q(R)$.*

Proof The kernel is $h^{-1}(J_q(S))/J_q(R)$. If $h(r) \in J_q(S)$, $h(r) = \sum t_j(s_j^q - \sigma(s_j)) + ps'$ for some $t_j, s_j, s' \in S$. Writing $s_j = h(\bar{s}_j)$, $t_j = h(\bar{t}_j)$, $s' = h(\bar{s}')$, $\bar{r} = \sum_j \bar{t}_j(\bar{s}_j^q - \sigma(\bar{s}_j)) + p\bar{s}'$, we obtain $r - \bar{r} \in I$, $\bar{r} \in J_q(R)$, so $r \in I + J_q(R)$. □

Definition 10.4 *Let X be a difference scheme. Let $\mathcal{J}_q$ be the difference-ideal sheaf on X generated by the presheaf: $\mathcal{J}_q(U) = J_q(\mathcal{O}_X(U))$. $M_q(X)$ is the closed subscheme corresponding to $\mathcal{J}_q$.*

The above lemmas show that for a difference ring R without zero - divisors, $M_q(\mathrm{Spec}^\sigma R) = \mathrm{Spec}^\sigma M_q(R)$.

Let us say that a *q-Frobenius difference scheme* is a difference scheme in which σ coincides with $x \mapsto x^q$ on every local ring. Any ordinary scheme together with a map into $\mathrm{Spec}\,\mathbb{F}_q$ admits a canonical q-Frobenius difference scheme structure.

M_q yields a functor on difference -schemes into q-Frobenius difference schemes.

Remark 10.5

Let K be a number field, σ an automorphism, R a finitely generated difference subring of K. Then R is also finitely generated as a ring. $\mathrm{Spec}\, M_q(R)$ may be empty, and may be reducible. It is always a reduced scheme: if $d = [K : \mathbb{Q}]$, then $\sigma^d(x) = x$ on R, so $x^{q^d} = x$ on $M_q(R)$, hence $M_q(R)$ has no nilpotents.

Remark 10.6

Let R be as above, S a finitely generated domain over R of positive but finite transcendence degree d, and let σ be the identity automorphism. Then for all but finitely many q, $M_q(S)$ is reducible. Indeed every homomorphism of S into K_q factors through $M_q(S)$. We will see later that the number of homomorphism of $M_q(S)$ into K_q is $O(1)q^d$; these all fall into the field $GF(q)$. The relevant Galois group has order d, so the number of prime ideals of $M_q(S)$ is at least $O(1)q^d/d$, and in particular > 1. But $M_q(S)$ is finite, hence it is reducible.

On the other hand, if $S = \mathbb{Z}[X], \sigma X = X + 1$, then $M_p(S)$ is a domain for all primes p.

Even when the fraction field of S is a "regular" extension of K, in the sense that it is linearly disjoint over K from any difference field of finite transcendence degree over K, and $M_q(R)$ is a field, $M_q(S)$ can be reducible. Example: $R = \mathbb{Z}[X, Y : \sigma(X) = YX]$.

It would be interesting to study this question systematically.

Notation 10.7 *Let X be a difference scheme over Y, q a prime power, y a point of $M_q(Y)$ valued in a difference field L. We denote $X_{q,y} = (M_q(X))_y$ (cf. 3.12). Thus if $X = \mathrm{Spec}^\sigma R$, $Y = \mathrm{Spec}^\sigma D$, R a D-algebra, then $X_{q,y} = \mathrm{Spec}\, M_q(R) \otimes_{M_q(D)} L$.*

10.2 Dimensions and Frobenius reduction

We will see that the functors M_q take transformal dimension to ordinary algebraic dimension. When the transformal dimension of X is zero, $M_q(X)$ will be a finite scheme; in this case the total dimension is a (logarithmic) measure of the rate of growth of $M_q(X)$ with q. Transformal multiplicity becomes, in a similar sense, geometric multiplicity, while transformal degree is related to the logarithm of the projective degree.

By the *size* $|Y|$ of a 0-dimensional scheme Y over a field, we will mean the number of points, weighted by their geometric multiplicity. Thus for a k-algebra R, the size of $\operatorname{Spec} R$ over $\operatorname{Spec} k$ is just $\dim_k(R)$.

Lemma 10.8 *Let X be a difference scheme of finite type over a Noetherian difference scheme Y.*

1. *Assume X has transformal dimension $\leq d$ over Y. Then for all prime powers q, and all points y of $M_q(Y)$, $X_{q,y}$ has dimension at most d as a scheme over L_y.*
2. *If X has reduced total dimension $\leq e$ over Y, then there exists $b \in \mathbb{N}$ such that for all large enough prime powers q, and all points $y \in M_q(Y)(L)$, L a difference field, the zero-dimensional scheme $X_{q,y}$ over L_y has at most bq^e points.*

Remark The statement of 10.8(2) refers to the number of points, multiplicities ignored. We will see later that it remains true if multiplicities are taken into account. The proof of this refinement is more delicate in that when passing to a proper difference subscheme X' of X one cannot forget X; even if a point is known to lie on X' one must take into account the multiplicity of X , not of X', at that point.

Proof The lemma reduces to the case $X = \operatorname{Spec}^\sigma(R)$, $Y = \operatorname{Spec}^\sigma(D)$, R a finitely - generated D-algebra. We may assume Y is irreducible and perfectly reduced, i.e. D is a difference domain. We wish to reduce the lemma further to the one-generated case, $R = D[a, \sigma(a), \ldots]$. We use Noetherian induction on $\operatorname{Spec}^\sigma D$ and, with D fixed, on $\operatorname{Spec}^\sigma S$, and thus assume the lemma is true for proper closed subsets.

Claim If $D \subset S \subset R$ and the lemma holds for $\operatorname{Spec}^\sigma R_{U'}$ over U' and for $\operatorname{Spec}^\sigma S_U$ over U, whenever U', U are Zariski open in $\operatorname{Spec}^\sigma S$, $\operatorname{Spec}^\sigma D$ respectively, then it holds for $\operatorname{Spec}^\sigma R$ over $\operatorname{Spec}^\sigma D$.

Proof We prove this for (1); the proof for (2) is entirely similar. Let $p \in \operatorname{Spec}^\sigma R$, $p \cap D = p_D$, $p \cap S = p_S$. Let d' the relative dimension of S_{p_S} over D_{p_D}, d'' the relative dimension of R_p over S_{p_S}. Then $d' + d'' \leq d$. There exists an nonempty open affine neighborhood U' of p_S in $\operatorname{Spec}^\sigma S$ such that $\operatorname{Spec}^\sigma R$ has relative dimension $\leq d''$ over any point of U'. There exists an nonempty open affine neighborhood U of p in $\operatorname{Spec}^\sigma D$ such that $\operatorname{Spec}^\sigma S$ has relative dimension $\leq d'$ over any point of U. Let q be a large prime, $y \in M_q(\operatorname{Spec}^\sigma D)$. Then as the lemma is assumed to hold in those cases, $(\operatorname{Spec}^\sigma S'_U)_{q,y}$ has dimension at most d'' over U', while $(\operatorname{Spec}^\sigma R_U)_{q,y}$ has dimension at most d'. It follows that for any $p \in \operatorname{Spec}^\sigma R$ with $p \cap R \in U'$ and $p \cap D \in U$, the lemma holds. By additivity of transcendence degree in extensions, the lemma holds for $\operatorname{Spec}^\sigma R$ over $\operatorname{Spec}^\sigma D$ at a neighborhood of any such p. Outside of U, or of U', the lemma still holds by Noetherian induction. □

We continue to use Noetherian induction on $\operatorname{Spec}^\sigma D$. Let k be the field of fractions of D. We may assume R is perfect, since factoring out the perfect ideal generated by 0 changes neither the transformal or reduced total dimension nor the physical size of the zero-dimensional schemes $X_{q,y}$. In this case O is the intersection of finitely many transformally prime ideals p_i of R, and it suffices to prove the lemma separately for each R/p_i. Thus we may assume R is a difference domain.

By the Claim, we may assume R is generated a single element a as a difference D-algebra.

If R is isomorphic to the difference polynomial ring over D , the transformal dimension of $\mathrm{Spec}^\sigma R$ over $\mathrm{Spec}^\sigma D$ equals 1, and we must show the dimension of $specM_q(S)$ over $specM_q(D)$ is everywhere at most 1. This is clear since $M_q(R)$ is generated over $M_q(D)$ by a single element.

Otherwise, there exists a relation $G(a, \sigma(a), \ldots, \sigma^m(a)) = 0$ in F, G over D. Moreover in case (2), we can choose $m \leq e$. We may write $G(a) = \sum c_i a^{\nu_i(\sigma)}$, where $c_i \in S, c_i \neq 0$, the ν_i are finitely many integral polynomials, of degree at most the total dimension r of $\mathrm{Spec}^\sigma R$ over $\mathrm{Spec}^\sigma S$.

Note that for all sufficiently large q, the values $\nu_i(q)$ are distinct, indeed $\nu_i(q) < \nu_j(q)$ if $i < j$. The set Y' of transformally prime ideals of D containing one of the nonzero coefficients of G is a proper closed subset of $\mathrm{Spec}^\sigma S$; by induction the lemma is true over Y'. On the other hand if $y \in \mathrm{Spec}^\sigma S \setminus Y'$, then the monomials of $M_q(G)$ are distinct and their coefficients are non-zero in L_y; it follows that $M_q(R) \otimes L_y$ is 0-dimensional over L_y, and moreover has at most $\max_i \ \nu_i(q) \leq O(q^e)$ points. This finishes the proof of the lemma. □

10.3 Multiplicities and Frobenius reduction

The order of magnitude of multiplicity upon Frobenius reduction will be shown to be bounded by the transformal multiplicity. We will begin with two elementary and purely algebraic lemmas regarding multiplicities.

10.3.1 An explicit example

We begin with an example in one variable; it shows explicitly the distribution of multiplicities of the points of $M_q(X)$, where X is a difference scheme defined by $F(X, X^\sigma, \ldots, X^{\sigma^n}) = 0$. If F is irreducible, there will be at most about q^{n-l-1} points with multiplicity of order q^l. The general picture is similar (cf. 5.9), but we have not succeeded in reducing it to the example (because of tricky additivity behavior of transformal multiplicity), and will give an independent proof.

The Hasse derivatives $D^\nu f$ of a polynomial f are defined by the Taylor series expansion:

$$f(X + U) = \sum_\nu (D^\nu f)(X) U^\nu$$

Here $X = (X_0, \ldots, X_n)$, $U = (U_0, \ldots, U_n)$, $X + U = (X_0 + U_0, \ldots, X_n + U_n)$, ν is a multi-index $(\nu_0, \ldots, \nu_k)$, $U^\nu = \Pi_i {U_i}^{\nu_i}$. Write *sup supp* ν for the highest i such that $\nu_i > 0$.

Example 10.9 *Let K be a field, $f \in K[X_0, \ldots, X_n]$ an irreducible polynomial, $f \notin K[X_1, \ldots, X_n]$. For $k \geq 1$, let V_k be the Zariski closed subset of $\mathbb{A}^{n+1}$ defined by the vanishing of all $D^\nu f$ with sup supp $\nu < k$.*

1. *$V_k = \mathbb{A}^k \times U$, $U \subset \mathbb{A}^{n+1-k}$, $\dim(U) < n - k$.*
2. *Assume K has characteristic $p > 0$, and let q be a power of p with $q > \deg_{X_i}(f)$ for $i < k$. Let $F(X) = f(X, X^q, \ldots, X^{q^n})$. Let S be a subscheme of $\mathbb{A}^1$ defined by an ideal*

I with $F \in I$, and let $a \in S$, $(a, \dots, a^{q^n}) \notin V_k$. Then the multiplicity of S at a is bounded by

$$\mathrm{Mult}_a(S) \leq \sum_{i<k} \deg_{X_i} f q^{k-1} \leq (\deg f) q^{k-1}$$

Proof (1) If $(a_0, \dots, a_n) \in V_k$, then using the Taylor series expansion, we see that all polynomials $D^\nu f$ are constant on $\mathbb{A}^k \times \{(a_k, \dots, a_n)\}$; so $\mathbb{A}^k \times \{(a_k, \dots, a_n)\} \subset V_k$. Thus $V_k = \mathbb{A}^k \times U$ for some U. Since f is irreducible, and involves X_0, V_k is a proper subset of the zero set $V(f)$. Thus $\dim V_k < \dim V(f) = n$. So $\dim U < n - k$.

(2) We may assume $a = 0$. Let g be the sum of monomials of f involving only the variables $X_0, \dots, X_{k-1}$. Then $D^\nu(f-g)(0) = 0$ if *sup supp* $\nu < k$. As $0 \notin V_k$, $D^\nu f(0) \neq 0$ for some such ν. Thus $g \neq 0$. Let $G(X) = g(X, \dots, X^{q^{(k-1)}})$. A monomial $\Pi_i X_i^{a_i}$ of g turns into the monomial $X^{\sum a_i q^i}$ of G; as by assumption $a_i < q$ for each i, no cancellation occurs; so $G \neq 0$. Write $G(X) = X^m H(X)$, $H(0) \neq 0$. Then for some I,

$$F(X) = G(X) + X^{q^k} I(X) = X^m (H(X) + X^{q^k - m} I(X))$$

Note $m < q^k$. So $F(X)/X^m$ is a polynomial not vanishing at 0, thus m bounds the multiplicity of F at 0. □

10.3.2 Algebraic lemmas

Lemma 10.10 *Let k be a field of characteristic p, $[k : k^p] = p^e$ (in characteristic 0, let $e = 1$.) Let S be a k-algebra, generated by m elements, M a maximal ideal of S with $dim_k(S/M) < \infty$. Then M is generated by at most $m + e + 1$ elements.*

Proof Let $h : S \to K$ be a surjective homomorphism with kernel M, K a finite extension field of k. Let $s_1, \dots, s_m$ be generators for S as a k-algebra.

Consider first the case $e = 0$, i.e. K/k separable. Using the primitive element theorem, K is generated over k by one element $a = h(s_0)$; and one relation, the minimal monic polynomial $P(a) = 0$, $P \in k[X]$. We have $h(s_i) = Q_i(a) = h(Q_i(s_0))$ (some $Q_i \in k[X]$), so $s_i - Q_i(s_0) \in M$; and clearly M is generated by the $m+1$ elements $P(s_0), s_i - Q_i(s_0)$.

In general, K/k is generated by $e+1$ elements, and $e+1$ relations. (Let $K \leq K_s \leq K$, where K_s/k separable, K/K_s purely inseparable. Then K_s/k can be presented by one generator a_0 and one relation, as above. K/K_s is generated by $\leq e$ generators $a_1, \dots, a_e$ and $\leq e$ relations; these can be viewed as a relation over k between a_0 and the a_i.) As above, the maximal ideal M is generated by the preimages of the above relations. □

Lemma 10.11 *Let k be a field, R a finite dimensional local k - algebra, with maximal ideal M. Let S be a finitely generated R-module. Then*

$\dim_k S \leq (\dim_k R)(\dim_k(S/MS))$

Proof Let A be a k-subspace of S with $\dim_k(A) = \dim_k(S/MS)$ and $A + MS = S$. Let T be the k-span of RA. So T is an R-submodule of S, and $\dim_k(T) \leq (\dim_k R)(\dim_k(S/MS))$. Also $T + MS = S$. By Nakayama (applied to the R-module S/T), $T = S$.

Lemma 10.12 *Let $f : X \to Y$ be a morphism of finite schemes over* $\mathrm{Spec}\,(k)$*. Let $p \in$* $\mathrm{Spec}\,Y$*, X_p the fiber of X above p, and $q \in X_p$. Then*

$\mathrm{Mult}_q X \leq (\mathrm{Mult}_q X_p)(\mathrm{Mult}_p Y)$

Proof Let R, S be the local rings of Y, X at p, q respectively. Then the local ring of X_p at q is S/MS, M being the maximal ideal of R. We must show that $(\dim_k S) \leq (\dim_k R)(\dim_k(S/MS))$. This follows from 10.11.

Lemma 10.13 *Let S a local ring, with finitely generated maximal ideal M. Suppose S/M^r has length $n < r$. Then $M^n = 0$.*

Proof We have $S \supset M \supset M^2 \supset \ldots \supset M^r$. M^i/M^{i+1} is an S/M-space of some finite dimension. As the length of S/M^r is $n < r$, we must have $M^i = M^{i+1}$ for some $i \leq n$. By Nakayama, $M^i = 0$. So $M^n = 0$.

N.B. Let I be an ideal of a ring R with $p \cdot 1_R = 0$, and let q be a power of p. Let $\phi_q(x) = x^q$. Then I^q might mean one of three things: the ideal RI^q generated by q-fold products of elements of I; the ideal $\phi_q(I)$ of $\phi_q(R)$; or possibly the ideal $R\phi_q(I)$ of R. We will refrain from writing I^q altogether, and use the above notations.

Remark 10.14 *Let S be a ring, with ideal M generated by $s_1, \ldots, s_b$. Assume $p = 0$ in S, and let q be a power of p. Then $SM^{bq} \subset S\phi_q(M) \subset SM^q$*

Proof The ideal SM^{bq} has generators of the form ${s_1}^{m_1} \cdot \ldots \cdot {s_b}^{m_b}$ with $m_1 + \ldots + m_b = bq$. Necessarily $m_i \geq q$ for some i, so ${s_i}^{m_i}$ and hence the product are in $\phi_q(M)$. □

Lemma 10.15 *Let S be a ring of characteristic $p > 0$, with maximal ideal M, and let q be a power of p. Assume M is finitely generated. Then $\phi^q(S)$ has a unique maximal ideal $\phi_q(M)$, with residue field $k = \phi_q(S)/\phi_q(M)$. Note S is an S^q-module, so $S/S\phi_q(M)$ is a k-space. Assume* $\dim_k(S/S\phi_q(M)) < q$*, or even just*

(#) $\dim_k(S/SM^q) = n < q$

Then S has length $\leq n$ (as an S-module.)

Proof Clearly S/SM^q has length $\leq n$ as an S-module. By Lemma 10.13, $SM^n = 0$, so $SM^q = 0$. Thus $S = S/SM^q$ has length $\leq n$. □

Let S be a local ring of characteristic $p > 0$, with maximal ideal M, generated by b elements. For r a power of p, write $S_r = \phi_r(S)$, $M_r = \phi_r(M)$; so S_r is local, with maximal ideal M_r (the elements of $S \setminus M$ are units; hence $\phi_r(S \setminus M)$ consists of units of S_r, while $\phi_r(M)$ is a proper ideal of S_r. These sets are thus disjoint, while their union equals S_r.)

Lemma 10.16 *Let S be a local ring of characteristic $p > 0$, with maximal ideal M, generated by b elements. Let r, r' be powers of p. Assume $\phi_r(S/SM_{r'})$ has length n as an S_r-module, and $(nb+1)r < r'$. then $SM^{rb^{n+1}} = 0$. In particular $(S_r M_r)^{b^{n+1}} = 0$, so S_r has length $< b^{b^{n+1}}$.*

Proof The assumption that $\phi_r(S/SM_{r'})$ has length n refers to the ring $S/SM_{r'}$, and the function $\phi_r(x) = x^r$ in that ring. Being a quotient of S, $S/SM_{r'}$ is an S-module, and $\phi_r(S/SM_{r'})$ is an S_r-module. Note

$$\phi_r(S/SM_{r'}) = (\phi_r(S) + SM_{r'})/(SM_{r'}) \simeq_{S_r} \phi_r(S)/(SM_{r'} \cap \phi_r(S))$$

Thus $S_r/(SM_{r'} \cap S_r)$ has length n as an S_r-module. (It is an $S_r/(M_r)^{r'}$-module; this ring has finite length as a module over itself, so all nontrivial finitely generated modules below have nonzero finite length.) Consider the S_r- modules $N_i = (SM^{r'} \cap S_r) + (S_r M_r)^i$. They lie between S_r and $SM_{r'} \cap S_r$, and form a descending chain. Consider $i = 1, 1+b, 1+b+b^2, \ldots$. As the length of $S_r/(SM_{r'} \cap S_r)$ is n, we can find $1 \leq i$, $bi < j \leq b^{n+1}$, with $N_i = N_j$.

By Remark 10.14, $SM^{rb} \subset SM_r$, so $SM^{rbi} \subset S(M_r)^i \subset SN_i$. On the other hand, $SN_j \subset SM^{rj}$, provided $rj < r'$. But $rb^{n+1} < r'$, and $N_i = N_j$; so $SM^{rbi} \subset SM^{rj} \subset SM^{rbi+1}$. By Nakayama, $SM^{rbi} = 0$. So $SM^{rb^{n+1}} = 0$.

In particular $(S_r M_r)^{b^{n+1}} = 0$. As M_r is also generated by b elements, $(S_r M_r)^l$ can be generated by $\leq b^l$ elements, so S_r has length $\leq$ $\mathit{length}(S_r/M_r) + \mathit{length}(M_r/(S_r M_r)^2) + \ldots \leq 1 + b + b^2 + \ldots + b^{b^{n+1}-1} < b^{b^{n+1}}$. □

10.3.3 Relative reduction multiplicity

We first observe a natural relationship between transformally radicial extensions and *relative* reduction multiplicity. The proof is less straightforward than it ought to be; one reason is that the notion of multiplicity of a point on a scheme is somewhat delicate, and does not easily permit devissage. Recall that when q is a power of the prime p, K_q denotes a (large) algebraically closed field of characteristic p, endowed with the map $x \mapsto x^q$. For any scheme Y, $Y(K_q)$ denotes the set of K_q-valued points of Y, i.e. difference scheme maps $y : \mathrm{Spec}^\sigma K_q \to Y$. Recall also that $X_{q,y}$ denotes $M_q(X) \times_{f,y} \mathrm{Spec}^\sigma (K_q)$.

Definition 10.17 *Let $f : X \to Y$ be a morphism of difference schemes.* X has reduction multiplicity $\leq k$ over Y *if for some integer B, for all large prime powers q, and all $y \in Y(K_q)$, and $z \in X_{q,y}$,*

$\mathrm{Mult}_z X_{q,y} \leq Bq^k$

Definition 10.18 *Let X be a difference scheme of finite type over a finitely generated difference field K. We will say that X is generically of k-bounded reduction multiplicity over K if there exists a finitely generated difference domain $D \subset K$, and a difference scheme X_0 of finite type over D, with $X \simeq X_0 \times_{\mathrm{Spec}^\sigma D} \mathrm{Spec}^\sigma K$, such that X_0 is of reduction multiplicity $\leq k$ over* $\mathrm{Spec}^\sigma D$.

Lemma 10.19 *Let X be a difference scheme of finite type over a Noetherian difference scheme Y. Assume X is transformally radicial over Y, and of total dimension $\leq e$ over Y. Then X has reduction multiplicity $\leq e$ over Y*

Proof

We assume inductively the lemma holds for $X_{Y'}$ over Y' for any proper closed $Y' \subset Y$. Note that if the lemma holds for $X_{Y'}$ as well as for $X_{Y \setminus Y'}$, then it holds for X. Thus it suffices to prove the lemma for any open subset of Y. We may thus assume Y is irreducible; and indeed that $Y = \mathrm{Spec}^{\sigma} R$, R a difference domain; and it suffices to prove the lemma for some difference localization R' of R in place of R.

The case of X is more delicate; we may still assume the lemma is true for any proper closed subset; but if the lemma holds for a closed subset and its complement, it is not clear that it holds for X. Still if X is a union of *open* subsets, it suffices to prove the lemma for each separately. Thus we may take $X = \mathrm{Spec}^{\sigma} S$, S a finitely generated transformally radicial R-algebra. We may assume S is well-mixed.

Claim 1 Let R' be a difference R-subalgebra of S, finitely generated as an R-module. If the conclusion of the lemma holds for X over $\mathrm{Spec}^{\sigma} R'$, then it holds for X over $\mathrm{Spec}^{\sigma} R$.

Proof For any $y \in M_q(\mathrm{Spec}^{\sigma} R)(L)$ and $y_1 \in (\mathrm{Spec}^{\sigma} R')_{q,y}$, there exists $d \in \mathbb{N}$ with

$$\mathrm{Mult}_{y_1}(\mathrm{Spec}^{\sigma} R')_{q,y} \leq d$$

By assumption, for some b_1,

$$\mathrm{Mult}_z X_{q,y_1} \leq b_1 q^e$$

So by 10.12, with $b = db_1$, $\mathrm{Mult}_z X_{q,y} \leq bq^e$. □

Find a sequence of subrings $R = S_0 \subset S_1 \subset \ldots \subset S_n = S$ such that $S_{k+1} = S_k[a_k]$ and $\sigma(a_k) \in S_k$. We will also use induction also on the length n of this chain.

Let L be the field of fractions of R,$S_L = L \otimes_R S$. Effecting a finite localization of R, we may assume S embeds into S_L. (The ideal of polynomials over L vanishing at $(a_1, \ldots, a_k)$ can be taken to be generated by polynomials over R.) We will further use induction on the Noetherian rank of S_L.

Consider the ideal I_L of $S_L = L \otimes_R S$ generated by elements a with $\sigma(a) = 0, a^2 = 0$. I_L is generated by finitely many elements $b_1, \ldots, b_m$. After a finite localization of R, we may assume that the b_i lie in S. Let $R' = R[b_1, \ldots, b_m]$. R' is finitely generated over R as an R-module. By the claim, it suffices to prove the lemma for S over R'. Every prime of R' must contain the b_i, so this amounts to proving the lemma for $S' = S/I_L$. If $I_L \neq 0$, then $S'_L = S' \otimes_R L$ has smaller Noetherian rank than S_L; so using induction on this ordinal, we have the lemma for S' and hence for S. Thus we may assume $I_L = 0$.

So S_L has no nonzero elements a with $\sigma(a) = a^2 = 0$. (Since S embeds in S_L, S has no such elements either.)

Next, suppose S_L has a zero-divisor c with $\sigma(c) \in L$. If $\sigma(c) \neq 0$, then passing to a finite localization we may assume $\sigma(c)$ is a unit in R; thus if $cd = 0$ then $\sigma(d) = 0$; so replacing c by d if necessary we may assume $\sigma(c) = 0$. Let $J = \{s \in S_L : cs = 0\}$. Then $S_L c$,J are nontrivial difference ideals of S_L, and after further localization of R, $J \cap S$ is nontrivial. The lemma holds for $S/(J \cap S)$ and for S/Sc by Noetherian induction. Moreover if $e \in S_L c \cap J$ then $e^2 = 0$ and $\sigma(e) = 0$, so $e = 0$ by the previous paragraph. So $S_L c \cap J = 0$. The lemma follows for S.

Thus we may assume there are no zero divisors $c \in S_L$ with $\sigma(c) \in L$.

Consider first the case that there exists a polynomial $F \in R[X]$ with nonzero leading coefficient r, such that $F(a_1) = 0$. By the remark preceding the claim, it suffices to prove the lemma for $S[r^{-1}]$ over $R[r^{-1}]$; so we may assume F is monic. In this case S_1 is a finitely generated R-module. By the induction hypothesis, the lemma holds true for S over S_1. But now by the Claim it holds for S over R.

Assume now there is no such F; so S_1 as an R-algebra is the polynomial ring in one variable. Moreover for $f \in R[X]$, $f(a_1)$ is not a zero-divisor in S. And S has Krull dimension e over R. It follows algebraically that - after replacing R by a finite localization - *every* fiber of the map of algebraic schemes $\operatorname{Spec} S \to \operatorname{Spec} R[a_1]$ has dimension $\leq e-1$.

Let $Y' = \operatorname{Spec}^\sigma S_1$. Clearly for any q and any $y \in M_qY$, and $y' \in Y'_{q,y}$,

$\mathrm{Mult}_{y'} Y'_{q,y} \leq q$

Using the case $e-1$ assumed inductively, there exists b such that for any large q, and any $y' \in Y'(L)$, and $z \in X_{q,y'}$,

$\mathrm{Mult}_z X_{q,y'} \leq bq^{e-1}$

Let $y \in M_qY(L)$, $z \in X_{q,y}$; let $y' \in Y'_{q,y}$ be the intermediate. We then have $\mathrm{Mult}_z X_{q,y'} \leq bq^{e-1}$ and $\mathrm{Mult}_{y'} Y'_{q,y} \leq q$; by 10.12,

$\mathrm{Mult}_z X_{q,y} \leq bq^e$

the lemma follows. □

10.3.4 Transformal multiplicity and reduction multiplicity

Lemma 10.20 *Let X be a difference scheme of finite type over a finitely generated difference domain D. Assume X is of generic transformal multiplicity $\leq n$. Then there exists a nonempty open $Y \subset \operatorname{Spec}^\sigma D$ and a bound b_0 such that for all inversive difference fields L and L -valued points y of Y and x of $B_{n+1}(X_y)$, for every local ring R of $(X_y)_x$, $\sigma^n(R)$ is a finite-dimensional L-space, of dimension $\leq b_0$.*

Proof Let K be the inversive closure of the field of fractions of D. We may take $X = \operatorname{Spec}^\sigma S$. Let $S_K = S \otimes_D K$.

We first consider the case of generic $y \in \operatorname{Spec}^\sigma D$, i.e. of difference field extensions L of K. Let $S_L = S \otimes_D L$ (so $\operatorname{Spec}^\sigma S_L = X_y$). Let x be an L-valued point of $B_{n+1}(X_y)$, i.e. $x : \sigma^{n+1}(S_L) \to L$ a difference ring homomorphism. Let $x_k = x|\sigma^{n+1}(S_K)$.

We have a natural homomorphism

$$S_K \otimes_{x_k} L \to S_L \otimes_x L$$

and it is easily seen to be surjective. By 5.12, $\sigma^n(S_K \otimes_{x_k} L)$ has total dimension 0. Thus the homomorphic image $\sigma^n(S_L \otimes_x L)$ also has total dimension 0. By 5.4, $S_L \otimes_x L$ is a finitely generated L-algebra, hence so is $\sigma^n(S_L \otimes_x L)$. By 4.5, $\dim_L(\sigma^n(S_L \otimes_x L)) < \infty$.

It follows by compactness that for some $d_0 \in D$ and b_0, for all inversive difference fields L and difference ring homomorphisms $y : D \to L$ with $y(d_{0)} \neq 0$, and all L-difference algebra homomorphisms $x : \sigma^{n+1}(S_y) \to L$, $\dim_L(\sigma^n(S_y \otimes_x L)) < b_0$. To see this, let F_0 be a finite set of generators of S as a difference D -algebra, $F = F_0 \cup \ldots \cup \sigma^n(F_0)$. Then the image of F generates $S_y \otimes_x L$ as an L-algebra, for any x, y. Let $F_1 = F, \ldots, F_{m+1} = FF_m$. Then $\dim_L(\sigma^n(S_y \otimes_x L)) < \infty$ iff for some m,

(*) for any $c, d \in \sigma^n F_m$, the image of cd in $S_y \otimes_x L$ is in the L-span of the image of $\sigma^n F_m$.

Let y_i, x_i, L_i be a sequence of such triples, with y_i approaching the generic point of $\mathrm{Spec}^\sigma D$; let $S_i = S_{y_i} \otimes_{x_i} L_i$; we have to show that $\dim_{L_i}(\sigma^n(S_i))$ is bounded. Let $(y_\infty, x_\infty, L_\infty, S_\infty)$ be an nonprincipal ultraproduct of the (x_i, y_i, L_i). Then $\dim_L(\sigma^n(S_{y_\infty} \otimes_{x_\infty} L_\infty)) =< \infty$, so a fortiori the image of this ring in S_∞ has finite dimension over L_∞. Thus for some m, (*) above holds for the image of F_m in S_∞, hence also for the image of F_m in S_i, for almost all i. This proves that $\dim_{L_i}(\sigma^n(S_i))$ is bounded, as required.

The existence of the open difference subscheme Y and the bound b_0 now follow by a standard compactness argument.

□

Corollary 10.21 *Let $f : X \to Y$ be a a morphism of difference schemes of finite type, of relative transformal multiplicity $\leq n$. Then there exists $b \in \mathbb{N}$ such that for all prime powers $q > b$, for any perfect field L and any $y \in (M_q Y)(L)$, every closed point $x \in M_q(X_y)$ has multiplicity $\leq bq^n$ on $M_q(X_y)$.*

Proof We may assume $X = \mathrm{Spec}^\sigma R$, $Y = \mathrm{Spec}^\sigma D$, R a finitely generated D-algebra. By lemma 10.10, the local rings of $M_q(X_y)$, $y : M_q(D) \to L$, L a perfect field, have maximal ideals generated by a bounded number of elements. (D itself is generated by a finite number e of elements, so that quotient domains of D have fields of fractions k with $[k^q : k] \leq q^e$, but we just consider perfect L, so actually 10.10 applies with $e = 0$.) Thus Lemma 10.16 applies, and translates Lemma 10.20 to say that $x' = x|\sigma^{n+1}(R \otimes_y L)$ has bounded multiplicity on $M_q(B_{n+1}(X_y))$. By Lemma 10.19, $X/(Y \times B_{n+1}(X))$ has reduction multiplicity $\leq n$. So $Mult_x M_q(X)_{y,x'} \leq O(q^n)$. By Lemma 10.12, x has multiplicity $O(q^n)$ on $M_q(X_y)$. □

The corollary 10.22 could also be obtained using Lemmas 6.2 and 11.23 (Bezout theorem methods.) One could take things up from that point with 10.19. However, the above methods permit estimating the multiplicity on X of points on $X' \subset X$, where X' has smaller total dimension, as in 10.24. This does not seem apparent from the Bezout approach.

Corollary 10.22 *Let X be a difference scheme of finite type over a Noetherian difference scheme Y, of total dimension d over Y. Then there exists b such that for all large enough prime powers q, and all $y \in M_q(Y)$, the zero-dimensional scheme $X_{q,y}$ over L_y has size at most bq^d.*

Proof By the usual Noetherian induction on Y we may assume $Y = \mathrm{Spec}^\sigma D$, D a difference domain with field of fractions K; and we may pass to a localization of D by a finite set. Moreover base change will not change the total relative dimension or the size, so we may replace D by a finitely generated extension within the inversive hull K' of K^{alg}. Let $X' = X \otimes_Y \mathrm{Spec}^\sigma K'$. We enlarge D within K so that the $Mlt_k X'$ and their components $X_{k,j}$ are defined over D. The components $X_{k,j}$ have reduced total dimension $\leq d - k$; so by 10.8, for all $y \in Y$ and all sufficiently large q, $(X_{k,j})_{y,q}$ has $O(q^{d-k})$ points, and by 5.9, away from

$Mlt_{k+1}(X)$, these points have multiplicity $\leq O(q^k)$ on M_qX. Thus $X_{q,y}$ is divided into d groups $(Z_kX)_{q,y}$, the k'th having at most $O(q^{d-k})$ of multiplicity (on M_qX) at most $O(q^k)$; so altogether there are $O(q^d)$ points, multiplicity counted.

Corollary 10.23 *Let X_0 be a difference scheme of finite type over a difference field K. Assume X_0 has total dimension d, and that no weak component of X_0 of total dimension d has transformal multiplicity > 0. Then there exist a finitely generated subdomain D of K, a difference scheme X over D with $X_0 = X \times_{\mathrm{Spec}^\sigma D} \mathrm{Spec}^\sigma K$, and an integer b, so that for all large enough prime powers q, and all $y \in M_q(\mathrm{Spec}^\sigma D)$, the zero-dimensional scheme $X_{q,y}$ over L_y has size at most bq^d; moreover all but bq^{d-1} points of $X_{q,y}$ (counting multiplicities) have multiplicity $< b$.*

Proof As in the proof of 10.22, we consider reduced subschemes Y of $Z_kX = Mlt_k(X) \setminus Mlt_{k+1}(X)$. Here however, for $k \geq 1$, we use 5.9(3) to conclude that the reduced total dimension of Y is $\leq d-k-1$. This gives an $O(q^{d-k-1})$-bound on multiplicities, by 5.9. As we are considering Y away from Mlt_{k+1}, we obtain $O(q^{d-1})$ points counted with their multiplicities on X. □

Corollary 10.24 *Let X_0 be as in Corollary 10.23. For any proper subscheme X' of X, of total dimension $d' < d$, the number of points of X', counted with their multiplicities on X, is $O(q^{d-1})$.*

Proof By Corollary 10.23, only $O(q^{d-1})$ have high multiplicity on $X_{q,y}$. These can therefore be ignored. The rest have multiplicity $O(1)$ on $X_{q,y}$, and by Corollary 10.22, their number is $O(q^{d'})$. Thus even with multiplicity, the number is $O(q^{d-1})$.

Part II

Intersections with Frobenius

11 Geometric Preliminaries

We show here that Theorem 1B is invariant under birational changes (11.25) and in the appropriate sense under taking finite covers (11.27). We use this together with de Jong's version of resolution of singularities to reduce to the smooth case, where intersection theory applies. We deal with the possible inseparability in 1B, and note a crude initial bound (11.23) on the set whose size we are trying to estimate.

Our basic reference, here and in later sections, is [Fulton]. We will begin with some basic lemmas on proper and improper intersections, degrees and correspondences.

11.1 Proper intersections and moving lemmas

Let X be a nonsingular variety over an algebraically closed field. We write $\mathrm{codim}_X(U)$ for $\dim(X) - \dim(U)$. A variety (or cycle) is said to have pure (co)dimension k if each irreducible component has (co)dimension k. Two subvarieties U,V of X of pure codimension k,l respectively are said to meet *properly* if either $U \cap V = \emptyset$, or $\operatorname{codim}_X(U \cap V) \geq k + l$. By the dimension theorem for smooth varieties, each component of $U \cap V$ must have codimension at most $k + l$; thus the intersection is proper iff each (nonempty) component of $U \cap V$ has codimension $k + l$ in X. The properness of the intersection is equivalent to the properness of intersection of all irreducible components of U with those of V.

Lemma 11.1 *Let U, V, W be pure-dimensional varieties. Assume U, V meet properly, and $U \cap V, W$ meet properly. Then U meets $V \cap W$ properly.*

Proof Let U, V, W have codimensions k, l, m. Then $U \cap V = \emptyset$ or $\operatorname{codim}_X(U \cap V) = k + l$; so $U \cap V \cap W = \emptyset$ or $\operatorname{codim}_X(U \cap V \cap W) = k + l + m$; this last condition is equivalent to the assumptions, and is symmetric in U, V, W.

Lemma 11.2 *Assume X, Y are smooth varieties. Let T be an irreducible subvariety of $X \times Y$, and let $\pi[k]T$ be the subvariety of X, whose points are $\{p \in X : \dim(T \cap (\{p\} \times Y)) \geq k\}$. If U is a subvariety of X meeting each component of each $\pi[k]T$ properly, then $U \times Y$ meets T properly. If $U \times Y$ meets T properly, then U meets πT properly.*

Proof Note that $\operatorname{codim}_{X \times Y}(U \times Y) = \operatorname{codim}_X(U)$; let c be the common value. Let W be an irreducible component of $(U \times Y) \cap T$. Say $\dim(W) = \dim(\pi W) + k$. Then $\pi W \subset \pi[k]T$. So $\pi W \subset \pi[k]T \cap U$,and $\dim(\pi W) \leq \dim(\pi[k]T \cap U) \leq \dim(\pi[k]T) - c \leq \dim(T) - k - c$ Thus $\dim(W) \leq \dim(T) - c$. This shows that $U \times Y$ meets T properly.

For the remaining statement, let $k = \dim(T) - \dim(\pi T)$. Then $\pi T = \pi[k]T$. So $\dim(T) - \operatorname{codim}_X(U) \geq \dim((U \times Y) \cap T) \geq k + \dim(U \cap \pi T)$; hence $\dim(\pi T) - \operatorname{codim}_X(U) \geq \dim(U \cap \pi T)$.

Lemma 11.3 *Let $U \subset (X \times Y)$, $V \subset (Y \times Z)$ be complete varieties, all of pure dimension d. Let $W = pr_{XZ}((U \times Z) \cap (X \times V))$. Let T be a subvariety of $X \times Z$. Assume*

U meets each component of $X \times pr_Y[l]V$ and of $pr_X[l]T \times Y$ properly for each l;

V meets each component of $pr_{YZ}[l]((T \times Y) \cap (U \times Z))$ properly;

Then $U \times Z$ meets $X \times V$ properly, and W meets T properly.

Proof

1. By the first assumption and 11.2, $U \times Z$ meets $X \times V$ properly.
2. Similarly, $T \times Y$ meets $U \times Z$ properly.
3. By the second assumption, $X \times V$ meets $((T \times Y) \cap (U \times Z))$ properly.
4. By 11.1, and (2),(3), $T \times Y$ meets $(U \times Z) \cap (X \times V)$ properly.
5. By the last statement of 11.2, T meets W properly.

Notation 11.4 *Let X be a smooth algebraic variety over an algebraically closed field. A* cycle *is a formal sum, with integer coefficients, of irreducible subvarieties of X.*

$B^i(X)$ denotes the group of cycles whose components have codimension i, up to algebraic equivalence ([Fulton] 10.3). $B^*(X)$ is the direct sum of the $B^i(X)$.

To each subscheme U of X, one associates a cycle $[U]$; it is the sum of the irreducible components of U, with certain nonnegative integer coefficients ([Fulton], 1.5).

For any cycle I, write $I = \Sigma_k \{I\}_k$, where $\{I\}_k$ is a k-dimensional cycle.

$B^i(X)$ denotes the group of cycles whose components have codimension i, up to algebraic equivalence ([Fulton] 10.3). $B^*(X)$ is the direct sum of the $B^i(X)$. An operation $\cdot$ can be defined on $B^*(X)$, making it into a commutative graded ring with unit. This operation is determined by the fact that $[U] \cdot [V] = [U \cap V]$ when U, V are irreducible subvarieties of X meeting properly.

Most of the time, we will be able to use the finer notion of rational equivalence; but since we are primarily interested in intersection theory, the difference between two algebraically equivalent cycles will not concern us. In particular, a 0-cycle $\sum n_i p_i$ is determined up to algebraic equivalence by $\sum n_i \in \mathbb{Z}$; we will identify the group of 0-cycles with $\mathbb{Z}$.

Lemma 11.5 *Let Y be a smooth projective variety over an algebraically closed field k, and let V be an irreducible subvariety of Y, and $f : Y \to V$ a flat morphism. There exists a purely transcendental extension K of k, and a cycle V' on Y defined over K, such that V, V' are rationally equivalent, and for any subvariety W of Y defined over k, each component C of V' meets W properly, and $f(C) = V$.*

Proof This is a version of the moving lemma described in [Fulton] 11.4.1, [Hyot], and the classical proof works. Say $Y \subset \mathbb{P}^N$; let G be the Grassmanian of linear subspaces of $\mathbb{P}^N$ of codimension one more than the codimension of U in Y. Pick a point p of G generic over k, representing a linear subspace L, and also $g \in PGL_N$, generic over $k(p)$. Given any k-subvariety U of Y, form the cone $C(L, U)$, and consider also $gC(L, U)$. Write $C(L, U) \cap Y = U + \sum m_i U_i''$, $gC(L, U) \cap Y = \sum m_j' U_j'$, and let $U[1] = \sum m_j' U_j' - \sum m_i U_i''$. If U is not a variety but a cycle, a formal sum $\sum U = n_i U_i$ of varieties of the same dimension, define $U = \sum n_i U_i[1]$. So $U[1]$ is rationally equivalent to U, is defined over a purely transcendental extension of k, and for any subvariety W of Y defined over k, each component C of $U[1]$ either meets W properly, or satisfies $\dim(C \cap W) < \dim(C' \cap W)$ for some component C' of U. Let $U[2] = U[1][1], U[3] = U[2][1]$, etc. Then $U[m]$ is rationally equivalent to U, is defined over a purely transcendental extension of k, and for any subvariety W of Y defined over k, each component C of $U[m]$ either meets W properly, or satisfies $\dim(C \cap W) \leq \dim(C' \cap W) - m$ for some component C' of U. It follows that $U[m+1]$ satisfies our requirements. The argument that f maps each component onto V is given separately, in 11.6; to apply it, embed $\mathbb{P}^N$ into some larger $\mathbb{P}^{N'}$ first by the d-uple embedding, so that any four distinct points of $\mathbb{P}^N$ are linearly independent in $\mathbb{P}^{N'}$. □

Lemma 11.6 *Let $U \subset Y \subset \mathbb{P}^N$ be projective varieties, with no three or four points of Y linearly dependent in $\mathbb{P}^N$. Let $f : Y \to V$ a flat morphism to an irreducible variety,* $\dim(Y) = n$, $\dim(U) = \dim(V) = d$.

1. *Let G be the Grassmanian of linear subspaces of $\mathbb{P}^N$ of codimension $n+1$. For $L \in G$, let $C(L, U)$ be the cone on U with center L. Then for generic $L \in G$, and any component U' of $C(L, U)$ other than U, $f(U') = V$.*
2. *Let W be any subvariety of $\mathbb{P}^N$ with of codimension $n-d$, and let $g \in PGL_N$ be generic. Then for any component U' of $gW \cap Y$, $f(U') = V$.*

Proof (Note that the hypotheses imply $n+1 < N$, or assume this.) Let

$$M_2 = \{(L, y, y', p, p') \in G \times (Y \setminus U)^2 \times (\mathbb{P}^N)^2 : p \neq p' \in L, y \in C(p, U), y' \in C(p', U), fy = fy'\}$$

Then M_2 projects to $G \times Y^2$, and the image of the projection contains

$$M = \{(L, y, y') \in G \times (Y \setminus U)^2 : y \neq y' \in C(L, U), fy = fy'\}$$

Indeed if $(L, y, y') \in M$, then there exist $p, p' \in L$ and $u, u' \in U$ with y, u, p and y', u', p' colinear. Since no three or four distinct points of Y are linearly dependent in $\mathbb{P}^N$, we must have $p \neq p'$.

So $\dim(M) \le \dim(M_2)$.

Let $G_{p,p'} = \{L \in G : p, p' \in L\}$; we have $\dim G_{p,p'} + 2N \le \dim(G) + 2(N - (n+1))$.

We will also use: if $y \in C(p, U)$, then $p \in C(y, U)$, and $\dim C(y, U) = \dim(U) + 1$.

Now compute, using the maps $(L, y, y', p, p') \mapsto (y, y', p, p) \mapsto (y, y') \mapsto fy$:

$$\dim(M_2) \le \dim(V) + 2(\dim(Y) - \dim(V)) + 2(d+1) + \dim G_{p,p'} =$$

$$= d + 2(n-d) + 2(d+1) + \dim(G) - 2(n+1) = d + \dim(G).$$

Thus we see that for generic $L \in G$, $\dim\{(y, y') : (L, y, y') \in M\} \le d$. But let $\widetilde{U} = U' \setminus U$; if $\dim f(U') < d$, then $\dim\{(y, y') \in \widetilde{U}^2 : fu = fu'\} \ge (\dim f(U')) + 2(\dim(U') - \dim f(U')) > d$, a contradiction. So $\dim f(U) \ge f(U') = d$, and thus $f(U) = V$.

For (2), a similar computation shows that $\dim\{(g, y, y') : y \neq y' \in Y \cap gW, fy = fy'\} \le d$, and we conclude as above.

□

Lemma 11.7 *Let V be a smooth projective variety over a difference field k, and let S be a subvariety of $Y = (V \times V^\sigma)$, $\dim(S) = \dim(V) = d$, with $\mathrm{pr}_0 : S \to V, \mathrm{pr}_1 : S \to V^\sigma$ dominant.*

There exists a cycle $S' = \sum_i m_i[U_i]$ on Y defined over k, such that S, S' are rationally equivalent, and $X = [\sigma]_K U_i \star \Sigma$ has transformal dimension 0 ; in fact (cf. 6.6) $pr_V[1](U_i) \cap X = \emptyset$ (as a difference scheme).

Moreover, each U_i as well as the varieties involved in the rational equivalence of S with S' all have dominant projections to V and to V^σ.

Proof Using 11.5, find a purely transcendental extension $k(s)$ of k and a cycle S' on Y, rationally equivalent to S, such that any subvariety of Y defined over k^{alg} meets S' properly, and such that each component of S' maps dominantly to V. We will show that S' has the required properties. It then follows easily, by specializing s into k (avoiding finitely many proper Zariski closed sets), that such a cycle and such a rational equivalence exist over k too.

We may assume that k is inversive.

Let $s_0 = s$, and $K = k(s_0, s_1, s_2, \ldots)$, $\sigma^n(s_0) = s_n$ (with $s_0, s_1, \ldots$ algebraically independent over k.) Let $V_n = \sigma^n(V)$; note V_n is defined over k.

Let S_0 be any component of S, $S_n = \sigma^n(S_0)$. Then any subvariety of $\sigma^n(Y)$ defined over k^{alg} meets S_n properly . Since $\sigma^n(Y)$ and S_n are defined over $k(s_n)$, and $k(s_0, \ldots, s_{n-1}, s_{n+1}, s_{n+2}, \ldots)^{alg}$ is linearly free from $k^{alg}(s_n)$ over k^{alg}, in fact any subvariety of $\sigma^n(Y)$ defined over $k(s_0, \ldots, s_{n-1}, s_{n+1}, s_{n+2}, \ldots)^{alg}$ meets S_n properly.

Claim Let W_1 be any proper subvariety of S_0, defined over $k(s_0)$. Then $W_1 \star \Sigma = \emptyset$.

To prove the claim, let $c = (d - \dim(W_1))/2 > 0$. We will inductively define W_n satisfying:

1. $W_n \subset V_0 \times \ldots \times V_{2^n-1}$
2. $\dim(W_n) \leq d - 2^n c$ (or $W_n = \emptyset$).
3. W_n is defined over $k(s_0, \ldots, s_{2^n-2})$
4. $(W_1 \star \Sigma)[2^n - 1] \subset W_n$

When $2^n c > d$, (2) forces $W_n = \emptyset$, so by (4), $([\sigma]W_1 \star \Sigma) = \emptyset$. For $n = 1$, (1-4) hold by assumption. Assume they hold for n. Then $(W_n)^{\sigma^{2^n}} \subset V_{2^n} \times \ldots \times V_{2^{n+1}-1}$, has dimension $\leq d - 2^n c$, and is defined over $k(s_{2^n}, \ldots, s_{2^{n+1}-2})$.

Thus $W^* = W_n \times (W_n)^{\sigma^{2^n}}$ has dimension $\leq 2d - 2^{n+1}c$, and is defined over $k(s_0, \ldots, s_{2^n-2}, s_{2^n}, \ldots, s_{2^{n+1}-2})$. Let $\pi : (V_0 \times \ldots \times V_{2^{n+1}-1}) \rightarrow (V_{2^n} \times V_{2^n+1})$ be the projection. Then $\pi[k]W$ is defined over $k(s_0, s_{2^n-2}, s_{2^n}, s_{2^{n+1}-2})$, hence meets S_{2^n-1} properly; by 11.2, $\pi^{-1}S_{2^n-1}$ meets W^* properly. Let $W_{n+1} = \pi^{-1}S_{2^n-1} \cap W^*$. Then $\dim W_{n+1} \leq \dim(W^*) - d = d - 2^{n+1}c$, and (1)-(4) hold. This proves the claim.

In particular, as $pr_0 : S_0 \rightarrow V$ dominantly, and $\dim(S_0) = \dim(V)$, $pr_0[1](S_0) \neq V$; so letting $W_1 = S_0 \cap {pr_0}^{-1}(pr_0[1](S_0))$, we have $\dim(W_1) < d$, and we can apply the claim.

□

Remark 11.8

An alternative treatment of a moving lemma for difference schemes can be given by transposing the proof of [Fulton],11.4.1 to difference algebra, almost verbatim. The methods used there - projective cones, moving via the group of automorphisms of projective space, "counting constants" - work very well for difference varieties and transformal dimension in place of varieties and dimension. One can also use the dimension growth sequence to get finer results.

11.2 Degrees of cycles

In general, the intersection product gives no direct information about improper intersections. In products of projective spaces, we can obtain such information, as in [Fulton], Example 8.4.6.

Notation 11.9 *Let H, H' be hyperplane divisors on $\mathbb{P}^n$, $\mathbb{P}^m$, respectively, $s = pr_1^* H$, $t = pr_2^* H'$. For a subvariety U of $\mathbb{P}^n \times \mathbb{P}^m$, $U = \Sigma a_{ij} s^i t^j$ as a cycle up to rational equivalence; the a_{ij} are called the bidegrees. ([Fulton], Example 8.4.4). If U is of pure dimension k and $i + j = k$, we have $a_{ij} = (U \cdot s^{n-i} t^{m-j})$. By taking a representative in general position, it follows that the bidegrees a_{ij} are non-negative integers.*

In a product of more than two projective spaces, multi-degrees are defined analogously.

A divisor H on a variety X is said to be very ample *if it is the pullback of a hyperplane divisor of projective space, under some projective embedding of X. The projective degree of $U \in B^i(X)$ under this embedding can then be expressed as $U \cdot_X H^i$, and written* $\deg_H(U)$.

We will denote projection from a product $X \times X' \times \ldots$ to some of its factors $X \times X'$ by $pr_{X,X'}$.

Notation 11.10 *If A, B are cycles on P, a multi-projective spaces, write $A \leq B$ if $B - A$ is effective. Equivalently, if $\{V\}^{ijk} = V \cdot h_1^i h_2^j h_3^k$ are the multi-degrees of a cycle V on P,*

$$\{A\}^{ijk} \leq \{B\}^{ijk}$$

for each i, j, k. Note that this partial order is preserved by the intersection product.

Let H_i be the hyperplane divisor on $\mathbb{P}^n$, $H = pr_1^* H_1 + pr_2^* H_2$. Then H is very ample. (It corresponds to the Segre embedding of $\mathbb{P}^n \times \mathbb{P}^n$ in $\mathbb{P}^N$.)

The following lemmas will be used in §12.

Lemma 11.11 *Let D_i be a very ample divisor on a projective variety X_i, $Y = X_1 \times X_2$,* $\dim(X_i) = d_i$, $D = \mathrm{pr}_1{}^* D_1 + pr_2{}^* D_2$. *Let $\pi = pr_1 : Y \to X_1$ be the projection.*

1. *Let U be a k-dimensional subvariety of X_1. Then* $\deg_D(U \times X_2) \leq \binom{k+d_2}{k} \deg_{D_2}(X_2) \deg_{D_1}(U)$
2. *Let W be an l-dimensional subvariety of Y. Then* $\deg_{D_1}(\pi_* W) \leq \deg_D(W)$

Proof

1. $\dim(U \times X_2) = k + d_2$, and $\deg_D(U \times X_2) = (\mathrm{pr}_1{}^* D_1 + pr_2{}^* D_2)^{k+d_2}(U \times X_2) ==$ $\sum_{i+j=k+d_2} \binom{i+j}{i} (D_1{}^i D_2{}^j \cdot U \times X_2)$. The only nonzero factor is $i = k, j = d_2$.
2. $\deg_D(W) = \sum_{i+j=l} \binom{l}{i} (pr_1{}^* D_1{}^i pr_2{}^* D_2{}^j) \cdot W \geq pr_1{}^* D_1{}^l \cdot W = D_1{}^l \cdot pr_* W$, using the projection formula for divisors ([Fulton] 2.3). The last quantity equals $\deg_{D_1}(\pi_* W)$.

□

We can obtain some information about a proper intersection in a smooth variety Y by viewing it as an improper intersection in projective (or multiprojective) space. Note that two rationally equivalent cycles on Y are a fortiori rationally equivalent on the ambient projective space, so have the same degrees.

Lemma 11.12 *Let Y be a smooth subvariety of projective space $\mathbb{P}^m$. Let U, V be properly intersecting subvarieties of Y, $W = U \cdot_Y V$. Then* $\deg(W) \leq \deg(U) \deg(V)$

Proof We may represent W by an effective cycle $\sum m_i W_i$, where W_i are the components of $U \cap V$ and $m_i = i(W_i, U \cdot V; Y)$ are the intersection multiplicities.

Let L be a generic linear subspace of $\mathbb{P}^m$ of dimension $m - \dim(Y) - 1$, and let C be the cone over U with vertex L, cf. [Fulton] Example 11.4.1. C is a subvariety of $\mathbb{P}^m$, "union of all lines meeting U and L".

Note that $\deg(C) = \deg(U)$: take a generic linear space J of dimension complementary to C in $\mathbb{P}^m$, so that J meets C transversally in $\deg(C)$ points. The cone E on J with center L is (over the original base field) a generic linear space, and so meets U transversally, in $\deg(U)$ points. But both numbers equal the number of pairs $(p, q) \in J \times L$ such that p, q, r are colinear for some $r \in U$.

By [Fulton] Example 11.4.3, each W_i is a proper component of $C \cap W$, and $m_i = i(W_i, C \cdot V; \mathbb{P}^m)$.

By the refined Bezout theorem [Fulton] Example 12.3.1, $\sum_i m_i \deg(W_i) \leq \deg(C) \deg(V)$. Thus $\deg(W) \leq \deg(U) \deg(V)$.

Notation 11.13 *Let D be a very ample divisor on Y, and let U be a cycle, or a rational equivalence class of cycles. Define*

$$|U|_D = \sup_S \inf_{U',U''} \deg_D(U') + \deg_D(U'')$$

where S ranges over all cycles of Y , and U',U'' range over all pairs of effective cycles such that U is rationally equivalent to $U' - U''$, and U',U'' meet S properly.

Corollary 11.14 *Let Y be a smooth variety, D a very ample divisor on Y. Let $|X| = |X|_D$. Let U,V be cycles on Y. Then $|U \cdot V| \leq |U||V|$*

Proof Let W be an arbitrary cycle on Y, of dimension $\dim(Y) - \dim(U) - \dim(V)$. Write $V = V' - V''$, with V, V' effective and meeting W properly, and $|V| = \deg(V') + \deg(V'')$. Write $U = U' - U''$ similarly, with U', U'' meeting $V' \cap W, V'' \cap W$ properly. Let X_1 be an effective cycle representing $U' \cdot V'$, supported on the components of $U' \cap V'$; and similarly X_2, X_3, X_4 for $U' \cap V''$,...,$U'' \cap V''$. Then by 11.1, the X_i meet W properly. By 11.12, $\deg(X_1) \leq \deg(U')\deg(V')$, etc. Thus $\sum_i |\deg(X_i)| \leq |U||V|$. Since W was arbitrary, $|U \cdot V| \leq |U||V|$.

Lemma 11.15 *For any cycle R on Y, and very ample H on Y, $|R|_H$ is finite.*

Proof Say D, Y, U are defined over an algebraically closed field k. Since k is an elementary submodel of an algebraically closed field K with $tr.deg._k K$ infinite, in the definition 11.13 of $|R|_H$, one can restrict the $sups_S$ to range over cycles S defined over k, while allowing the U', U'' to be defined over K. The finiteness is now immediate from Lemma 11.5. □

11.3 Correspondences

Let us recall the language of correspondences, (cf. [Fulton], 16.1). Let X,X' be smooth, complete varieties of dimension d over an algebraically closed field. A correspondence R on $X \times X'$ is a d-cycle on $X \times X'$; we will write somewhat incorrectly $R \subset (X \times X')$. The transpose of R is the corresponding cycle on $X' \times X$; it is denoted R^t. If T is a correspondence of $X' \times X''$, then one defines the composition by:

$T \circ R = pr_{X,X''}{}_*(pr_{X,X'}{}^*R \cdot pr_{X',X''}{}^*T)$

When the context does not make clear which product is in question, we will use $X \circ Y$ and $X^{\circ n}$ for the composition product and power, and $X \cdot Y$, $X^{\cdot n}$ for the intersection product and power.

Notation 11.16 *Suppose X and X' are smooth complete varieties of dimension d, given together with very ample divisors H,H' on X,X'. Let $R \subset (X \times X')$ be a correspondence. Write*

$\deg_{cor}(R) = [p \times X'] \cdot R$

where p is a point of X

$\deg_{cor}(R)$ is intended to denote the degree of R as a correspondence. If R is an irreducible subvariety R and pr_X restricts to a dominant morphism $\pi : R \to X$, then $\deg_{cor}(R)$ is the degree of π. In particular that if R is the graph of a function, then $\deg_{cor}(\Phi) = 1$. By contrast, we will write $\deg(R)$ or $\deg_{cy}(R)$ for $\deg_{pr_1{}^*H+pr_2{}^*H'} R$. If $R = R_1 - R_2$ with R_1, R_2 effective (irredundantly), we let $\deg_{|cy|}(R) = \deg_{cy}(R_1) - \deg_{cy}(R_2)$.

Note that $\deg_{cor}(\Phi) = 1$, $\deg_{cor}(\Phi^t) = q^d$, $\deg_{cor}(S) = \delta$.

Lemma 11.17 *Let $R \subset X \times Y$ and $T \subset Y \times Z$ be correspondences. Computing their degrees with respect to very ample divisors H_X,H_Y,H_Z on X,Y,Z respectively:*

1. $\deg_{cor}(T \circ R) = \deg_{cor}(T)\deg_{cor}(R)$.
2. $\deg_{cor}(R) \leq \deg_{|cy|}(R)$
3. $|T \circ R| \leq \binom{2d}{d}^2 |T||R|$
4. *Suppose $Z = X$ and $H_X = H_Z$. Then*

$$((T \circ R) \cdot_{X^2} \Delta_X) = (R \cdot T^t) \leq |T^t||R|$$

Proof

1. Let P_{XY} be the correspondence $p \times Y$, where p is a point of X. Then $\deg_{cor}(R) = (P_{XY} \cdot R) = \Delta_X \cdot (P^t_{XY} \circ R)$. (See (4) below for the last equality.) For R the cycle of an irreducible variety, and hence in general, $P^t_{XY} \circ R$ is a multiple of P^t_{XX}. So

   $$P^t_{XY} \circ R = \deg_{cor}(R) P^t_{XX}$$

   Similarly $P^t_{YZ} \circ T = \deg_{cor}(T) P^t_{YY}$, so $P^t_{XZ} \circ T = P^t_{XY} \circ P^t_{YZ} \circ T = \deg_{cor}(T) P^t_{XY}$. Thus:

   $(P^t_{XZ} \circ T \circ R) = \deg_{cor}(T)(P^t_{XY} \circ R)$

   The desired formula follows. upon intersecting this with Δ_X.

2. Here we may assume $R \geq 0$, so that $\deg_{|cy|}(R) = \deg_{cy}(R)$. On X, $H^{\cdot n} = c[pt]$ for some $c \geq 1$; so

   $$\deg_{cor}(T) = (P_{XY} \cdot R) = c^{-1}(pr_X{}^* H)^n \cdot R \leq \deg_{cy}(R)$$

3. Both $R \circ T$ and $|R|$,$|T|$ depend on R,T only up to rational equivalence, so we may change R and T within their rational equivalence class. By 11.3, R,T may be replaced so that $\deg_{|cy|}(R) \leq |R|$, $\deg_{|cy|}(T) \leq |T|$, $(R \times Z) \cap (X \times T)$ is a proper intersection, and $pr_{X,Z}((R \times Z) \cap (X \times T))$ meets properly a given cycle S. Then $R \circ T = pr_{X,Z*}((R \times Z) \cap (X \times T))$, and by 11.11(1),11.12, and 11.11(2), $\deg_{|cy|}(W) \leq (\binom{2d}{d} \deg_{|cy|}(R) \binom{2d}{d} \deg_{|cy|}(T)| \leq \binom{2d}{d}^2 |R||T|$; and W meets S properly. Taking supremum over S, the statement follows.

4. The equality $(T \circ R) \cdot \Delta_X = (R \cdot T^t)$ follows from the projection formula, and the definition of composition of correspondences; both are equal to the triple intersection product $(R \times X) \cdot (X \times T) \cdot \Delta_{13}$.

   The inequality is clear from 11.12. □

11.4 Geometric statement of the uniformity

To formulate the theorem geometrically, we will need to consider V, q, and $S \subset V \times V^{\phi_q}$ all varying separately. Below, B will be the base over which V varies. B' will be the base for S; it need not be a subscheme of $B \times B$, though one loses little or nothing by thinking of that case. A scheme *over* B is a scheme S together with a morphism $\alpha : S \to B$. We will not always have a notation for α; Instead, if L is a field and $b \in B(L)$, we will write S_b for the fiber of S over the point b. If $U \to V$ is a dominant map of irreducible varieties over k, then the (purely inseparable) degree of U/V is the (purely inseparable) degree of $[k(U) : k(V)]$.

Notation 11.18 *1. B and B' are reduced, irreducible, separated schemes over $\mathbb{Z}[1/m]$ or $\mathbb{F}_p$. (Not necessarily absolutely irreducible.) A (base change) map $\beta : B' \to B^2$ is also given. We will refer to the cases as "characteristic 0" and "characteristic p", respectively.*

Let V be a scheme over B. View V^2 as a scheme over B^2. Let S be a ($B'-$) subscheme of $V^2 \times_{B^2} B'$.

For L a field and $b \in B'(L)$, denote $\beta(b) = (b_1, b_2) \in V^2(L)$. We will assume that V_{b_1},V_{b_2} and S_b are varieties, and view S_b as a subvariety of $V_{b_1} \times V_{b_2}$

2. *We further assume that for $b \in B'$, S_b,V_{b_1},V_{b_2} are absolutely irreducible; the projection map $S_b \to V_{b_2}$ is a quasi-finite map; $\dim(V_{b_1}) = \dim(V_{b_2}) = \dim(S_b) = d$, $\deg(S_b/V_{b_1}) = \delta < \infty$, and if B is over $\mathbb{F}_p$, the purely inseparable degree of S_b over V_{b_2} is δ'_p.*

3. *Let q be a prime power (power of p), $a \in B(K_q)$, $b \in B'(K_q)$ such that $\beta(b) = (a, \phi_q(a))$. In this situation, we denote:*

$$V_b(S, q) = \{c \in V_a(K_q) : (c, \phi_q(c)) \in S_b(K_q)\}$$

In this language, the asymptotic version of ACFA 1 is the following:

Theorem 1B' *Let B be a reduced, separated scheme of finite type over $\mathbb{Z}[1/m]$ or $\mathbb{F}_p$. Let B', V, S, β be as in 11.18. For any sufficiently large q, if $b \in B'(K_q)$, and $b_2 = \phi_q(b_1)$, then there exists $c \in V_{b_1}(K_q)$ with $(c, \phi_q(c)) \in S_b(K_q)$.*

While we need a mere existence statement, we see no way to prove it without going through a quantitative estimate. We formulate this estimate as follows. Note the similarity to the Lang-Weil estimates; these are the special case when the S_b are the diagonals.

Theorem 1B

Let assumptions be as in Theorem 1B'. Then there exists an open ($\mathbb{Z}$ or $\mathbb{F}_p$) subscheme B'' of B' and constants ρ and $\delta^* > 0$ such that if $b \in B''(K_q)$, $\beta(b) = (a, \phi_q(a))$, then $V_b(S, q)$ is finite, of cardinality

$$\#(V_b(S, q)) = \delta^* q^d + e \qquad \text{with } \ |e| \le \rho q^{d-\frac{1}{2}}$$

In fact $\delta^* = \delta/\delta'_p$. Each point occurs with multiplicity δ'_p, so that the number of points counted with multiplicity is $\delta q^d + O(q^{d-\frac{1}{2}})$.

Theorem 1B' follows from Theorem 1B by restriction to an open subset of V, and using Noetherian induction on B'.

11.5 Separability

Lemma 11.19 *In Theorem 1B, we may assume that $\delta'_p = 1$, i.e. that for $b \in B'$, S_b/V_{b_2} is separable.*

Proof If B is over $\mathbb{Z}$ (or $\mathbb{Z}[1/m]$,) this may simply be achieved by replacing B' by an open subscheme B'' so that for $b \in B''$, $deg(S_b/V_{b_2})$ is a constant δ', and then further by the open subscheme $B''' = B'' \otimes_{\mathbb{Z}[1/m]} \mathbb{Z}[1/(m\delta')]$. Suppose then that B is over $\mathbb{F}_p$, and $\delta'_p = p^l$. We will modify B' and S so as to obtain a similar situation with $\delta'_p = 1$; the modified objects will be denoted by a $^\sim$ but B and V are left the same.

Let $F_B : B^2 \to B^2$ be the map (Id, ϕ_{p^l}). Let $\widetilde{B}' = B' \times_{B^2} B^2$, where the implicit map $B^2 \to B^2$ is F_B. (Thus if B' is a subscheme of B^2, $\widetilde{B}' = F_B^*(B')_{red})$ is the reduced scheme underlying the pullback of B' by F_B^*). Now F_B induces a map $F_2 : \widetilde{B}' \to B'$. Also let F_1 denote the map $(Id, \phi_{p^l}) : V^2 \to V^2$, and let

$$F = (F_1 \times_{F_B} F_2) \ : \quad V^2 \times_{B^2} \widetilde{B}' \to V^2 \times_{B^2} B'$$

Finally let $\widetilde{S} = F^*(S)_{red}$ be the reduced pullback of S via F. Also if q is a power of p, let $\widetilde{q} = q/p^l$. We now claim that $\widetilde{\delta}'_p = 1$, $\widetilde{\delta} = \delta p^{ld}/\delta'_p$, and that if $\widetilde{b} \in \widetilde{B}'(K_q)$, $b = F(\widetilde{b})$, then $V_b(S, q)$ and $V_{\widetilde{b}}(\widetilde{S}, \widetilde{q})$ coincide as sets. This is easily seen by going to fibers: if $\widetilde{\beta}(\widetilde{b}) = (b_1, \widetilde{b}_2)$ then $\beta(b) = (b_1, b_2)$ with $b_2 = \phi_{p^l}(\widetilde{b}_2)$. If further $(a_1, \widetilde{a}_2)$ is a generic point of $\widetilde{S}_{\widetilde{b}}$, and $a_2 = \phi_{p^l}(\widetilde{a}_2)$, then (a_1, a_2) is a generic point of S_b. We have $\delta = (K_q(a_1, a_2) : K_q(a_1))$, $\widetilde{\delta} = (K_q(a_1, \widetilde{a}_2) : K_q(a_1))$, and the degree computations are elementary. The truth of Theorem 1B for S now follows from the same for $\widetilde{S}$.

11.6 Rough bound (Bezout methods)

We use Bezout's theorem (we need Fulton's "refined" version) to give a rough upper bound on the number of points on the intersection of a variety with Frobenius. It appears difficult to obtain the distribution of multiplicities of these points by this method; or to find the number of points hidden by a positive-dimensional component. We will thus use a different method, and the present section will not really be needed. On the other hand it might used in other asymptotic applications,e.g. to powers other than Frobenius.

Definition 11.20 *Let C be an irreducible subvariety of a variety P, $\mathcal{F}$ a set of hypersurfaces of P. We say C is* weakly cut out *by $\mathcal{F}$ if it is a component of the intersection scheme $\cap\mathcal{F}$.*

More generally, if the cycle $[\cap\mathcal{F}] = \sum n_i C_i$ with C_i the cycle of an irreducible variety, and $0 \leq m_i \leq n_i$, we say that the cycle $\sum m_i C_i$ is weakly cut out *by $\mathcal{F}$.*

Lemma 11.21 *Let S be a $N-k$-dimensional subvariety of projective space $\mathbb{P}^N$, weakly cut out by hypersurfaces of projective degree b. Then S is weakly cut out by k hypersurfaces of degree b.*

More generally, if $C = \sum m_i C_i$ is weakly cut out by l hypersurfaces of degree b, and $\dim(C_i) = N - k$ *for each i, then C is weakly cut out by k hypersurfaces of degree b.*

Proof Say S is weakly cut out by $P_j = 0$ $(j = 1, \ldots, l)$, with P_j a homogeneous polynomial of degree b. Let F be a field of definition for the various P_j. Let M be a $k \times l$ matrix, with coefficients generic over F. Let $(Q_1, \ldots, Q_k) = M(P_1, \ldots, P_l)$. Evidently each Q_i vanishes on S. Let $S[i]$ be the scheme cut out by $Q_1, \ldots, Q_i$. Then each component C of $S[i]$ has dimension $\geq N - i$. In particular, if $i < k$, then $C \not\subset S$, hence not all P_j vanish on C. By genericity of the $i + 1$'st row of M, over F and the previous rows, Q_{i+1} does not vanish on C. Thus, inductively, each component of $S[i]$ has dimension precisely $N - i$. Since $S \subset S[k]$, it follows that S is a component of $S[k]$.

The case of a cycle $C = \sum_i m_i[C_i]$, where S is an irreducible variety, follows by the same construction. It is merely necessary to note in the end that if $S = C_i$, since $S \subset S[k] \subset Z$ as schemes, where Z is the scheme cut out by the P_j, and S is a component of both $S[k]$ and Z, the geometric multiplicity of S in $S[k]$ is bounded by the multiplicity of S in Z. □

The proof of the next lemma will use positivity properties of intersection theory in multi-projective space. Let $\mathbb{P} = \mathbb{P}^{n_1} \times \mathbb{P}^{n_2} \times \mathbb{P}^{n_3}$ be a product of three projective spaces. Let T_i be divisors on $\mathbb{P}$, with multi-degrees (a_i, b_i, c_i) with respect to the standard divisors H_i, pullbacks of the hyperplane divisors on $\mathbb{P}^{n_i}$. Then

$$\{[T_1 \cap \ldots \cap T_m]\}_{\dim(\mathbb{P})-m} \leq T_1 \cdot \ldots \cdot T_m$$

Here on the left side we have the proper part of the cycle corresponding to the scheme-theoretic intersection, and on the right-hand side the intersection cycle in the sense of intersection theory. The inequality $\leq$ has the sense of 11.10. More explicitly, one considers the intersection of $T_1 \times \ldots \times T_m$ with the diagonal embedding of $\mathbb{P}$ in $\mathbb{P}^m$. This intersection has proper components $W_1, \ldots, W_l$, and other distinguished varieties Z_j. By [Fulton] 6.1, we have the canonical decomposition

$$T_1 \cdot \ldots \cdot T_m = \sum_i m_i[W_i] + \sum_j m_j \alpha_j$$

where α_j are certain cycles on the W_ν. We push forward to $\mathbb{P}$ to obtain the same equality for cycles on $\mathbb{P}$. Now the tangent bundle to multi-projective space is generated by global sections ([Fulton] 12.2.1(a,c)), hence ([Fulton] 12.2(a)) each α_j is represented by a non-negative cycle. So

$$T_1 \cdot \ldots \cdot T_m \geq \sum_i m_i[W_i]$$

On the other hand, by [Fulton] 7.1.10, m_i is precisely the geometric multiplicity of W_i in the intersection scheme $T_1 \cap \ldots \cap T_m$; so $\{[T_1 \cap \ldots \cap T_m]\}_{\dim(\mathbb{P})-m} = \sum m_i[W_i]$. The claimed Bezout inequality follows.

Lemma 11.22 *Let H be the hyperplane divisor on $\mathbb{P}^n$, $H_{12} = pr_1{}^*H + pr_2{}^*H$ on $\mathbb{P}^n \times \mathbb{P}^n$. Let U be a k-dimensional subvariety of $\mathbb{P}^n \times \mathbb{P}^n$, weakly cut out by hypersurfaces with H_{12}-degrees $\leq b$. Then*

$$\deg_{H_{12}}[U \cap \Phi_q]_l \leq C(q+1)^{k-l}$$

where C depends only on b,k,l,n (and is explicilty estimated in the proof.).

Proof The subvariety $\Phi_q \subset \mathbb{P}^n \times \mathbb{P}^n$ is cut out by $\binom{n}{2}$ hypersurfaces of bidegree $(q, 1)$ (namely, those defined by $X_i{}^q Y_j = X_j{}^q Y_i$). The proof of 11.21 shows that $U \cap \Phi_q$ is weakly cut out by $2n - k$ hypersurfaces of H_{12}-degree $\leq b$, and $k - l$ hypersurfaces of H_{12}-degree $O(q)$; and there are boundedly many choices of ω. Thus

$$0 \leq [U \cap \Phi_q]_l \leq \sum_\omega [S_\omega]_l$$

Fix one ω, and write the cycle $[S_\omega]$ as $\sum m_i C_i$; m_i is the geometric multiplicity of C_i in S_ω. By by [Fulton] 7.1.10, m_i is equal to the intersection multiplicity along C_i of the $2n - l$ hypersurfaces cutting out S_ω. By [Fulton] 12.3.1, $\sum m_i \deg_{H_{12}}(C_i)$ is bounded by the product of the degrees of these hypersurfaces. The lemma follows.

Lemma 11.23 *Let assumptions be as in 11.18 (1),(3). Assume further that for any $b \in B'$, the projection maps $S_b \to V_{b_2}$ has finite fibers. Then for all large enough q, and any $b \in B'$, $V_b(S, q)$ is finite. Moreover, for some constant ρ, $\#V_b(S, q) \leq \rho q^d$.*

Proof Note that the assumptions remain true if V is replaced by a proper subscheme U and S by any component of $S \cap U^2 \times_{B^2} B'$. Hence using Noetherian induction, we may assume the lemma holds for proper closed subschemes of V.

As in 11.19, we may assume the projection $S_b \to V_{b_2}$ is generically separable. Let $b \in B'(K_q)$; note that $V_b(S, q)$ is the set of points of a scheme over K_q, and let C be an irreducible component over K_q. Let a_1 be a generic point of C, $a_2 = \phi_q(a_1)$. Then $(a_1, a_2) \in S_b$ so the field extension $[K_q(a_1, a_2) : K_q(a_2)]$ is finite and separable. Since it is also purely inseparable, $K_q(a_1) = K_q(a_1^q)$, so $K_q(a_1)$ is a perfect, finitely generated field extension of K_q. It follows that $K_q(a_1) = K_q$, so that C is finite. Since C was an arbitrary component, $V_b(S, q)$ must be finite.

For the quantitative bound, we may (again using Noetherian induction and stratifying if needed) embed V_a in a projective space $\mathbb{P}_n$; the dimension and degree of the embedding are bounded irrespective of $a \in B$, as well as the bidegrees (f_1, f_2) of the resulting embedding of S_b in $\mathbb{P}_n \times \mathbb{P}_n$. Now $V_b(S, q)$ is the projection to V_{a_1} of the intersection of S_b with the graph Φ_q of the Frobenius map $\phi_q : \mathbb{P}_n \to \mathbb{P}_n$. By 11.22 the intersection has size $O(q^d)$, and the lemma follows. □

Remark 11.24

It's easy to see - and well known, cf. [Lang59], [Pink92] - that if $S_b \to V_{b_2}$ is étale, then S_b and Φ_q intersect transversally. (Look at the tangent spaces: $T\Phi_q$ is vertical at each point, while $TS_b \subset TV_a \times TV_b$ is a linear isomorphism between the tangent spaces to the factors.) (To see that they meet properly, if $S \subset V \times V^\sigma$ are varieties over k, and $(a, b) \in (S \cap \Phi_q)(k^{alg})$ with $a \in k(b)^{sep}$, then $k(a) \subset k(a^{1/q})^{sep}$, so $k(a)^{sep}$ is perfect, and it follows that $a \in k^{sep}$. See also 5.2.)

11.7 Finite covers

Lemma 11.25 *(Birational invariance). Let B,V,B',S satisfy the hypotheses of Theorem 1B. Suppose $\widetilde{V}$ is an open subscheme of V, so that for a the generic point of B, $\widetilde{V}_a$ is an* *open subvariety of V_a. Let $\widetilde{S} = S \cap \widetilde{V}^2 \times_{B^2} B'$. Let $\widetilde{B}'$ be an open subscheme of B' such that the hypotheses of 11.18 hold. Then Theorem 1B is true for B,V,B',S iff it holds for B,$\widetilde{V}$,$\widetilde{B}'$,$\widetilde{S}$.*

Proof The invariants δ,δ'_p are the same for the two families. The sets $V_b(S,q)$ and $\widetilde{V}_b(\widetilde{S},q)$ are also the same, except for points in $F_b(S^*,q)$, where F is a closed subscheme of V complementary to $\widetilde{V}$ at the generic fiber, and S^* is the restriction of S to F. By 11.23, the size of $F_b(S^*,q)$ is of the order of the error term in the expression for $V_b(S,q)$ in 1B. □

The first item of the next lemma will be used to reduce to smooth varieties. The second will not be used, but originally gave some indication of the correctness of the form of Theorem 1B, since it shows that the desired lower bound in this theorem follows from the upper bound. (Once once one knows the theorem for the easy case of single difference equations),

11.26

$$\#(V_b(S,q)) \leq \delta^* q^d + e \qquad \textit{with} \;\; e \leq \rho q^{d-\frac{1}{2}}$$

We will write such statements as

$$\#(V_b(S,q)) \leq \delta^* q^d + O(q^{d-\frac{1}{2}})$$

The point is that the implicit coefficient depends only on the system (B', V', S, β) and not on the choice of b or of q.

Lemma 11.27 *Let B,V,B',S satisfy the hypotheses of 1B. Let $h_B : \widetilde{B} \to B$ and $h_V : \widetilde{V} \to V$ be compatible quasi-finite maps (base extension for B, and finite cover of V). Let $\widetilde{B}' = B' \times_{B^2} \widetilde{B}^2$. We get an induced map $h : (\widetilde{V})^2 \times_{\widetilde{B}^2} \widetilde{B}' \to V^2 \times_B^2 B'$. Let $\widetilde{S}(j)$ be the various irreducible components of $h^{-1}S$, that are nonempty in the generic fiber.*

(a) Suppose the conclusion of 1B holds for each $\widetilde{B}$,$\widetilde{V}$,$\widetilde{B}'$,$\widetilde{S}(j)$. Then it holds also for the original system.

(b) Conversely, suppose the conclusion of 1B hold for (B, V, B', S), and the conclusion of 11.26 holds for each $\widetilde{B}$,$\widetilde{V}$,$\widetilde{B}'$,$\widetilde{S}(j)$. Then the conclusion of 1B holds for each $\widetilde{B}$,$\widetilde{V}$,$\widetilde{B}'$,$\widetilde{S}(j)$.

Proof We can first apply the base change from B to $\widetilde{B}$ without changing V (i.e. replacing V by $V \times_B \widetilde{B}$); this clearly makes no difference to 1B. Thus we may assume $\widetilde{B} = B$, and $\widetilde{B}' = B'$. Let a be the generic point of B. Let $h_a : \widetilde{V}_a \to V_a$ be the map induced by h on the fibers. Outside of a proper subvariety of V_a, all fibers of h_a have the same size s. This proper subvariety may be removed using 11.25. Similarly we may assume the various $\widetilde{S}(j)$ are disjoint. Let b be a generic point of B'. Let

$$H = h_{b_1} \times h_{b_2} : \; \widetilde{V}_{b_1} \times \widetilde{V}_{b_2} \to V_{b_1} \times V_{b_2}$$

We have $H^{-1}S_b = \cup_j S(j)_b$. Thus

$$H^{-1}V_b(S,q) = \cup_j V_b(\widetilde{S}(j), q)$$

By the constant size of the fibers of h_a,

$$s\#V_b(S,q) = \Sigma_j \#V_b(\widetilde{S}(j), q)$$

A simple exercise in degrees of field extensions shows:

$$\delta^* s = \sum_j \delta_j^*$$

Where δ_j^* refers to $\widetilde{S}(j)$, cf. 1B. (Because of 11.19, only the case $\delta_p' = 1$ need be considered in the applications).

Now (a) if 1B holds $\widetilde{S}(j)$, we have

$$\#V_b(\widetilde{S}(j), q) = \delta_j^* q^d + O(q^{d-\frac{1}{2}})$$

Summing over j , using the previous two displayed equalities, we obtain

$$s\#V_b(S, q) = \delta^* s q^d + O(q^{d-\frac{1}{2}})$$

and division by s yields (a). (b) follows similarly, using the principle that if $\sum_i a_i = \sum_i a_i'$ and each $a_i \leq a_i'$, then each $a_i = a_i'$ (here all up to $O(q^{d-\frac{1}{2}})$.) □

11.8 Reduction to smooth varieties

Lemma 11.28 *In Theorem 1B, we may assume that V is an open subscheme of $\bar{V}$, $S = \bar{S} \cap V^2$, with $\bar{V}$ smooth and projective over B.*

Proof In characteristic zero, we can use Hironaka's resolution of singularities; the generic fiber V_b of V is birational to an open subvariety of a smooth, proper variety, and by 11.25 the theorem is invariant under birational changes of this type. Similarly, in general, we can use 11.27 and the following theorem of de Jong ([deJong]):

For any variety V over a field k there exists a finite extension $\widetilde{k}$ of k, and a smooth projective variety $\widetilde{V}$ over $\widetilde{k}$, and a finite, dominant morphism from an open subvariety of $\widetilde{V}$ to V. □

Remark 11.29

For theorem 1B' the reduction to the smooth case is not needed on this (geometrical and numerical) level; a similar but easier reduction shows directly (using [deJong]) that the instances of the axiom scheme ACFA referring to open subvarieties of smooth complete varieties imply the rest of the axioms.

12 The virtual intersection number

This section is devoted to a proof of theorem 12.2. It is formally similar to Theorem 1B; but it refers to the virtual intersection number, in the sense of intersection theory, and not to the actual intersection. By virtue of 11.28,11.19 we can now assume that the varieties V_a are smooth and complete, and S_b/V_{b_2} separable. But as a result of the completion process, the intersection may have infinite components. The effect of these components on the intersection number will have to be estimated in the next section.

Lemma-Notation 12.1 *Let μ, ν be bounds for the absolute degree and the sum of Betti numbers of V_b, so that:*

- *For $b \in B$, there exists a very ample divisor H_b on V_b, such that $|S_b| = |S_b|_{H_b} \leq \mu$*
- *For $b \in B$, letting X be the variety $\bar{V}_b$ base-extended to an algebraically closed field, we have $\Sigma_{i \leq 2\dim(X)} \dim H^i(X, \mathbb{Q}_l) \leq \nu$.*

Proof To show the existence of either μ or ν, it suffices to show that the numbers in question are bounded on a Zariski open subset of V; then one may use Noetherian induction. Let b be a generic point of V, and pick a very ample H_b on V_b. The construction of 11.15 can be described by a first-order formula concerning b. This formula must remain valid for a Zariski neighborhood of b. Thus $|S_{b'}|$ is uniformly bounded on this neighborhood. The same argument, using the constructibility of the higher derived images of constant sheaves, shows that $\Sigma_{i \leq 2\dim(X)} \dim_{Z/lZ} H^i(X, Z/lZ)$ is bounded on a Zariski neighborhood; this number bounds also the sum of the Betti numbers.

Theorem 12.2 *Let notation be as in 1B and 11.28. Let $b \in B'(K_q)$, $\beta(b) = (a, \phi_q(a))$, $X = \bar{V}_a$, $X' = \bar{V}_{\phi_q(a)}$. Let $\Phi_b \subset (X \times X')$ be the graph of Frobenius. Then*

$$S_b \cdot_{X \times X'} \Phi_b = \delta q^d + e \qquad \textit{with} \ \ |e| \leq \binom{2d}{d}^2 \mu\nu q^{d-\frac{1}{2}}$$

Example 12.3

In the case of projective space, i.e. $\bar{V}_{b_1} = \mathbb{P}_d$, we can immediately prove Theorem 12.2 cycle-theoretically (without transcendental cohomology.)

Let $\bar{S}_b \subset {\mathbb{P}_d}^2$ be the closure of S_b. Consider the intersection of $\bar{S}_b$ with Φ_q, the graph of Frobenius on $\mathbb{P}^d$. In the intersection theory sense, it can be computed in terms of the bidegrees:

$$[\Phi_q] = \sum_i q^i h_1^i h_2^{d-i}$$

$$[S] = \sum_i a_i {h_1}^i {h_2}^{d-i}$$

with $a_0 = \delta$. So

$$[\Phi_q \cdot S] = \sum_i a_i q^{d-i} = \delta q^d + O(q^{d-1})$$

□

For curves, Weil's "positivity" proof of the Riemann Hypothesis ([Weil48]) works here too. (The notation follows §11.3.)

Example 12.4 *Let C be a smooth, complete curve of genus g over a field K of characteristic $p > 0$. Let $q = p^m$. Let ϕ_q be the q-Frobenius, and let $C' = C^{\phi_q}$. Let $S \leq C \times C'$ be an irreducible subvariety. Then*

$$S \cdot \Phi_q = q \deg_{cor}(S) + deg_{cor}(S^t) + e$$

with

$$|e| \leq ((2g)(2\deg_{cor}(S)\deg_{cor}(S^t) - |S \cdot S^t|))^{1/2} q^{1/2} (\leq O(1) q^{1/2})$$

Proof Let A_0 be the group of divisors on $C \times C'$, up to rational equivalence, $A = \mathbb{Q} \otimes A_0$. Define a symmetric bilinear form:

$$\beta(X, Y) = \deg_{cor}(X) \deg_{cor}(Y^t) + \deg_{cor}(Y) \deg_{cor}(X^t) - X \cdot Y$$

If X,Y are divisors on $C \times C^{'}$, let $X^t Y$ denote the intersection-theoretic composition $X^t \circ Y$. It is a correspondence on $C \times C$; and we have:

$$(\Delta.(X^t Y)) = (X, Y), \quad deg(X^t Y) = deg(Y) deg(X^t)$$

Thus $\beta(X, Y) = \beta_C(X^t Y, \Delta)$ where β_C is defined like β, but on $C \times C$.

Let $Z = \Phi_q$. We have $Z^t Z = Z Z^t = q\Delta$ (where Δ denotes the diagonal of C, respectively C'.) By associativity of composition, for any X and Y,

$$(Z^t X)^t (Z^t Y) = X^t (Z Z^t) Y = q X^t Y$$

Hence $\beta(X, Y) = \beta_C(X^t Y, \Delta) = (1/q)\beta_C((Z^t X)^t, (Z^t Y))$. Now according to Weil, $\beta_C(W^t, W) \geq 0$; hence $\beta(X, X) \geq 0$. This non-negativity extends to $A \otimes \mathbb{R}$.

By Cauchy- Schwartz, we have $\beta(X, S) \leq \beta(X, X)^{1/2} \beta(S, S)^{1/2}$. Now

$$\beta(X, X) = \beta_C(q\Delta, \Delta) = q\beta_C(\Delta, \Delta) = q(2 - \Delta \cdot \Delta) = q(2 + (g - 2)) = 2gq$$

$$(X \cdot S) = q deg(S) + deg(S^t) - \beta(S, \Phi_q)$$

So $e^2 \leq \beta(S, S)(2gq)$, and the estimate follows. □

But in general, we will have to decompose $[S]$ and Φ_q cohomologically and not cycle-theoretically.

It will be convenient to renormalize the norm, and write $||S|| = \binom{2d}{d}^2 |S|$. Then (by 11.17) $||S \circ T|| \leq \binom{2d}{d}^4 |S| \, |T| = ||S|| \, ||T||$

When dealing with the étale cohomology groups, we will always work over an algebraically closed field k, fix a prime $l \neq char(k)$, and fix an isomorphism of $\mathbb{Z}_l[1]$ with $\mathbb{Z}_l$. See the section of [Deligne 74] on "orientations", and [Kleiman68]. (For the groundwork, see [Deligne 77] and the references there, or [Freitag-Kiehl] or [Milne1980]).

Notation 12.5 *If R is a correspondence on $X \times X'$, we let $\eta_i(R) : H^i(X, \mathbb{Q}_l) \to H^i(X', \mathbb{Q}_l)$ be the endomorphism of $H^i(X, \mathbb{Q}_l)$ induced by R ; cf [Kleiman68], 1.3. When $X = X'$, we let $\tau_i(R)$ the trace of this endomorphism.*

When R is a correspondence on X as in 12.5, the Lefschetz trace formula ([Kleiman68], 1.3.6) relates the intersection number of R with the diagonal to the endomorphisms $\eta_i(R)$ as follows:

12.6 $R \cdot \Delta_X = \Sigma_{i=1,\ldots,2n} (-1)^i \tau_i(R)$

For the top cohomology, we have

12.7 *Let $n = \dim(X)$. Then $\dim H^{2n}(X) = 1$, and $\tau_{2n}(R) = \deg_{cor}(R)$ is the degree of S as a correspondence (11.16).*

(The first statement is in [Kleiman68] 1.2; hence τ_{2n} acts as multiplication by a scalar, and this scalar can be seen to be $\deg_{cor}(R)$ by computing the effect of $\eta_{2n}(R)$ on the image of a point under the cycle map.)

We will make weak use of Deligne's theorem ([Deligne 74]).

12.8 *Let X be a smooth projective variety over k, the algebraic closure of a finite field. Suppose X descends to $\mathbb{F}_q$. Let Φ be the graph of the Frobenius correspondence $x \mapsto x^q$ on $X \times X$. Then all the eigenvalues of $\eta_i(\Phi^t)$ are algebraic numbers, of complex absolute value $q^{i/2}$.*

We wish to study eigenvalues of compositions of two correspondences S,T. This can probably be done using the remark 12.14 below. However we will take another route, starting with a slight generalization of [Lang59], V.3 lemma 2.

Lemma 12.9 *Let $b, t > 0$ and ω_j, c_j be complex numbers ($j = 1, \ldots, s$). Suppose*

$$\limsup t^{-n} \left|\Sigma_j c_j \omega_j^n\right| \leq b$$

Then one may partition $\{1, \ldots, s\}$ into disjoint sets J_1,J_2 such that

$$\Sigma_{j \in J_2} c_j \omega_j^n = 0$$

for all n, and

$$|\omega_j| \leq t$$

for each $j \in J_1$.

Proof Let M be the maximal norm of an ω_j. Dividing the ω_j and t by M, we may assume $M = 1$.

Consider first the ω_j lying on T, the group of complex numbers of norm one. Say they are $\omega_1, \ldots, \omega_{s'}$; note that summing only over $j \leq s'$ does not effect the $\limsup$ in the hypothesis. Let S be the closed subgroup of $T^{s'}$ generated by the point $\alpha = (\omega_1, \ldots, \omega_{s'})$. Now S is a compact group. In every neighborhood of S, there are infinitely many powers α^m; for these we have $(c_1, \ldots, c_{s'}) \cdot \alpha^m$ arbitrarily small; hence there exists a point γ in the closure of this neighborhood, with $(c_1, \ldots, c_{s'}) \cdot \gamma = 0$. So such points are dense in S, and thus $(c_1, \ldots, c_{s'}) \cdot \gamma = 0$ for all $\gamma \in S$. We may thus put all the ω_i, $i \leq s'$, into J_2. Now removing them from the sum has no effect, so the hypothesis applies to $\omega_{s'+1}, \ldots, \omega_s$, and we may proceed by induction. □

Lemma 12.10 *Let R and T be commuting correspondences on a smooth projective variety X of dimension d over an algebraically closed field K, $T \circ R = R \circ T$. Fix an embedding of $\mathbb{Q}_l$ into $\mathbb{C}$. Let c_{ij}, e_{ij} be the corresponding eigenvalues of the endomorphisms $\eta_i(R)$, $\eta_i(T)$, considered as complex numbers. Let J_1 be the set of (i,j) such that e_{ij} has complex norm at most $||T||$, and J_2 the rest. Then for any m, k*

$\Sigma_{(i,j) \in J_2} (-1)^i c_{ij}^k e_{ij}^m = 0$

Proof The following claim is immediate from 11.17 (3) and (4):

Claim For all m, $(R \circ T^{\circ m} \cdot \Delta_X) \leq |R| \, ||T||^m$

Let $\gamma = \sum_{(i,j)\in J_1} |c_{ij}| + |R|$. By definition of J_1,

$$\left|\Sigma_{(i,j)\in J_1}(-1)^i c_{ij} e_{ij}^m\right| \leq (\gamma - |R|)||T||^m$$

By 12.6,

$$(R \circ T^{\circ m}) \cdot \Delta_X = \Sigma_i (-1)^i \tau_i(RT^m) = \Sigma_{ij}(-1)^i c_{ij} e_{ij}^m$$

Thus

$$\left|\Sigma_{(i,j)\in J_2}(-1)^i c_{ij} e_{ij}^m\right| \leq \gamma ||T||^m$$

By 12.9, one can partition J_2 into disjoint sets J_3,J_4 such that $\Sigma_{(i,j)\in J_3}(-1)^i c_{ij} e_{ij}^m = 0$ for all positive integers m, and $|e_{ij}| \leq ||T||$ for each $(i,j) \in J_4$. But then $J_4 \subseteq J_1$ by definition of J_1, so $J_4 = \emptyset$. Thus $J_2 = J_3$, and we have shown what we wanted for $k = 1$. Applying this to R^k in place of R yields the lemma. □

The following Proposition does not mention Frobenius, and could be stated over arbitrary fields, but the proof we give uses [Deligne 74].

Proposition 12.11 *Let R be a correspondence on a smooth projective variety X over k, the algebraic closure of a finite field. Then every eigenvalue of $\eta_i(R)$ is an algebraic number, of complex absolute value at most $||R||$.*

It suffices to show that in every complex embedding, every eigenvalue has absolute value at most $||r||$. Fix therefore a complex embedding, and view the eigenvalues of $\eta_i(R)$ as complex numbers e_{ij}. Let $T = \Phi_q^t$ be a Frobenius correspondence on X, with q large enough so that $T \circ R = R \circ T$ as correspondences. Let J_2 index the set of e_{ij} of absolute value $> |R|$. Let c_{ij} be the eigenvalues of T, with corresponding indices. By 12.10, for any m and k,

$$\Sigma_{(i,j)\in J_2}(-1)^i c_{ij}^m e_{ij}^k = 0$$

Suppose for contradiction J_2 is nonempty. Let $i_{\max}$ be the highest index represented in J_2, and let $J_3 = \{j : (i_{\max}, j) \in J_2\}$. Let $d_{ij} = (-1)^i c_{ij} q^{-i_{\max}/2}$, and let $d_j = d_{i_{\max},j}$, $e_j = e_{i_{\max},j}$. Then $|d_j| = 1$, $|d_{ij}| < 1$ for $i < i_{\max}$, and $\Sigma_{(i,j)\in J_2} d_{ij}^m e_{ij}^k = 0$ for any k,m. For fixed m,k we have $\Sigma_{(i,j)\in J_2} d_{ij}^{m+m'} e_{ij}^k = 0$ Let m' approach infinity, becoming highly divisible; then $d_{ij}^{m'}$ approaches 0 for $i < i_{\max}$, while $d_j^{m'}$ approaches 1. We obtain

$$\Sigma_{j\in J_3} d_j^m e_j^k = 0$$

for any k,m. Dually, let $J_4 = \{j \in J_3 : |e_j| = L\}$, where L is the highest value attained by an e_j. Then as above we get

$$\Sigma_{j\in J_4} d_j^m (e_j/L)^k = 0$$

for any k,m. But now letting both m and k approach infinity, each d_j^m and each $(e_j/L)^k$ approach 1, and we obtain $|J_4| = 0$, a contradiction. Thus $J_2 = \emptyset$ and we are done.

Corollary 12.12 *Let X be a smooth complete variety over the algebraic closure k of a finite field. Let X' be the image of X under the q-Frobenius automorphism of K, and let $T = \Phi_q^t \subset X' \times X$ be the transposed Frobenius correspondence. Let $R \subset X \times X'$ be a correspondence. Then every eigenvalue of $\eta_i(T \circ R)$ is an algebraic number, of complex absolute value at most $q^{i/2}\, ||R||$.*

Proof For some m, X and S are defined over $\mathbb{F}_{q^m}$. Taking m'th powers, it suffices to prove that every eigenvalue of $\eta_i((T \circ R)^m)$ is an algebraic number, of absolute value at most $q^{im/2}||R||^m$. Let X_i be the image of X under ϕ_q^i, and $R_i \subset X_i \times X_{i+1}$ the image of R. Then

$(T \circ R)^m = \Phi_{q^m}^t \circ R_{m-1} \circ \ldots \circ R$

Let $R^* = R^{m-1} \circ \ldots \circ R$. Then (11.17) $||R^*|| \leq \Pi_{i=0,\ldots,m-1} ||R_i|| = ||R||^m$. R^* is a correspondence on X. By 12.11, every eigenvalue of $\eta_i(R^*)$ is algebraic of absolute value at most $|R|^m$. Now note that $\Phi_{q^m}^t$ is a correspondence on X, commuting with R^*, and (12.8) with algebraic eigenvalues of absolute value $q^{im/2}$. Thus $\eta_i((T \circ R)^m) = \eta_i(\Phi_{q^m}^t \circ R^*)$ has algebraic eigenvalues of absolute value $(||R||q^{i/2})^m$. The conclusion follows.

Lemma 12.13 *In Theorem 12.2 we may assume K_q is the algebraic closure of the prime field; in other words it suffices to prove the result for b chosen from a finite field.*

Proof We may assume V (and $V \times_B V$) are flat over B. Let T be an irreducible component of $B' \cap \beta^{-1}\Phi_B$, where $\Phi_B \subset B^2$ is the graph of Frobenius on B. By [Fulton] (10.2), the number

$$\bar{S}_b \cdot_{\bar{V}_{b_1} \times \bar{V}_{b_2}} \Phi_b$$

is constant for $(b_1, b_2) \in T$. Thus one can replace X, X' by V_{b_1}, V_{b_2}, where b is a point of T rational over a finite field. □

proof of Theorem 12.2

By 12.13, we may assume K is the algebraic closure of a finite field. By 12.6,

$S_b \cdot \Phi_b = ((\Phi_b)^t \circ S_b) \cdot \Delta_X = \Sigma_{i \leq 2d} \tau_i((\Phi_b)^t \circ S_b)$

By 12.12,

$$|\Sigma_{i<2d} \tau_i((\Phi_b)^t \circ S_b)| \leq \binom{2d}{d}^2 \mu\nu q^{d-\frac{1}{2}}$$

By 12.7, $\tau_{2d}((\Phi_b)^t \circ S_b) = \deg_{cor}((\Phi_b)^t \circ S_b) = \deg_{cor}(S_b)q^d$. □

We conclude the section with the rationality lemma mentioned but not used above.

Lemma 12.14 *Let X be a smooth projective variety over a field K, and let S, T be correspondences on X. Write a power series in two variables:*

$$F(s,t) = \Sigma_{m,n \geq 0} S^m \cdot T^n s^m t^n$$

Then F is a rational function, with denominator of the form $f(s)g(t)$.

Proof Using Grothendieck's cohomological representation, it suffices to prove more generally that if V is a complex vector space, and S, U two linear transformations of V, then $\Sigma_{m,n \geq 0} tr(S^m U^n) s^m t^n$ is a rational function of two variables.

We will actually show that

$$\Sigma_{m,n} S^m U^n s^m t^n \in End(V)[[s,t]]$$

lies in $End(V) \otimes_{\mathbb{C}} C(s,t)$, so that all the matrix coefficients are rational functions. (With denominators as specified.) Now

$$\Sigma_{m,n\geq 0} S^m U^n s^m t^n = (\Sigma_m S^m s^m)(\Sigma_n U^n t^n)$$

So it suffices to consider $\Sigma_m S^m s^m$. We can sum the geometric series:

$$\Sigma_{m\geq 0} S^m s^m = (1 - Ss)^{-1} = (det(1 - Ss)^{-1}) S'$$

For a certain matrix S' with coefficients polynomial in s. □

13 Equivalence of the infinite components: asymptotic estimate

While Theorem 1.1promises actual points of intersection, Theorem 12.2 refers rather to an intersection number, i.e. to the intersection with a better placed S' rationally equivalent to S. We will now bridge the gap by giving an asymptotic estimate of the difference between the number of isolated points of these intersections. Here we do so numerically and roughly (up to lower order of magnitude); this requires only tying together some previous threads. This suffices for the proof of Theorem 1.1. However our methods are essentially precise and motivic, and later we will describe the language needed to bring this out.

We will actually bound at once both the equivalence of the infinite components, and the contribution of isolated points of high multiplicity.

We work with a single fiber of the data 11.18: for some $b \in B'(K_q)$ with $\phi_q(b_1) = b_2$, we let $\widetilde{X}_1 = V_{b_1}$, $\widetilde{X}_2 = V_{b_2}$, $\widetilde{S} = S_b$. Thus we have a d-dimensional variety $\widetilde{X}_1$ over K_q, $\widetilde{X}_2 = \widetilde{X}_1^{\phi_q}$, and an irreducible subvariety $\widetilde{S} \subset \widetilde{X}_1 \times \widetilde{X}_1$. We will bound a certain quantity associated with the intersection of $\widetilde{S}$ with the graph of Frobenius. We will write $O(q^l)$ for a quantity bounded by Cq^l, where C is a constant independent of q, and dependent on $\widetilde{X}_1$, $\widetilde{X}_2$, $\widetilde{S}$ in a way that remains bounded when the varieties are obtained as above from the data 11.18, with b varying.

We assume (cf. 11.19,11.28) that the projection from $\widetilde{S}$ to $\widetilde{X}_2$ is étale, and that $\widetilde{X}_1$ is an open subvariety of a smooth, complete variety X_1. (And thus $\widetilde{X}_2$ of $X_2 = X_1^{\phi_q}$). Let $\Phi = \Phi_q \subset Y = X_1 \times X_2$ denote the graph of $\phi_q : X_1 \to X_2$, and let S be the Zariski closure of $\widetilde{S}$ in $X_1 \times X_2$.

Proposition 13.1 *Let $e = (\Phi_q \cdot S) - |\{\Phi_q \cap \widetilde{S}\}|$. Then $|e| \leq O(q^{d-1})$*

Observe that the points of $\Phi_q \cap \widetilde{S}$ are simple, by by 11.24; so it makes no difference whether $|\{\Phi_q \cap \widetilde{S}\}|$ is counted with multiplicities.

Almost all Frobenius specializations Recall that a Frobenius field K_q is an algebraically closed difference field of characteristic $p > 0$ made into a difference field via $x \mapsto x^q$ ($q = p^m, m \geq 1$). $|Y|$ denotes the size of a 0-dimensional scheme, i.e. the number of points with multiplicities.

We will consider difference schemes Y of finite type over a difference field k. Y arises by base extension from a scheme Y_D over a finitely generated difference domain $D \subset K$. We will say: "for almost all (q, h), $\cdots$ " to mean: for some such D and Y_D, for all sufficiently large prime powers q, and all $h : M_q(D) \to K_q$ ....

When K is infinite, by enlarging D, one sees that a simpler formulation is equivalent:

"for some D, Y_D, for all prime powers q and all $h : h : M_q(D) \to K_q \ldots$."

When concerned with only almost all (q, h), we will not mind increasing D within K, and so can assume that some such Y_D has been fixed, and write $Y_{q,h}$ for $M_q(Y_D) \otimes_h K_q$. Different choices of Y_D will give the same $Y_{q,h}$ for almost all (q, h).

Let $D \subset k$ be a finitely generated difference domain, and let $D[t]'$ be a finitely generated sub-D-difference algebra of $k[\breve{t}]_\sigma$. By enlarging $D[t]'$, and D, we may obtain the form $D[t]' = D[t^{\sigma - n}, f^{-1}]_\sigma$, where $f \in D[t]_\sigma$ and $f(0)$ is invertible in D. Consider morphisms $h : M_q(D) \to K_q$ into Frobenius fields K_q with $q \geq n$. For any such h, let ${h_t}^* : D[t]' \to K_q(t)$ be the unique difference ring morphism extending h and with ${h_t}^*(t^{\sigma - n}) = t^{q-n}$. Sometimes we will consider it as a map $D[t]' \to K_q[t, f^{-1}(t)]$ or $D[t]' \to K_q[\breve{t}] := K_q[t, g^{-1} : g(0) \neq 0]$; in this case we will write $\breve{h}_t$.

When Y is a difference scheme of finite type over $k[\breve{t}]_\sigma$, we fix $D[t]'$ as above; given q and $h : M_q(D) \to K_q$, we define $Y_{q,h} = Y_{q,{h_t}^*}$.

Let k be a perfect, inversive difference field, V an absolutely irreducible, projective (or complete) algebraic variety over k, $\dim(V) = d$.

Let $\mathcal{C}(V)$ be the set of subvarieties S of $V \times V^\sigma$, such that any component C of S (over k^{alg}) has dimension d, and the projections $\mathrm{pr}_0^S : S \to V, \mathrm{pr}_1^S : S \to V^\sigma$ are dominant. Let $\mathbb{Z}\mathcal{C}(V)$ denote the free Abelian group generated by $\mathcal{C}(V)$. For any subscheme W of $V \times V^\sigma$, let $[W]$ be the formal sum of the components of W of dimension d, weighted by their geometric multiplicities. (cf. [Fulton].)

Write $S \in \mathcal{C}(V)_{\mathrm{sep}}$ if in addition, $k(S)$ is a separable finite extension of $k(V^\sigma)$. In this case, the projection pr_1^S is étale above a nonempty Zariski open subset of V^σ. Let $V_{et}(S)$ be the largest smooth open subvariety U of V such that pr_1^S is étale above U^σ.

In general, for $S \in \mathcal{C}(V)$, let $\mathrm{pr}_1[1](S) = \{a \in V : \dim {\mathrm{pr}_1^S}^{-1}(a) > 0\}$. Let $V_{fin}(S)) = (V \setminus \mathrm{pr}_0[1](S))^{\sigma^{-1}}$, and $\widetilde{S} = S \cap (V_{fin}(S) \times V_{fin}(S)^\sigma)$. Thus $V_{fin}(S)$ is the largest open subvariety of V such that, with this definition, $\mathrm{pr}_1 : \widetilde{S} \to \widetilde{V}^\sigma$ is quasi-finite.

Let $X(S) = (\widetilde{S} \star \Sigma)$. So $X(S)$ has total dimension d.

Let $\mathcal{C}_f(V) = \{S \in \mathcal{C}(V) : \mathrm{pr}_0[1](S) \cap ([\sigma]_k S \star \Sigma) = \emptyset\} = \{S \in \mathcal{C}(V) : X(S) = S \star \Sigma\}$. (cf. 6.6 .) $\mathbb{Z}\mathcal{C}_f(V)$ is the free Abelian group generated by $\mathcal{C}_f(V)$.

The $X(S)$ with $S \in \mathcal{C}_f(V)$ are directly presented, and of total dimension d; and moreover of transformal multiplicity 0.

Asymptotic equivalence Given $S \in \mathcal{C}(V)$, defined over a finitely generated difference domain D, let $X = X(S)$, let $\eta_0(S; q, h)$ be the number of points in $X_{q,h}(K_q)$, and let $\eta(S; q, h) = \delta'_p \eta(S; q, h)$ (cf. 11.18).

We say $S, S' \in \mathcal{C}(V)$ are *asymptotically equinumerous* if if there exists a finitely generated difference domain $D \subset k$ (containing any given finite set), such that for some b, for almost all q, for all $h : M_q(D) \to K_q$, $|\eta(S; q, h) - \eta(S'; q, h)| \leq bq^{d-1}$.

We extend the terminology to $\mathbb{Z}\mathcal{C}(V)$ by additivity.

(We could instead count the points of $X(S)_{q,h}$ with their geometric multiplicities, or with appropriate intersection multiplicities; asymptotically this makes no difference, and so we chose the approach involving least irrelevant complications.)

Let $\mathbb{Z}\mathcal{C}(V)_{\mathrm{rat}}$ be the group of cycles in $\mathbb{Z}\mathcal{C}(V)$ that are rationally equivalent to 0, by a rational equivalence involving cycles in $\mathbb{Z}\mathcal{C}(V)$; i.e. $\mathbb{Z}\mathcal{C}(V)_{\mathrm{rat}}$ is generated by cycles $\underline{S}_0 - \underline{S}_\infty$ where $\underline{S} \subset (V \times V^\sigma) \times \mathbb{P}^1$ is a $d+1$-dimensional variety, with $\underline{S}_0, \underline{S}_\infty \in \mathcal{C}(V)$. Similarly define $\mathbb{Z}\mathcal{C}(V)_{\substack{\mathrm{sep}\\ \mathrm{rat}}}$ within $\mathbb{Z}\mathcal{C}(V)_{\mathrm{sep}}$. Recall 11.7 :

Lemma 13.2 *For any $S \in \mathcal{C}(V)$ there exists $S' \in \mathbb{Z}\mathcal{C}_f(V)$ with $S - S' \in \mathbb{Z}\mathcal{C}(V)_{\mathrm{rat}}$.*

Lemma 13.3 *Let $S \in CV$, and let Z be a proper Zariski closed subset of V. Then $|X_{q,h} \cap Z| = O(q^{d-1})$.*

Proof By 6.6 , 10.24 . □

Lemma 13.4 *Let $\underline{S}$ be an irreducible $d+1$-dimensional variety of $(V \times V^\sigma) \times \mathbb{P}^1$. Let t denote a generic point of $\mathbb{P}^1$ over k. Assume $S_t = \underline{S}_t$ and $S_0 = \underline{S}_0$ are in $\mathcal{C}(V)_{\mathrm{sep}}$. Then $S_t \sim_a S_0$.*

(Here we identify S_0, a variety over k, with the same variety viewed by base change as a variety over $k(t)$. The separability assumption could be dispensed with, as in 13.5 , 11.19 .)

Proof Let $V_{\breve{A}_k} = V \times_{\mathrm{Spec}^\sigma k} \breve{A}_k$. $V_t = V \times_k k(t)_\sigma$, $V_0 = V$ denote the fibers of this map over the generic point of $\breve{A}_k$ and the point $t = 0$ respectively.

Let X_t be the closure of $X(S_t)$ in V_t (i.e. the smallest $k(t)_\sigma$- definable closed difference subscheme containing $X(S_t)$.) Let X be the closure in $V_{\breve{A}_k}$ of X_t; it is a k- closed difference subscheme of $\underline{V} = [\sigma]_k V \times_k \breve{A}_k$, flat over $\breve{A}_k$, whose fiber over $k(t)_\sigma$ is X_t.

Let $X_{\to 0}$ be the pathwise specialization of X_t, as in 9.7 .

Let $W = V_0 \setminus V_{et}(S_0)$; and let ${}^*W$ be as in Proposition 9.10 (a formula in the language of transformal valued fields over $k(t)_\sigma$.) By definition of $X(S_t)$, and 6.6 , condition (1) of 9.10 holds; hence it holds also for the closure X_t. (Cf. 4.6). Since $S_0 \in \mathcal{C}(V)$, condition (2) of 9.10 holds too. Thus (4) holds, i.e. ${}^*W$ has inertial dimension $< \dim(V)$.

Consider a (large) prime power q. Fix a nontrivial valuation of $L = L_q = K_q(t)^a$ over K_q (all are isomorphic); let R_q denote the valuation ring. We will show that $|X(S_t)_{q,h}(L)| = |X(S_0)(L)| \bmod O(q^{d-1})$. Since q is large enough, $X(S_0)_{q,h}$ is a finite scheme; so $X(S_0)_{q,h}(L) = X(S_0)_{q,h}(K_q)$.

Since V_t is projective, $V_t(R_q) = V_t(L)$, and the residue homomorphism induces a map $r_V : V_t(L) \to V_0(L)$; restricting to $r_0 : (X_t)_{q,h}(L) \to (X_0)_{q,h}(K_q)$. Let $Y = (X_t)_{q,h} \setminus^* W$. Let $r_1 = r_0|Y$. By definition of ${}^*W$, $r_1 : Y \to V_{et}(S_0)$.

Injectivity of r_1 follows from 8.3. An alternative argument: by 11.24 , $X(S_0)_{q,h}$ has only simple points on $V_{et}(S_0)$; by 7.22 , r_1 must be injective.

Surjectivity. We can use a Hensel lemma approach (analogous to 8.4 ; here only algebraic and not transformal valuation fields are in question, 8.5 .)

But we prefer to vary the visualization. Let ϕ denote the q-Frobenius, $\Phi \subset V \times V^{\phi}$ the graph of ϕ, $\underline{\Phi} = \Phi \times \mathbb{P}^1$. Then $\underline{S} \cap \underline{\Phi}$ is an algebraic set of dimension 1, forming a system of algebraic curves and perhaps points. But by the dimension theorem, any isolated point of the intersection must be singular. Let $p \in (X_0)_{q,h}(K_q)$, $p \notin W$; then $(p, 0)$ is a smooth point of $(V \times V^{\sigma}) \times \mathbb{P}^1$, so it must lie on one of the curves C of $\underline{S} \cap \underline{\Phi}$. Let $(p', t) \in C$ with t a generic point of $\mathbb{P}^1$. Then $p' \notin \mathrm{pr}\,_1[1](S_t)$ (since otherwise, as (p', t) specializes over K_q to $(p, 0)$, we would have $p \in \mathrm{pr}\,_1[1](S_0) \subset W$.) So $p' \in X(S_t)$. Thus there exists a place of $K_q(t)^a$ into K_q with $p' \mapsto p$. For the valuation corresponding to this place, p lies in the image of $X_t(L)$ under the residue map; but all such valuations are isomorphic. This proves the surjectivity of r_1.

We thus exhibited a bijection between $X(S_0)_{q,h}(K_q) \setminus W = X(S_0)_{q,h}(K_q) \cap V_{et}(S)$ and $X(S_t)_{q,h}(L) \setminus {}^*W(L)$. Now by 13.3 , at most $O(q^{d-1})$ of $X(S_0)_{q,h}(K_q)$ lies on W (even when counted with multiplicities on $X(S_0)_{q,h}(K_q)$.) By 8.18 (cf. 8.19), ${}^*W(L)$ has $O(q^{d-1})$ points. This proves the lemma. □

Lemma 13.5 *Let $S \in \mathbb{Z}\mathcal{C}(V)_{\mathrm{rat}}$. Then $S \sim_a 0$.*

Proof Since 13.4 was stated for $\mathcal{C}(V)_{\mathrm{sep}}$, we wish to reduce to $S \in \mathbb{Z}\mathcal{C}(V)_{\substack{\mathrm{sep}\\ \mathrm{rat}}}$. To this end, let $\phi(x) = x^{p^l}$, where $p^l = \max_T [k[T] : k[V^{\sigma}]]_{\mathrm{insep}}$. Here T runs over all components of S, and any additional components needed to witness the rational equivalence of S to 0. Define τ by $\sigma = \tau\phi$; let $V' = V^{\phi}$; consider the map $g : (V \times V^{\sigma}) \to (V' \times V^{\sigma})$, $g(v, w) = (\phi(v), w)$. Then $g_* S \in \mathbb{Z}\mathcal{C}(V)_{\mathrm{sep}}$. By [Fulton] §1.4, rational equivalence is preserved under proper push-forwards; (straightforwardly), so that $g_* S \in \mathbb{Z}\mathcal{C}(V)_{\substack{\mathrm{sep}\\ \mathrm{rat}}}$. Thus as in 11.19, the statement for $g_* S$ implies the same for S, and we may assume $S \in \mathbb{Z}\mathcal{C}(V)_{\mathrm{sep}}$. (An alternative here, if we wished to use multiplicities, would be to use the projection formula for intersection multiplicities, $|T \cap \Phi_q| = |g_* T \cap \Phi_{q/p^l}|$, and a bound for the geometric multiplicities showing them to be equal.)

We may take $S = [\underline{S}_0] - [\underline{S}_{\infty}]$, $\underline{S}$ an irreducible $d+1$-dimensional variety of $(V \times V^{\sigma}) \times \mathbb{P}^1$. Let t denote a generic point of $\mathbb{P}^1$ over k. Then by 13.4 , $[\underline{S}_t] \sim_a [\underline{S}_0]$ and $[\underline{S}_t] \sim_a [\underline{S}_{\infty}]$. So $[\underline{S}_0] - [\underline{S}_{\infty}] \sim_a 0$, as required. □

We repeat the statement of 13.1 in this language, and prove it.

Corollary 13.6 *Let $S \in \mathcal{C}(V)$. For almost all (q, h), the numbers $|X(S)_{q,h}|$, $(S_{q,h} \cdot \Phi_q)$, $\Phi_q \cap \widetilde{S}$ all differ by at most $O(q^{d-1})$.*

Proof By 13.2, there exists $S' \in \mathbb{Z}\mathcal{C}_f(V)$ with $S - S' \in \mathbb{Z}\mathcal{C}(V)_{\mathrm{rat}}$. By 13.5, $S \sim_a S'$. Since rational equivalence preserves intersection numbers, $S'_{q,h} \cdot \Phi_q = S_{q,h} \cdot \Phi_q$. Thus we may assume $S = S'$, i.e. we may assume $S \in \mathcal{C}_f(V)$. As in 13.5, 11.19, we may assume in addition that $S \in CV sep$. By 6.6, the intersection $S_{q,h} \cap \Phi_q$ on $(V \times V^{\sigma})_{q,h}$ is proper; so $(S_{q,h} \cdot \Phi_q)$ is the number of points of $S_{q,h} \cap \Phi_q$, counted with intersection multiplicities. On $V_{et}(S)_{q,h}$ the geometric multiplicity of the intersection is 1; while on $Z = V \setminus V_{et}(S)$, by 13.3, the total number of points of intersection (with geometric multiplicities) is $O(q^{d-1})$. Thus counting the points of $X(S)_{q,h}$ with intersection or with geometric multiplicities, or without multiplicities agree up to $O(q^{d-1})$. Since the intersection multiplicities are sandwiched in between ([Fulton] Prop. 7.1), and have value 1 whenever the geometric multiplicity is 1, they yield the same result. $\square$

14 Proofs and applications

14.1 Uniformity and Decidability

Throughout this paper we have been dealing with a variable Frobenius on an essentially fixed variety. When the variety is not itself defined over the fixed field of the Frobenius, this is slightly strange; given the algebraic variety $Y = X \times X^{\phi}$ there is (at most) one Frobenius on it. We have explained the uniformity in three different ways:

1. By fixing, not Y and S, but only the *degrees* of their projective completions. (As in Theorem 1.1)
2. Using the formalism of 11.18 (As in 1B)
3. Using difference schemes. In this language, one could state Theorem 1A thus:

Theorem 1.1C *Let D be a finitely generated difference domain, X an absolutely transformally integral difference scheme of total dimension d over* $\mathrm{Spec}^{\sigma} D$. *Then there exist $b \in \mathbb{N}$ and $c \in D$ such that for prime power $q > b$,and any $y \in M_q(D[c^{-1}])$, $X_{q,y}$ is nonempty, indeed has $\delta^* q^d + O(q^{d-1/2})$ distinct points.*

Here if K is the field of fractions of D and L is the transformal function field of X, then $\delta^* = \deg_{lim}(L/K)/\iota'(L/K)$, where $\deg_{lim}$ is the limit degree, and ι' is the purely inseparable dual degree, cf. 5.2.

Uniformity statements such as Theorem 1.1 amount to the same in any of the formulations. To demonstrate this, we show (2) implies (3) implies (1), and (2).

proof that Theorem 1.1B is equivalent to Theorem 1.1C

In effect Theorems 1.1B and 1.1C make the same statement about points of a difference scheme $\mathcal{X}$; but Theorem 1.1B assumes the generic fiber X of $\mathcal{X}$ is directly presented, while Theorem 1.1C assumes instead that X is irreducible. One may pass from the directly presented to the general case using 6.2; and from the irreducible case to the general case using 6.4.

proof of Theorem 1B

We make the assumption stated in 11.28, and use the notation there. (We also accept the conclusion of 11.19.) By Theorem 12.2, $\bar{S} \cdot \Phi_q$ has degree as stated in 1B. Theorem 1B now follows from Proposition 13.1.

$\square$

proof that Theorem 1.1B implies Theorem 1

Suppose given a family X_i of affine varieties over algebraically closed fields k_i of characteristic p_i, and $S_i \subset (X \times X^{\phi_{q_i}})$, with the assumptions of Theorem 1.1 holding, and the degrees of the projective completions of the X_i and S_i bounded, but with $(q^{-(d-\frac{1}{2})}(|S(k) \cap \Phi_q(k)|) - aq^d$ unbounded. Let (L, σ) be an ultraproduct of the (k_i, ϕ_{q_i}). Since the degrees and dimensions are bounded, the corresponding ultraproducts of the X_i and S_i are ordinary subvarieties X,$S \subset X \times X^\sigma$ of finite dimensional affine space over L. Let D be a finitely generated difference sub-domain of L, with X,S defined over L; replacing D by a localization, we may assume X,S remain absolutely irreducible varieties of dimension d when reduced modulo any prime ideal of D. This puts us in the situation of Theorem 1.1B. For almost every i, one has a difference ring homomorphism $D \to k_i$, mapping X to X_i. By Theorem 1.1B, $X_i \cap \Phi_q$ has about aq^d points. A contradiction. □

Corollary 14.1 *Let k be an algebraically closed difference ring, X a difference variety over K of transformal dimension l. Then there exists a finitely generated difference domain $D \subset l$, such that X descends to D, and for almost all q, for all $h : D \to K_q$,* $\dim(X_{q,h}) = l$.

Proof X admits a definable map to $[\sigma]_K \mathbb{A}^l$, whose image contains the complement of a proper difference subvariety of $[\sigma]_K \mathbb{A}^l$. Thus for almost all q, h,$X_{q,h}$ admits a dominant map to $\mathbb{A}^l$, and so has dimension $\geq l$. The converse is 10.8. □

We note one more form of Theorem 1.1; where in effect quasi-finiteness is replaced by the weaker assumption of finite total dimension.

Theorem 14.2 *Let V be a quasi-projective variety over K_q, and let $S \subset (V \times V^{\phi_q})$ be an irreducible subvariety. Assume* $\dim(S) = \dim(V) = d$, *and* $\mathrm{pr_V|S}$ *is a generically quasi-finite map of degree a. If the scheme $S \cap \Phi_q$ is 0-dimensional, then it has size $\leq aq^d + O(q^{d-\frac{1}{2}})$*

Proof By taking ultraproducts, we find (V, S) over a difference domain D, with $X = [\sigma]S \cap \Sigma$ of finite total dimension; and we must show that $|X_{q,h}| \leq aq^d + O(q^{d-\frac{1}{2}})$. By Theorem 1.1, it suffices to show that X has total dimension d, and that every weak component of X of total dimension d is Zariski dense in V. This follows from 16.7. □

ACFA *Proof* of Theorem 1.4 That T_∞ contains ACFA follows from Theorem 1.1. Now it is shown in [Chatzidakis-Hrushovski] that ACFA is nearly complete: the full elementary theory of a model (K, σ) of ACFA is determined by the isomorphism type of $(K_0, \sigma|K_0)$, where K_0 is the subfield of K consisting of points algebraic over the prime field. By the Cebotarev density theorem, every isomorphism type of $(K_0, \sigma|K_0)$ can occur with σ_0 an ultrapower of Frobenius maps. So every completion of ACFA is consistent with $T_\infty = ACFA$. Thus every sentence of T_∞ is a consequence of ACFA.

Decidability The decidability referred to in this paper, in particular in Theorem 1.6, is in the sense of Gödel; it corresponds to the dichotomy: finite/infinite, and not to: finite/bounded, or to distinctions between different degrees of boundedness. *The proof of Proposition 13.1 can routinely be seen to be effective in this sense.* We mention two instances. Suppose one has a smooth variety Y over a finite field, and two subvarieties V,X on

Y, meeting improperly. One knows there exists Y' rationally equivalent to Y, intersecting X transversally. Then without further work, this is effective: it is merely necessary to search (over some finite field extension) for a Y' and for data demonstrating the rational equivalence. To take another example, suppose a sentence is shown to be true in every (boolean-valued) ω-increasing transformal valuation field, by whatever methods. Since the class of such transformal valuation fields is an elementary class, one can search for an elementary formal proof, with assured success; the proof will only use that the valuation is m-increasing, for some m. Thus the statement will be true in every valuation field with the automorphism $\sigma(x) = x^q$, as soon as $q \geq m$; and we have found m effectively.

Theorems I or 1B follow from Propositions 12.2 and 13.1. To use Theorem 12.2, we require an effective bound on the Betti numbers $\dim H^i(X, \mathbb{Q}_l)$ of a smooth projective variety X over an algebraically closed field. In characteristic zero, by Artin's comparison theorem, one can compute the Betti numbers using singular cohomology. For this it suffices to search and find a triangulation. In positive characteristic, the situation is somewhat less clear a priori, but the proof in SGA4 or in [Freitag-Kiehl] of the finiteness of these numbers is everywhere effective. A more explicit reference would be nice but I was unable to find one. We thus state:

Fact 14.3 *Let $b(n,m)$ be the maximum possible Betti number of a smooth subvariety of $\mathbb{P}^n$ of degree m. Then $b(n,m)$ is bounded by a recursive function.*

Proof of Theorem 1.6.
Let T be the theory of all K_p. The completions of T are:

- The theories T_p of the individual K_p; these are just the theory of algebraically closed fields of characteristic p, with an axiom stating $(\forall x)(\sigma(x) = x^p)$. Call this axiom α_p.
- The completions of ${T^{\infty}}_0$, the extension of ACFA stating that in addition, the field has characteristic zero.

Let β_n be the disjunction of α_p for $p \leq n$.

Thus T is axiomatized by sentences of the form:

$$\beta_n \quad \vee \quad \phi$$

with ϕ an axiom of ACFA; the problem is only to know, given ϕ, a value of n for which this disjunction is true. Now the proof of 12.2 estimates an intersection number $\bar{S} \cdot \Phi_q$ as aq^d, with an error of $\leq b(q^{d-\frac{1}{2}})$, with b bounded effectively and independently of q. 13.1 gives effectively another constant c, such that $|S \cap \Phi_q| - (\bar{S} \cdot \Phi_q) \leq cq^{d-1}$. Thus it suffices to choose n such that $cq^{d-1} + b(q^{d-\frac{1}{2}}) < aq^d$ for $q > n$.

This shows that the axioms above are recursively enumerable. On the other hand, every sentence is equivalent, modulo these axioms, to a bounded Boolean combination of existential sentences describing a finite extension of the prime field. The decidability problem for the theory is thus reduced to the same problem for such sentences, and this is settled by Cebotarev. □

14.2 Finite simple groups

In the case of *algebraic* simple groups G of dimension n, over an algebraically closed field, the proof is classical (cf. e.g. Humphreys, Linear Algebraic Groups): any conjugacy class X of G generates G in at most $\dim(G)+1$ steps. One considers X_n, the set of n-fold products of elements of X. Then $\dim(X_n)$ is nondecreasing, and must eventually stabilize: $\dim(X_{n-1}) = \dim(X_n)$, $n \leq \dim(G)$. At this point, one considers the *stabilizer* of the Zariski closure of X_n, and concludes it is a group of dimension equal to $\dim(X_n)$, and containing a conjugacy class. Simplicity of G implies that the stabilizer equals G.

Boris Zilber realized that the proof generalizes to groups of finite Morley rank; the key to his proof was a different definition of the stabilizer, using the dimension theory directly without reference to closed sets.

Difference equations of finite total dimension do not have finite Morley rank, but they do have "finite S1-rank", cf [Chatzidakis-Hrushovski]. More precisely, such a dimension theory applies to their solution sets in a universal domain for difference fields; not necessarily to smaller fields.

In the presence of finite S1-rank, the stabilizer must be defined in a significantly different way than in finite Morley rank; it can however be defined, and enjoys similar properties, sufficient to make the proof go through. See [Hrushovski-Pillay].

The ultraproducts of the groups $G_n(q)$ are a priori a group defined in the same way over the ultraproduct U of the difference fields $GF(p)^{alg}, \phi_q)$. The proof presented in [Hrushovski-Pillay] would then go through, provided one knows that U has an appropriate dimension theory. The present paper completes the proof by showing that U is a universal domain for difference fields.

14.3 Nonstandard powers and a suggestion of Voloch's

We first restate and prove 1.5. Let $F = GF(p)^{\mathrm{alg}}$ be the algebraic closure of a finite field, and let G be the automorphism group. G is isomorphic to the profinite completion of Z, and we consider it with this topology.

Let $F = GF(p)^{\mathrm{alg}}$. Let Υ be the set of automorphisms σ of F such that (F, σ) has ACFA. Then Υ is co-meager in the sense of Baire.

Proof Fix a variety V over F. Then V is defined over $GF(q)$ for some power $q = p^l$ of p. There are l possibilities for $\sigma|GF(q)$; fix one of them, say ϕ_{p^j}; let $V' = V^{\phi_{p^j}}$; and then fix $S \subset V \times V'$ as in the statement of ACFA. There are countably many possible V, j, V', S altogether. It suffices therefore to show that for each such set of data, the set of σ meeting the particular instance of ACFA relative to V, V', S is comeager among all σ with $\sigma|GF(q) = \phi_{p^j}$. This set is an open set of automorphisms, so it suffices to show that it is dense. For this we must show that for any l' and j' with $l|l'$ and $j' = j(modl)$, there exists $\sigma \in \Upsilon$ with $\sigma|GF(p^{l'}) = \phi_{p^{j'}}$, and such that the relevant instance of ACFA holds. However, by Theorem 1.1, any given instance of ACFA holds for any sufficiently large (standard) power of the Frobenius automorphism; in particular one can pick a Frobenius power $p^{j''}$ large enough, and such that $j'' = j' \pmod{l'}$. Thus the open set in question is dense, and Υ is comeager. □

The examples of Cherlin-Jarden

Let l be a positive rational. The equation E_l : $\sigma(x) = lx, x \neq 0$ implies $\sigma^n(x) = l^n x$. Thus E_l has no solutions fixed by σ^n, hence no algebraic solutions. So $(\bar{\mathbb{Q}}, \sigma)$ is not a model of ACFA.

It can also be shown that if $a_l \in E_l$, with l varying over the positive rational primes, then the a_l are algebraically independent over $\mathbb{Q}$. This justifies the assertion made after the statement of 1.5: any model of ACFA of characteristic 0 must have infinite transcendence degree over $\mathbb{Q}$.

However,as a rule, an axiom of ACFA *is* true in $(\bar{\mathbb{Q}}, \sigma)$ for a generic σ. For instance, among order-one difference varieties, this rule admits only a few concrete families of exceptions, of the form $\sigma(x) = f(x)$ with f a fractional linear transformation, and $\sigma(x) = ax + b$, with x ranging over an elliptic curve.

Voloch's question on two primes

Let p be a rational prime, and let Ω_p be the roots of unity of order prime to p. $L_p = \mathbb{Q}(\Omega_p)$ be the maximal prime-to-p cyclotomic extension of $\mathbb{Q}$. Let p_1, p_2 be two distinct primes on the integers of L_p, lying above the rational prime p. Let V be a variety, say given by linear equations L over $\mathbb{Z}$. Voloch suggested investigating the solutions to L modulo *both* primes, with coordinates on Ω_p:

(1) $$V(v_1, v_2) = \{x = (x_1, \ldots, x_k) \in \Omega_p^k : L(x) \in p_1 \cap p_2\}$$

He noted that x_2 is conjugate to x_1 by an automorphism θ of L_p, and that on Ω_p, one may write: $\theta(x) = x^n$, with $n \in (\Pi_{l \neq p} \mathbb{Z}_l)^*$. The question is therefore equivalent to solving, in $GF(p)^{alg}$, the equations:

(2) $$(x_1, \ldots, x_k) \in V, (x_1^n, \ldots, x_k^n) \in V$$

Under what conditions is the number of solutions finite?

In a very special case, the present results are relevant. Assume $n = p^m + 1$, where $m \in (\Pi_{l \neq p} \mathbb{Z}_l)$; or more generally, that $n = f(p^m)$, f is a fixed nonconstant polynomial over $\mathbb{Z}$. Choose $m \in (\Pi_{l \neq p} \mathbb{Z}_l)$ *generically.* Then the structure $(GF(p)^{alg}, x \mapsto x^n)$ is interpretable in $(GF(p)^{alg}, x \mapsto x^{p^m})$. The latter is a model of ACFA by 1.5, hence we have a precise criterion for the answer to (2). In particular, it is easy to see that that if V is defined by a *single* linear form with at least 3 variables, then (1) has infinitely many solutions.

14.4 Consequences of the trichotomy

The Lefschetz principle suggests transferring the results of [Chatzidakis-Hrushovski] to statements about the Frobenius. The propositions below were found in this way; though in retrospect they can be proved using only a very small part of the results of the present paper. Nonetheless the general Lefschetz principle gives the propositions a general context and makes the proofs immediate.

The first and third propositions are direct translations of Theorem 1.10 in [Hrushovski01], using Theorem 1.4. Note that the numerical bounds obtained there can be construed as first-order statements.

Proposition 14.4 *Let $f \in \mathbb{Z}[T]$ be an integral polynomial in one variable, with no cyclotomic factors. Let $L(X_1, \ldots, X_k)$ be a linear form over $\mathbb{Q}$, $k \geq 3$. Let $X(q)$ be the roots of unity of order $f(q)$, in $\widetilde{\mathbb{F}_q}$. Let $Y(q)$ be the set of k-tuples $a = (a_1, \ldots, a_k)$ from $X(q)$ satisfying $L(a) = 0$, but such that no proper sub-sum is zero. Then there exists an absolute bound $B(f, L)$ and p_0 such that for all primes $p \geq p_0$ and all powers q of p, $|Y(q)| \leq B(f, L)$. The bounds p_0 and $B(f, L)$ are effectively computable from f and L.*

Proof If σ is an automorphism of a field K, let $X(\sigma) = \{a \in K : a^{f(\sigma)} = 1\}$, and let $Y(\sigma)$ be be defined analogously to $Y(q)$. By Theorem 1.10 in [Hrushovski01], the axioms of ACFA together with the axiom scheme of fields of characteristic 0, imply that $X(\sigma)$ is finite; say $|X(\sigma)| \leq B$ where $B = B(f, L)$ is finite. Hence a finite set of axioms of ACFA imply that if the field characteristic is larger than some p_0, then $Y(\sigma)$ is finite. By Theorem 1.4, the said axioms hold in all algebraically closed fields of large enough positive characteristic, when $\sigma(x) = x^q$, q any power of the characteristic. □

Remark 14.5

1. An upper bound for $B(f, L)$ can be given explicitly; it is doubly exponential in k and in the degree of f, with coefficients using the absolute values of the coefficients. This follows from [Hrushovski01], Proposition 1.11, by plugging in the parameters.
2. Using a positive - characteristic version of [Hrushovski01], Theorem 1.10, one can determine p_0; the proposition holds for all primes p such that $f(p^{a/b}) \neq 0$ for all rational a/b.
3. One could take L over a number field instead of over $\mathbb{Q}$, still with a uniform bound, independent of L.
4. One can also compute the set of numbers that actually occur as $|Y(q)|$, or and describe the set of q for which a particular number occurs.

By taking $f(T) = T - 2$, we obtain the following corollary; in hindsight it is easy to find a direct proof, but we keep it as an example. For a prime p, let $E_2(p)$ be the set of natural numbers e such that 2 is in the multiplicative subgroup of $(Z/eZ)^*$ generated by p. (This includes the e such that p is a primitive root mod e.) Let

$$R_2(p) = \{x \in GF(p)^{alg} : (\exists e \in E_2(p))(x^e = 1)\}$$

Corollary 14.6 *Let p be an odd prime. Let $L = \sum c_i X_i$ be a linear form in $k \geq 3$ variables. Let Y be the set k - tuples $(a_1, \ldots, a_k) \in R_2(p)^k$ with $\sum c_i X_i = 0$, but no proper sub-sum equals 0. Then $|Y| \leq 2^{k2^{(k-1)}}$*

Proof This follows from the proposition (and the remarks), since if $x \in R_2(p)$ then $x^{q-2} = 1$ for some power q of p.

Similar remarks will apply to the propositions below. First, a non-linear analog. For simplicity we consider one-variable q-nomials over $\mathbb{Z}$, but the results of [Chatzidakis-Hrushovski] allow an analysis of the behavior of any system of q-polynomials, in several variables, with respect to the question raised below; and the bounds obtained are independent of parameters, if the polynomial is over a larger field.

Let $h(X,Y) \in \mathbb{Z}[X,Y]$. Let us say that h is *special* if there exists a curve C, either $C = G_m$ or C an elliptic curve, a coset S of a subgroup of $C \times C$, correspondences $U, V \subset (C \times \mathbb{P}^1)$ (U,V are irreducible curves, projecting dominantly to C and to $\mathbb{P}^1$), and points (a,b,c,d) such that (a,b),(a,c),(b,d), are generic points of S, U,V respectively, and $h(c,d) = 0$.

(If h is special, then the roots of the q-nomial $h(X^q, X)$ are in uniform algebraic correspondence with either roots of unity as in the previous proposition, or an elliptic analog, or the points of $\mathbb{F}_q$.)

Let us say that an affine variety $U \subset \mathbb{A}^n$ is *special* if it is explicitly a product of curves and points; i.e. there exists a partition of the coordinates of $\mathbb{A}^n$, corresponding to an isomorphism $j : \mathbb{A}^n \to \mathbb{A}^{k_1} \times \ldots \mathbb{A}^{k_l}$, such that $j(U) = C_1 \times \ldots \times C_l$, with C_i a point or a curve.

Theorem 14.7 *Let $h_i(X,Y) \in \mathbb{Z}[X,Y]$ be non-special polynomials, and let $U \subset \mathbb{A}^n$ be an affine variety over $\mathbb{Z}$. Then either h is special, or there exists a finite union W of special subvarieties of U (with W defined over $\mathbb{Q}$) such that for all sufficiently large primes p, and all powers q of p, if $h_i(a_i, (a_i)^q) = 0$ and $(a_1, \ldots, a_n) \in U$, then $(a_1, \ldots, a_n) \in W$.*

Proof Assume h is not special. By the trichotomy theorem of [Chatzidakis-Hrushovski], in any model $\mathbb{Q} \leq M \models ACVF$,

$$X = \{x : h(x, \sigma(x))\}$$

is a stable, stably embedded definable set, whose induced structure is superstable of rank one, and *distintegrated*: the algebraic closure relation on $X(M) \setminus \mathbb{Q}^a$ is an equivalence relation.

Let $U \subset \mathbb{A}^n$ be an affine variety. If $(a_1, \ldots, a_n) \in U$ and $a_i \in X$, let $w_o(a) = \{i \leq n : a_i \in \mathbb{Q}^a\}$, $w_+(a) = \{i \leq n : a_i \notin \mathbb{Q}^a\}$ and let $\{w_k(a) : k = 1, \ldots, m\}$ be the classes of the equivalence relation on $w_+(a)$ defined by: $i \sim j$ iff $\mathbb{Q}(a_i)^a = \mathbb{Q}(a_j)^a$. Let $a(k) = (a_i : i \in w(k))$, so that we can write $a = (a(0), a(1), \ldots, a(m))$. Let $C(k)$ be the locus over $\mathbb{Q}^a$ of $a(k)$; so $C(0)$ is finite, and $C(k)$ is a curve for $k \geq 1$. Let $S(a) = \{a(0)\} \times C(1) \times \ldots \times C(m)\}$; it is a special variety, defined over $\mathbb{Q}^a$. By disintegration of X, the fields $\mathbb{Q}(a(0)), \ldots, \mathbb{Q}(a(m))$ are linearly free. Thus a is a generic point of $S(a)$. Since $a \in U$, we have: $S_a \subset U$. Now a was an arbitrary point of $U(M) \cap X(M)^n$, in a model M, and so by compactness, there are finitely many special S_i, defined over $\mathbb{Q}^a$, such that $U(M) \subset \cup_i S_i$. Further taking the union of all $\mathbb{Q}$-conjugates, we find a $\mathbb{Q}$- definable finite union S of special varieties, with $U \cap X^n \subset S$ (in any difference field extension of $\mathbb{Q}$.) The statement on the large primes follows immediately by compactness. □

Similar results exist for q-nomials of higher order, and for commutative algebraic groups. In particular, for semi-Abelian varieties, we have:

Proposition 14.8 *Let A be a commutative algebraic group scheme over a scheme Y of of finite type over $\mathbb{Z}$, with generic fiber A_Y a semi-Abelian variety. Let f be an integral* *polynomial in one variable, with no cyclotomic factors. Let V be a subscheme of A. For any closed point $y \in Y$, let A_y be the fiber of A over y. Let ϕ denote the Frobenius endomorphism of A_y, relative to the residue field k_y, and let ψ be any finite power of ϕ. Let $X(y, \psi, f)$ be the kernel of the endomorphism $f(\psi)$ on $A(\widetilde{k_y})$.*

Then there exists a Zariski open $Y_0 \subset Y$ and finitely many group subschemes B_i of A over Y_0, with $B_i \subset V$, and an effectively computable integer B, such that for any closed point $y \in Y$, $X(y, \psi, f) \cap V$ is contained in at most B translates of some of the B_i.

The proof is analogous to that of 14.7. □

15 Jacobi's bound for difference equations

Consider a system u of n difference equations $u_1, \ldots, u_n$ in n variables $x_1, \ldots, x_n$; assume u_i has order h^i_k with respect to the variable x_i. Assume that the system defines a difference scheme of finite total dimension d. What bound can one impose on d, based on the data (h^i_k)?

Jacobi considered this problem for differential equations, and gave as a bound:

$$\mathbf{j} = \max_{\theta \in Sym(n)} \Sigma^n_{k=1} h^{\theta(k)}_k$$

Jacobi's proof was criticized by Ritt [Ritt], and the problem is considered open. See [Ritt], [KMP], [Cohn79] for partial results.

The transposition to difference equations was made by Cohn, with analogous partial results; cf. [Lando], [KMP]. Jacobi's original idea was reduction to the linear case (the "equations of variation"), and then the the case of constant coefficients; and present results still follow this line. However when the equations have hidden singularities, the reduction to the linear case requires additional assumptions, or new ideas.

We show here that Jacobi's bound is correct for difference equations by a quite different method, Frobenius reduction. We also get an "explanation" for the curious determinant-like formula.

Cohn refined the problem by setting $h^i_k = -\infty$ if x_i does not occur in u_k; we accept this refinement. (It not only gives a better bound in many cases, but allows for greater flexibility in manipulating the equations.)

Consider first the analogous problem for *algebraic* equations. Let $U_1, \ldots, U_n$ be polynomials in n variables $x_1, \ldots, x_n$, over a field k; suppose U_i has degree H^j_i in the variable x_i. Let $Z(U)$ be the scheme cut out by $\{U_i\}$, and let Z_0 be the 0-dimensional part of $Z(U)$.

Lemma 15.1 $|Z_0| \leq \mathbf{J} = \sum_{\theta \in Sym(n)} \prod^n_{k=1} H^{\theta(k)}_k$

Proof Let $V = (\mathbb{P}^1)^n$ be the n'th Cartesian powers of the projective line. The tangent bundle of V is generated by its global sections, since this is true for each factor $\mathbb{P}^1$. The normal bundle to V embedded diagonally in V^n is hence also generated by its global sections. Thus by [Fulton] Theorem 12.2, each contribution to the intersection of V with a subvariety of V^n of codimension n is represented by a non-negative 0-cycle. The number of components of intersection, in particular the number isolated points, is therefore bounded by the intersection number (just as in [Fulton], Example 12.3.)

Each equation U_i determines a hypersurface $[U_i]$ in V (the closure in V of the subscheme of $(\mathbb{A}^1)^n$ cut out by U_i.) Z_0 is the 0-dimensional part of $[U_1]\cap\cdots\cap[U_n] \simeq ([U_1]\times\cdots\times[U_n])\cap V$. By the above discussion, $|Z_0| \leq [U_1]\cdot\ldots,\cdot[U_n]$. It remains to compute this intersection number.

Now $[U_i]$ is rationally equivalent to $\sum_{j=1}^n h_i^j D_j$, where D_j is the divisor defined by the equation $x_j = 0$, the pullback of a point on the j'th copy of $\mathbb{P}^1$. We have $D_j^2 = 0$, since $\{pt\}^2 = 0$ in $\mathbb{P}^1$. Thus $D_\theta(1)\cdot\ldots\cdot D_\theta(n) = 0$ if θ is not injective, while $D_\theta(1)\cdot\ldots\cdot D_\theta(n) = D_1\cdot\ldots\cdot D_n = \{pt\}$ if θ is a permutation. Thus

$$[U_1]\cdot\ldots,\cdot[U_n] = \prod_{i=1}^n \sum_{j=1}^n h_i^j D_j = \mathbf{J}$$

□

Theorem 15.2 *Let k be a difference field $u_i \in k[x_1,\ldots,x_n]_\sigma$ a differential polynomial of order h_i^j in x_j. Let W be a component of the difference scheme Y cut out by $h_1 = \cdots = h_n = 0$. Then the total dimension of W, if finite, is no larger than Jacobi's bound $\mathbf{j}$.*

Proof By assumption, W is transformally reduced, of total dimension w. Let v be a differential polynomial that does not vanish on W, but vanishes on every other component of Y. Adding a variable y and the equation $vy = 1$ does not change the Jacobi bound (using Cohn's convention, the new equation has order 0 in the variable y, and all other equations have order $-\infty$ in this variable.) Thus we may localize away from the other components; so we may assume that Y itself has finite reduced total dimension w. We must show that $w \leq \mathbf{j}$.

Let $D \subset k$ be a finitely generated difference domain, with $u_i \in D[x_1,\ldots,x_n]_\sigma$. Given a homomorphisms $\phi : D \to K_q$, let U_i be the polynomial obtained by replacing $\sigma(x)$ by x^q and d by $\phi(d)$ in u_i.

Let $H_i^j = \deg_{x_j} H_i$. Writing u_j as a sum of σ-monomials, they each involve u_j^ν for some $\nu \in \mathbb{N}[\sigma]$, and we have $H_i^j = \nu_j(q)$ for the highest such ν. This ν has the form $d\sigma^{h_i^j} + \cdots$ (lower terms), so $H_i^j = dq^{h_i^j} + O(q^{h_i^j - 1})$. Thus $\lim_{q\to\infty}\log_q H_i^j = h_i^j$.

By 10.8, for some D, for almost all q and ϕ, $Y_q = \{x : U_1(x) = \cdots = U_n(x) = 0\}$ is *finite*; and by Theorem 1.1C, for some $a, b \in \mathbb{Q}$ with $a > 0$, $|Y_q| = aq^w + e_q, |e_q| \leq bq^{w-1/2}$. On the other hand , by 15.1, $|Y_q| \leq \mathbf{J}$. So

$$aq^w + e_q \leq \sum_{\theta\in Sym(n)} \prod_{k=1}^n H_k^{\theta(k)}$$

Taking $\log_q$ and letting $q \to \infty$ we obtain

$$w \leq \max_{\theta\in Sym(n)} \sum_{k=1}^n h_k^{\theta(k)}$$

□

Remark 15.3 Ritt's conjecture for difference equations.

As Cohn observed ([Cohn87] for differential equations, personal communication for difference equations), the validity of the refined Jacobi bound implies the Ritt dimension conjecture: given a system of m difference equations in n variables, if $m < n$ then every component of the solution set has transformal dimension $\geq n - m$. This can also be deduced directly from Theorem 1.1, as in 16.1.

16 Complements

16.1 Transformal dimension and degree of directly presented difference schemes

Here is a variant of 6.4, valid for all components. The proof *assumes* Theorem 1.4, and illustrates the way the main result of this paper may be used in difference algebra.

Proposition 16.1 *Assume theorem 1.4. Let X be a smooth algebraic variety over a difference field K. Let $S \subset X \times X^\sigma$ be an absolutely irreducible subvariety,* $\dim(X) = d$, *and assume* $\dim S = e + d$. *Let*

$Z = \{x \in X : (x, \sigma(x)) \in S\}$

Then any component of Z has transformal dimension $\geq e$.

Proof Let $Z(1), \ldots, Z(m)$ be the components. Suppose for contradiction that trans. dim. $(Z(1)) < e$. We may assume K is the field of fractions of a finitely generated difference domain D. By the main theorem, for almost all q, and almost $y \in \operatorname{Spec} M_q(D)$, $M_q(Z)_y = \cup_j M_q(Z(j)_y)$, $M_q(Z(1))_y \not\subset \cup_{j>1} M_q(Z(j))_y$, and $\dim(M_q(Z(j))_y) \leq$ trans. dim. $(Z(1)) < e$. It follows that $M_q(Z)_y$ has a component of dimension $< e$. However $M_q(Z)_y = S_y \star \Phi_q$. Since (for almost all y) $S_y \subset X_y \times (X^\sigma)_y$ is irreducible of dimension $d + e$, and $X_y \times (X^\sigma)_y$ is smooth, the dimension theorem implies that every component of this intersection has dimension $\geq (d + e) + d - 2d = e$; a contradiction. □

Remark 16.2 *Moreover, in 16.1, if $h : Z \to W$ is a morphism of difference schemes, and W has transformal dimension* 0*, then each component of each fiber of h has transformal dimension $\geq e$.*

The proof is similar, using the fact that $M_q(W)_y$ is finite, and thus the components of the fibers of $M_q(h)_y$ are also components of $M_q(Z)_y$.

Note that 16.1 is equivalent to to following purely geometric statement:

Lemma 16.3 *Let U, V, S be quasi-projective varieties over an algebraically closed difference field k, with $V \subset U^\sigma$, and $S \subset (U \times V)$. Let $\mathcal{T}$ be the set of absolutely irreducible varieties $T \subset S$ such that $(pr_1 T)^\sigma = pr_2 T$. (Where $pr_i T$ is the Zariski closure of the i'th projection.) Let $e \geq 1$. Assume* $\dim(S) \geq \dim(U) + e$. *Then any maximal $T \in \mathcal{T}$ satisfies* $\dim(T) \geq \dim(pr_1 T) + e$.

Problem 16.4 *Is there a simple direct proof of 16.3? Is the smoothness assumption necessary, even if one just wants one component of transformal dimension e?*

The difficulty here is associated with singularities. If all varieties encountered were smooth, it would be easy to prove 16.3: Let $T_0 \in \mathcal{T}$. Let U' be a component of $pr_1[1](S)$ (cf. 6.6) containing $pr_1(T_0)$. Let $S' = (U' \times U) \cap S$; then $\dim(S') \geq \dim(S) - (\dim(U) - \dim(U')) \geq \dim(U') + e$. Thus one can proceed by induction.

In attempting to find a direct proof of 16.3, I considered projecting at the right moment to $\mathbb{P}^m$ by a finite map, and using the dimension theorem there. To do this, one needs to know Lemma 16.6 below; and in order to reduce to the case of purely positive transformal dimension, one finds oneself requiring a stronger statement than 16.3, namely 16.2 above. I did not carry through this proof to the end, but the lemmas that came up seem sufficiently suggestive in themselves to be stated here.

Definition 16.5 *An irreducible difference variety U over K is* purely positive transformal dimensional *if every differential rational morphism on X into a difference variety of transformal dimension* 0 *is constant.*

Proposition 16.6 *Let K be a difference field, U, V algebraic varieties over K, $f : U \to V$ a quasi-finite map. Let X be an irreducible difference subvariety of V, Zariski dense in V. Assume X is purely positive transformal dimensional. Then for some Zariski open $\widetilde{U} \subset U$, $f^{-1}(X) \cap U$ is an irreducible difference variety (and also purely positive transformal dimensional.)*

Proof [Chatzidakis-Hrushovski].

Corollary 16.7 *Let U, S be affine varieties over an algebraically closed difference field k, with $S \subset (U \times U^{\sigma})$. Let $X = S \star \Sigma \subset [\sigma]_k U$. Then either X has positive transformal dimension, or else it has total dimension* $\leq \dim(U)$.

In the latter case, if U is absolutely irreducible, every weak component of X of total dimension equal to $\dim(U)$ *is Zariski dense in U.*

Proof Let $(a_0, a_1, \ldots)$ be a generic point of a weak component C of U. (I.e $(a_0, a_1, \ldots, a_n) \to (a_1, \ldots, a_{n+1})$ is a specialization over k, but not necessarily an isomorphism; and $(a_0, a_1) \in S$.)

Let $d_n = tr.deg._k k(a_0, \ldots, a_n)$. If, for some n, $d_n < d_{n+1}$, let U_n, V_n, S_n be the k-loci of $(a_0, \ldots, a_n), (a_1, \ldots, a_{n+1}), (a_0, \ldots, a_{n+1})$ respectively. Then the hypothesis of 16.3 holds. So there exists a point c in a difference field extension of k, with $c_n = \sigma^n(c)$ satisfying: $(c_0, \ldots, c_{n+1}) \in S_n$, and of positive transformal dimension over k. In particular, $(c_0, c_1) \in S$, so $c \in X$ and X has positive transformal dimension.

Otherwise, $d_{n+1} \leq d_n$ for each n. So $d_n \leq d_0 \leq \dim(U)$ for each n. This shows that C has total dimension $\leq \dim(U)$. Moreover if c_0 is not a generic point of U, then C has total dimension $< \dim(U)$. □

16.2 Projective difference schemes

$\mathbb{N}[\sigma]$- Graded rings Let $\mathbb{N}[\sigma]$ denote the set of polynomials over $\mathbb{N}$, with indeterminate labeled σ. Consider difference rings R graded by the ring $\mathbb{N}[\sigma]$. In other words, we are provided with a decomposition of R as an Abelian group:

$$R = \oplus_{n \in \mathbb{N}[\sigma]} H_n$$

Multiplication in R induces maps $H_n \times H_m \to H_{m+n}$. The action of σ is assumed to carry H_n to $H_{n\sigma}$. Such a structure will be called a $\mathbb{N}[\sigma]$-graded ring.

We will assume R is generated as a difference ring by $H_0 \cup H_1$. A difference ideal P is called *homogeneous* if it is generated by the union of the sets $P \cap H_n$.

Projective difference schemes Let R be a $\mathbb{N}[\sigma]$-graded ring, with homogeneous components H_n. We define an associated difference scheme, $\mathrm{Proj}^{\sigma}\,(R)$, as follows. The underlying space is the set of homogeneous transformally prime ideals, not containing H_1. The topology is generated by open sets of the form: $W_a = \{p : a \notin p\}$, with $a \in H_k$ for some $k \in \mathbb{N}[\sigma]$. Such sets are called affine open. One assigns to W_a the difference ring:

$$R_a = \{R[a^{-n} : n \in \mathbb{N}[\sigma]]\}_0 = \{\frac{b}{a^n} : n \in \mathbb{N}[\sigma], b \in H_{nk}\}$$

and glues.

Explanation: the localized difference ring $R[a^{-n} : n \in \mathbb{N}[\sigma]]$ has a natural grading, with an element of the form b/a^n ($b \in R, n \in \mathbb{N}[\sigma]$) in the homogeneous component of degree $\deg(b) - n\deg(a)$. $\{R[a^{-1}]\}_0$ is the degree-zero homogeneous component of this $\mathbb{N}[\sigma]$-graded ring. To verify that this yields, uniquely, a difference scheme structure, observe that $\{R[(a_1a_2)^{-n} : n \in \mathbb{N}[\sigma]]\}_0$, the affine ring corresponding to $W_{a_1a_2}$, is the difference ring localization of the ring corresponding to W_{a_1} by the element $\frac{a_2^{\deg a_1}}{{a_1}^{\deg a_2}}$. Note also that the prime ideals of R_a can be identified with the elements of W_a (a homogeneous prime ideal of R is the kernel of a homomorphism $h : R \to L$, L a difference field, with $h(a) \neq 0$. Then h restricts and extends to R_a. Conversely, let $g : R_a \to L'$ be a homomorphism. Then g extends to a graded homomorphism $\bar{g} : R \to L'[t^{\mathbb{N}[\sigma]}]$, with $\bar{g}(a) = t$.)

We could alternatively directly describe the structure sheaf in terms of section, as in the definition of Spec^{σ} ; the local ring at p is defined as

$$\{\frac{b}{a} : (\exists k)\ a, b \in H_k,\ b \notin p\}$$

$\mathbb{N}[\sigma]$-graded ring associated to a graded ring Let D be a difference domain, and let R be a graded D-algebra in the usual sense, $R = \oplus_{i \in \mathbb{N}} R_i$. Assume the homogeneous component of degree 0 is a domain.

Lemma 16.8 *$[\sigma]_D R$ has a unique $\mathbb{N}[\sigma]$-grading, compatible with the grading of R. $\{[\sigma]_D R\}_0 = [\sigma]_D(R_0)$.*

Proof For $n \in \mathbb{N}[\sigma]$, $n = \sum_i k_i \sigma^i$, let H_n be the subgroup of $[\sigma]_D R$ generated by the products $\Pi_i \sigma^i(a_i)$ with $a_i \in R_{k_i}$. (We write here a_i also for the image of a_i in $[\sigma]_D R$.)

Clearly the H_n generate $[\sigma]_D R$ between them, and any $\mathbb{N}[\sigma]$-grading must make H_n the homogeneous component of degree n. Thus it is only a question of showing that the H_n are in direct sum. Let F be a free R_0-algebra, graded so that the free generators span F_1 as a R_0-module, and let $h : F \to R$ be a surjective homomorphism of graded R_0-algebras. Let $\bar{F} = [\sigma]F$; it is clear that the lemma holds for $\bar{F}$. Let J be the $\mathbb{N}[\sigma]$-ideal generated by $ker(h)$. Also J is generated by $J \cap \bar{F}_1$, so J is homogeneous, and $J \cap \bar{F}_0 = (0)$. The $\mathbb{Z}$-graded part of $(\bar{F})/J$ is $F/ker(h) \cong R$. Thus we obtain a map $R \to (\bar{F})/J$. This map extends to a difference ring map $[\sigma]_D R \to (\bar{F})/J$. We also have a map $\bar{F} \to [\sigma]_D R$. By the universal properties, these are isomorphisms, and $[\sigma]_D R$ is $\mathbb{N}[\sigma]$-graded.

Difference structure on projective schemes Let D be a difference domain.

Lemma 16.9 *Let R be a graded D-algebra. Then* $\mathrm{Proj}^\sigma\, [\sigma]_D R$ *is the difference scheme obtained by gluing together* $\mathrm{Spec}^\sigma\, [\sigma]_D R_a$*, where R_a runs over the various localized rings $R_a = \{R[a^{-1}]\}_0$, $a \in R_1$.*

Proof First, for any $a \in R_n, n \in \mathbb{N}$, $[\sigma]_D(R[a^{-1}]) = ([\sigma]_D R)[a^{-1}]$, the second localization being taken in the sense of difference rings. This is clear, since both are the universal answers to the following problem: a difference ring $\bar{R}$, a homomorphism $h : R \to \bar{R}$, with $h(a)$ invertible in $\bar{R}$. Next, one verifies that when R is graded, the two induced gradings on these rings are the same; in particular, $[\sigma]_D(R_a) = ([\sigma]_D R)_a$, using the notation of the definition of Proj^σ , and the previous lemma. □

Lemma 16.10 *There are canonical isomorphisms:*

- $M_q([\sigma]_D R) = R \otimes_D M_q(D)$
- $M_q(\mathrm{Spec}^\sigma\, [\sigma]_D R) = \mathrm{Spec}\,(R) \times_{\mathrm{Spec}\, D} M_q(\mathrm{Spec}\, D)$
- *When R is a graded D-algebra,* $M_q(\mathrm{Proj}^\sigma\, ([\sigma]_D R)) = \mathrm{Proj}\,(R \otimes_D M_q(D))$

Proof

- By the universal properties.
- Apply Spec to the previous.
- By 16.9 and the previous item.

Multi-projective varieties Suppose R is graded by $\mathbb{N}[\sigma]^k$, in place of $\mathbb{N}[\sigma]$. We can define $\mathrm{Proj}^{\mathbb{N}[\sigma]^k} R$ analogously to the case $k = 1$. Sometimes if the intended structure of R is clear we will just write $\mathrm{Proj}^\sigma\, R$ or $\mathrm{Proj}\, R$. The underlying space is the set of homogeneous difference ideals, not containing a homogeneous component R_j (for any $j \in \mathbb{N}[\sigma]^k$). The basic affine open sets are the sets of primes not containing $b = a_1 a_2 \dots a_k$, where a_i has degree $(0, \dots, 0, 1, 0, \dots, 0)$. The localization $R[b^{-1}, \sigma(b^{-1}), \dots]$ has a graded structure, and the associated ring is defined as the graded component of degree 0.

Note 16.11

Let $p \in \operatorname{Proj} R$. Then there exists a_i of degree $(0, \ldots, 0, 1, 0, \ldots, 0)$ with $a_i \notin p$. Since p is homogeneous, and prime, the a_i are algebraically independent over $R_0/(p \cap R_0)$. Thus if R is a domain, and a D-algebra, the transformal transcendence degree of R over D is k more than the transformal dimension of Proj R over D.

The ideals J_q used in defining M_q are not homogeneous. However, a ring graded by $\mathbb{N}[\sigma]^k$ can be (forgetfully) viewed as graded by $\mathbb{N}^k$, in many ways; any ring homomorphism from $\mathbb{N}[\sigma]$ to $\mathbb{N}$ will give such a way. Given q, let $h_q : \mathbb{N}[\sigma] \to \mathbb{N}$ be the homomorphism with $\sigma \mapsto q$, and also let $h_q : \mathbb{N}[\sigma]^k \to \mathbb{N}^k$ denote the product homomorphism. Use h_q to view R as $\mathbb{N}^k$ graded: the graded component of degree n is by definition the sum of the graded components of degree ν, over all ν with $h_q(\nu) = n$. Then J_q is homogeneous for this grading; the generators can be taken to be $\sigma(r) - r^q$ with r in some homogeneous component, and these elements are homogeneous in the h_q-induced grading. Therefore $M_q(R) = R/J_q(R)$ has a natural $\mathbb{N}^k$-graded structure. We can thus view M_q as a functor on $\mathbb{N}[\sigma]^k$-graded difference rings into $\mathbb{N}^k$-graded rings.

Notation 16.12 *The difference polynomial ring in one variable over a difference ring R is denoted $R[t^{\mathbb{N}[\sigma]}]$. The localization of this ring by t is denoted $R[t^{\mathbb{Z}[\sigma, \sigma^{-1}]}]$, and is called the $\mathbb{Z}[\sigma, \sigma^{-1}]$-polynomial ring over R. The k-variable version $R[t_1^{\mathbb{Z}[\sigma, \sigma^{-1}]}, \ldots, t_k^{\mathbb{Z}[\sigma, \sigma^{-1}]}]$ is called the $\mathbb{Z}[\sigma, \sigma^{-1}]^k$-polynomial ring over R.*

Note:

Lemma 16.13 *Let $R = \oplus\{R_a : a \in \mathbb{N}[\sigma]^k\}$ be an $\mathbb{N}[\sigma]^k$-graded ring. Let $a_i \in R$ have degree $(0, \ldots, 0, 1_{(i)}, 0, \ldots, 0)$. Let $S = R[{a_i}^{\sigma^n} : i = 1, \ldots, k,\ n \in \mathbb{Z}]$ be the localization by $a_1, \ldots, a_k$, and let S_0 be the degree-0 component. Then S is the $\mathbb{Z}[\sigma, \sigma^{-1}]^k$- polynomial ring over S_0:*

$$S = S_0[{a_1}^{\mathbb{Z}[\sigma, \sigma^{-1}]}, \ldots, {a_k}^{\mathbb{Z}[\sigma, \sigma^{-1}]}]$$

Proof An element c of degree $(n_1, \ldots, n_k)$ of this ring can be written as $b{a_1}^{n_1} \ldots {a_k}^{n_k}$, where $b = c{a_1}^{-n_1} \ldots {a_k}^{-n_k} \in S_0$. By homogeneity, any difference - algebraic relation among the a_i over S_0 implies a monomial relation among them. But $ba_1^{n_1} \ldots {a_k}^{n_k} = 0$ implies $b = 0 \in S_0$, since the a_i are invertible.

Lemma 16.14 *Let R be a $\mathbb{N}[\sigma]^k$-graded ring. Then $M_q(\mathrm{Proj}^\sigma\,(R)) \cong \mathrm{Proj}\,(M_q(R))$*

Proof To simplify notation we treat the case $k = 1$. Let $a \in R_1$. By 16.13, $R[a^{-1}, a^{-\sigma}, \ldots] = R_a[a^{\mathbb{Z}[\sigma]}]$ is a $\mathbb{Z}$-polynomial ring over R_a. Thus $M_q(R_a) = \{M_q(R_a[a^{\mathbb{Z}[\sigma]}])\}_0$.

Write $\bar{a}$ for the image of a modulo J_q. Then

$$M_q(R[a^{-1}, a^{-\sigma}, \ldots]) = M_q(R_a)[\bar{a}^{-1}]$$

Taking degree-0 components,

$$M_q(R_a) = \{M_q(R[a^{-1}, a^{-\sigma}, \ldots])\}_0 = \{M_q(R)[\bar{a}^{-1}]\}_0 = (M_q(R))_{\bar{a}}$$

Thus

$$M_q(\mathrm{Spec}^\sigma\,(R_a)) = \mathrm{Spec}\,(M_q(R)_{\bar{a}})$$

Now the difference schemes $\mathrm{Spec}^{\sigma}(R_a)$ glue together to give $\mathrm{Proj}^{\mathbb{N}[\sigma]^k} R$, while the schemes $\mathrm{Spec}\,(M_q(R)_a)$ glue together to give $\mathrm{Proj}\,(M_q(R))$. (As a runs through R_1, $\bar{a}$ runs through a generating set for $M_q(R)_1$.) The lemma follows after verifying that the gluing maps agree.

Remark 16.15

If D is a difference domain, $D[X,Y] = D[{X_1}^{\mathbb{N}[\sigma]}, \dots, X_n^{\mathbb{N}[\sigma]}, {Y_1}^{\mathbb{N}[\sigma]}, \dots, {Y_m}^{\mathbb{N}[\sigma]}]$, bi-graded so that $X_1^{a_1} \dots X_n^{a_n} Y_1^{b_1}, \dots, Y_m^{b_m}$ is homogeneous of degree $(\sum a_i, \sum b_i)$, then $\mathrm{Proj}\, D[X,Y]$ is the product over D of the σ- projective spaces of dimensions m,n. Conversely, Proj of any bi- or multi-graded ring can be viewed as a fiber product of simple Proj^{σ} of some component rings.

Morphisms between graded rings Let R be a $\mathbb{N}[\sigma]^k$-graded ring, and S a $\mathbb{N}[\sigma]^l$-graded ring. Let $h : R \to S$ be a surjective σ - ring homomorphism. Assume $h(R_c) \subset S_{\lambda(c)}$, where $\lambda : \mathbb{N}[\sigma]^k \to \mathbb{N}[\sigma]^l$ is an injective $\mathbb{N}[\sigma]$-linear map. Then one obtains a map $h^* : \mathrm{Proj}\, S \to \mathrm{Proj}\, R$, as follows. If q is a homogeneous transformally prime ideal on S, not containing a homogeneous component, then so is $h^{-1}(q)$ on R. Let $h^*(q) = h^{-1}(q)$. h^* is continuous: the open set W_a^R defined by an element $a \in R$ pulls back to the set $W_{h(a)}^S$. Finally h induces a difference ring homomorphism on the graded rings, $R[a^{-1}, a^{-\sigma}, \dots] \to S[h(a)^{-1}, h(a)^{-\sigma}, \dots]$, respecting the grading in the same sense as h; and in particular induces a map $R_a \to S_{h(a)}$.

16.3 Blowing up

We give two constructions of blowing up. We show that the transformal blowing-up reduces under Frobenius to a scheme containing the usual blowing-up, and of the same dimension, without resolving the interesting questions regarding their exact relation. One construction can be viewed as the result of a deformation along the affine t line (of transformal dimension 1) while the other is a limit of a deformation along $t = \sigma(t)$ (total dimension 1.) The latter seems to have no analog in the Frobenius picture. Nevertheless they are shown to give (almost) the same result.

Blowing up difference ideals Let R be a difference ring, $X = \mathrm{Spec}^{\sigma}(R)$, and J a finitely generated difference ideal. Consider the ordinary polynomial ring $R[t]$ over R, and set $\sigma(t) = t$. Let $R[Jt]$ be the subring of $R[t]$:

$$R[Jt] = \sum_{n \in \mathbb{N}} (Jt)^n$$

Observe that $R[Jt]$ is finitely generated as a difference ring, if R is. Indeed if V is a set of generators for J as a difference ideal, and Y a set of generators for R as a difference ring, then $Y \cup Vt$ is a set of generators for $R[Jt]$.

Moreover, the degree 0 difference subring of the difference ring localization $R[Jt][(at)^{-1}]$ is finitely generated as a difference ring. If $a \in V$, it is generated by $Y \cup \frac{Vt}{at} \cup \left\{ \frac{\sigma(a)t}{at}, \frac{at}{\sigma(a)t} \right\}$.

View $R[Jt]$ as a (non-Noetherian) graded ring with an endomorphism; form the ordinary scheme-theoretic Proj , [Hartshorne] II 2; and consider the difference subscheme described in

3.10 above. Let

$$\widetilde{X}_J = \mathrm{Fix}^\sigma \, \mathrm{Proj}\,(R[Jt])$$

A basic open affine of $\mathrm{Proj}\,(R[Jt])$ has the form $W_a = \mathrm{Spec}\,(R[\frac{J}{a}])$, with $a \in J$. The intersection of W_a with the set of difference ideals is contained in $\mathrm{Spec}\,(R[\frac{J}{\sigma^n(a)}])$ for each n. Thus one sees that $\widetilde{X}_J$ has an open covering by open subschemes of the form $\mathrm{Spec}^\sigma R[\frac{J}{a}]_\sigma$, $a \in J$, where the square brackets $[\ ]_\sigma$ here refer to localization of difference rings.

There is a natural map $\pi : \widetilde{X}_J \to X$, considered as part of the structure.

Compatibility with localization There are two (closely related) points here.

First, suppose $\pi : \widetilde{X}_J \to X$ is the blowing up at J, $X' = \mathrm{Spec}^\sigma R'$ is an open subscheme of X, $R' = R[a^{-1}, a^{-\sigma}, \ldots]$, with inclusion map $i : X' \to X$, and $J' = R'J$. Then $\widetilde{X'}_{J'} = \pi^{-1}(X')$. This is immediately verified.

Secondly, suppose J' is another finitely generated difference ideal, $J \subset J'$, and for every $p \in \mathrm{Spec}^\sigma R$, if R_p is the local ring, then $JR_p = J'R_p$. Then $\widetilde{X}_J = \widetilde{X}_{J'}$. Indeed, every prime ideal containing J must contain J', so $\widetilde{X}_{J'}$ is covered by the open affines $\mathrm{Spec}^\sigma R[\frac{J'}{a}]$ with $a \in J$. Next, one may find an open covering of $\mathrm{Spec}^\sigma R$, such that the restrictions of J and J' to each open set agree. It follows that $\mathrm{Spec}^\sigma R[\frac{J'}{a}] = \mathrm{Spec}^\sigma R[\frac{J}{a}]$ since they are equal locally (on R).

Now define $\widetilde{X}_{\mathcal{J}}$ and

$$\pi : \widetilde{X}_{\mathcal{J}} \to X$$

for a general difference scheme X and quasi-coherent σ - ideal presheaf $\mathcal{J}$ on X, by gluing.

If $i : W \to X$ is a closed subscheme of a scheme X, with corresponding difference ideal sheaf $\mathcal{J} = ker(\mathcal{O}_X \to i_* \mathcal{O}_W)$, we write

$$\widetilde{X}_W = \widetilde{X}_{\mathcal{J}}$$

We define the *exceptional divisor* E_W to be the difference scheme inverse image of W under this map.

A second construction: the closed blowing up of ideals In order to obtain an embedding of the blowing-up in σ-projective space, we use a different construction. For this construction, we will blow up an ideal (or quasi-coherent ideal presheaf) I instead of a difference ideal (or difference ideal sheaf) J. The two constructions coincide when $I = J$ is both a finitely generated ideal and a difference ideal, so there is no confusion in denoting both by $\widetilde{R}_I$ or $\widetilde{R}_J$. However, to emphasize the difference we will temporarily use a superscript ^{c} for the closed blowing-up of ideals.

Let R be a difference ring, and let I be a finitely generated ideal. We define the closed blowing-up ring $\widetilde{R}^c_I$ as a subring of $R[t^{\mathbb{N}[\sigma]}]$:

$$\widetilde{R}^c_I = \sum_{n \in \mathbb{N}[\sigma]} I^n t^n$$

where if $n = \sum_i \sigma^{k(i)}$, I^n is the subgroup of R generated by elements of the form $\Pi_i f_i$, with $f_i \in \sigma^{k(i)}(I)$.

$\widetilde{R}_I$ inherits a $\mathbb{N}[\sigma]$-grading from $R[t^{\mathbb{N}[\sigma]}]$, with R the homogeneous component of degree 0.

If $X = \mathrm{Spec}^\sigma R$, let

$$\widetilde{X}_I^c = \mathrm{Proj}^\sigma\,(\widetilde{R}_I^c)$$

and let $\pi : \widetilde{X}_I^c \to X$ be the natural map, corresponding to the inclusion $R \to \widetilde{R}_I^c$.

If X is any difference scheme, and $\mathcal{I}$ a quasi-coherent ideal presheaf of $\mathcal{O}_X$ (considered forgetfully as a sheaf of rings), one can check as above that the blowing ups of $\mathrm{Spec}^\sigma U$ at the ideals $\mathcal{I}(U)$, and their canonical maps to U, glue together to give a difference scheme over X, denoted $\widetilde{X}_{\mathcal{I}}^c$.

Definition 16.16 *Let X be a difference scheme, I a quasi-coherent ideal presheaf. Let $Y = \widetilde{X}_I^c$, $\pi : Y \to X$ the structure map. Then we obtain an ideal presheaf $\pi^* I$ on Y. We define the exceptional divisor to be the subscheme of Y defined locally by $\pi^* I$ (or equivalently by the difference ideal sheaf generated by $\pi^* I$.)*

Comparing the blowing-ups We include a result comparing the two blowing-ups as the finitely generated ideal I approaches J.

When J is already generated as a difference ideal by $I \cap \sigma^{-1}(I)$, the difference between the two blowing-ups is a proper closed subscheme Z of the exceptional divisor. It can be interpreted as follows. In either version of the blowing up, a point of the exceptional divisor corresponds to a "direction of approach" to the difference subscheme defined by J. However in the closed blowing up, directions in which $\sigma(y)/y$ approaches infinity as $y \to 0$ are allowed; in the open blowing up they are not.

If one blows up after applying M_q, I and J become identified, and one obtains only "directions of approach" in which $\sigma(y)/y$ approaches 0; thus these points are bounded away from Z.

There is more to be said here:

1. As $I \to J$, the closed blowing up $\widetilde{X}_I^c$ appears to change fairly gently. Perhaps the different $\widetilde{X}_I^c$ can be compared by transformally birational radicial morphisms. (cf. Definition 5.3.)
2. Outside Z, the closed blowing ups contains points representing directions in which $\sigma(y)/y$ is finite but nonzero; upon applying M_q, these points form a detachable divisor. (Is it possible to get rid of them before?)
3. On the other hand it may be interesting precisely to study blowing up one-generated difference ideals. Note that after M_q, such ideals become principal, so blowing them up has no effect on smooth varieties. For instance, it appears possible that one obtains a better definition of "irreducible difference variety" by demanding that the irreducibility persist to the strict transform under such principal blowing ups.

Proposition 16.17 *Let R be a difference ring, J a finitely generated difference ideal,*

V a set of generators of J, and I the ideal generated by $V \cup \sigma(V)$. Let $X = \mathrm{Spec}^\sigma R$. Then $\widetilde{X}_J$ is isomorphic to an open subscheme of $\widetilde{X}^c_I$. Specifically, let

$$Z(V) = \{p \in \mathrm{Proj}^\sigma\,(\widetilde{R}_I) : Vt \subset p\}$$

Then

$$\widetilde{X}_J = \widetilde{X}^c_I \setminus Z(V)$$

Proof Both difference schemes are obtained by gluing together affine schemes $\mathrm{Spec}^\sigma S$, where for some $a \in V$, S is a subring of $R[1/a, 1/\sigma(a), \ldots]$. Namely:

$$S = S_1 = \{\frac{c}{a^n} : n \in \mathbb{N}[\sigma], c \in I^n\}$$

according to the closed construction; for the other, one checks that

$$S = S_2 = \{\frac{c}{a^n} : n \in \mathbb{N}[\sigma], c \in J^n\}$$

Note that $\mathrm{Spec}^\sigma\,(S_1)$ (respectively $\mathrm{Spec}^\sigma\,(S_2)$) embeds naturally as an open subscheme of $\widetilde{X}^c_I$ (resp. $\widetilde{X}_J$). In the first case, by definition of the closed set $Z(V)$, the union of these open subschemes over $a \in V$ is precisely $\widetilde{X}^c_I \setminus Z(V)$. in the second case, it is $\widetilde{X}_J$, since V generates J as a difference ideal. We will now fix $a \in V$ and show that $S_1 = S_2$. The resulting isomorphisms of open subschemes are natural and glue together to give the required isomorphism.

Clearly $S_1 \subset S_2$. For the other direction, it suffices to check that $\frac{c}{a} \in S_1$ when $c \in J$. $\{c : \frac{c}{a} \in S_1$ forms an ideal of R. J is generated as an ideal by $\cup_k I^{\sigma^k}$, so we may take $c \in I^{\sigma^k}$ for some k, so that $\frac{c}{a^{\sigma^k}} \in S_1$. But also $a^\sigma \in \sigma(V) \subset I$, so $\frac{a^\sigma}{a} \in S_1$, hence by applying σ, $\frac{a^{\sigma^{i+1}}}{a^{\sigma^i}} \in S_1$ for $i < k$. Multiplying these $k+1$ elements, we obtain $\frac{c}{a} \in S_1$.

Blowing up and M_q-reduction

Lemma 16.18 *Let R be a difference ring, I an ideal. Let $\widetilde{R}^c_I$ be the closed blowing up ring. Let $S = M_q(R)$, and $\bar{I} = M_q(I) = (I + J_q(R))/J_q(R)$. Let $\widetilde{S}^c_{\bar{I}} = \sum_{n \in \mathbb{N}} (\bar{I}t)^n \subset S[t]$.*

There exists a natural surjective homomorphism of graded rings

$$j_q : M_q(\widetilde{R}^c_I) \to \widetilde{S}^c_{\bar{I}}$$

Proof Clearly $M_q(R[t^{\mathbb{N}[\sigma]}]) \cong S[t]$; by restriction, we get an isomorphism

$$\widetilde{R}^c_I\,/(J_q(R[t^{\mathbb{N}[\sigma]}]) \cap \widetilde{R}^c_I) \cong \widetilde{S}^c_{\bar{I}}$$

Now clearly $J_q(\widetilde{R}^c_I) \subset (J_q(R[t^{\mathbb{N}[\sigma]}]) \cap \widetilde{R}^c_I)$, yielding the surjective homomorphism

$$M_q(\widetilde{R}^c_I) = \widetilde{R}^c_I / J_q(\widetilde{R}^c_I) \to \widetilde{S}^c_{\bar{I}}$$

The naturality can be seen via the universal property of M_q. □

Lemma 16.19 *Let $X = \mathrm{Spec}^\sigma R$ be an affine difference scheme, I an ideal of R, $\bar{I} \subset M_q(R)$ the image of I under M_q. There exists a natural embedding of $\widetilde{M_q(X)}_{\bar{I}}$ as a closed subscheme of $M_q(\widetilde{X}^c_I)$.*

Moreover, if V is a sub-ideal of I, generating the same difference ideal, and $Z(V) = \{p \in \mathrm{Proj}^\sigma\,(\widetilde{R}^c_I) : Vt \subset p\}$, then the image of $\widetilde{M_q(X)}_{\bar{I}}$ is disjoint from $M_q(Z(V))$.

Proof From 16.18 we obtain a map

$$j_q^* : \operatorname{Proj} \widetilde{M_q(R)}_{\bar{I}}^{c} \to \operatorname{Proj} M_q(\widetilde{R}_I^c)$$

This can be composed with 16.14. The "moreover" is clear, since if $Vt \subset j_q^{-1}(p)$ then $j_q(Vt) \subset p$. However V and I generate the same difference ideal, so if $\bar{V}$ denotes the image of V in $M_q(R)$, then $\bar{V}$ and $\bar{I}$ generate the same ideal. Thus $\bar{V}t \not\subseteq p$ for a homogeneous ideal $p \in \widetilde{M_q(R)}_{\bar{I}}^{c}$. □

Corollary 16.20 *Let X be a difference scheme, I a quasi-coherent ideal presheaf. Let $\bar{I}$ be the M_q-image of I. There exists a natural embedding of $\widetilde{M_q(X)}_{\bar{I}}$ as a closed subscheme of $M_q(\widetilde{X}_I)$.*

Proof This reduces to the local case, 16.19, by gluing. □

Note 16.21

In 16.20, assume I is a sub-presheaf of a difference ideal sheaf J, and $I \cap \sigma(I)$ generates J. Then the image of $\widetilde{M_q(X)}_{\bar{I}}$ in $M_q(\widetilde{X}_I)$ is contained in the M_q-image of the open blowing up of X at J.

A geometric view of blowing up difference schemes When $X = \operatorname{Proj}(R)$ is itself a projective or multi-projective difference variety, blowing-ups of X have a natural multi-projective structure. Suppose R is graded by $\mathbb{N}[\sigma]^k$. Then $R[t^{\mathbb{N}[\sigma]}]$ is naturally graded by $\mathbb{N}[\sigma]^{k+1}$; this induces a grading of $\widetilde{R}_I^c$.

Let I be a homogeneous ideal of the $\mathbb{N}[\sigma]^k$-graded difference ring R. Let $\mathcal{I}$ be the corresponding ideal sheaf on $X = \operatorname{Proj}^{\mathbb{N}[\sigma]^k} R$. Then

$$\widetilde{X}_{\mathcal{I}} \cong \operatorname{Proj}^{\mathbb{N}[\sigma]^{k+1}} \widetilde{R}_I^c$$

naturally.

Let R be a difference ring; view it as a scheme $X = \operatorname{Spec} R$, endowed with an endomorphism. Let $I = I(0)$ be a finitely generated ideal of R. Let $B_1(R, I)$ be the affine coordinate ring of the blow-up of $\operatorname{Spec} R$ at I, in the sense of algebraic geometry ([Fulton]). Thus $B_1(R, I)$ may be viewed as a subring of $R[t]$:

$$B_1(R, I) = R + \sum_{i \geq 1} I^n t^n \subset R[t]$$

We can also view $B_1(R)$ as an extension of R. Let $I(1)$ be the ideal of $B_1(R, I)$ generated by $\sigma(I) \subset R \subset B_1(R, I)$. Proceeding inductively, let

$$B_{n+1}(R, I) = B_1(B_n(R, I), B_n(R, I)\sigma^{n+1} I)$$

$B_1(R, I)$ is naturally $\mathbb{Z}$-graded, with R the homogeneous component of degree 0. This grading inductively builds up to a $\mathbb{Z}^n$-grading on $B_n(R, I)$. We let

$$X_n = Proj^{\mathbb{Z}^n}(B_n(R, I))$$

On X_n we have the exceptional divisor E_n corresponding to the ideal $\sigma^n(I)B_n(R,I)$. Let $B_\infty(R,I)$ be the direct limit of the rings $B_n(R,I)$. Each of these rings is naturally embedded in $R[t,t^\sigma,\ldots]$, and $B_\infty(R,I)$ can be taken to be the union of the rings $B_n(R,I)$ within $R[t,t^\sigma,\ldots]$. While the individual $B_n(R,I)$ do not in general admit a difference ring structure, $B_\infty(R,I)$ is clearly a difference subring of $R[t,t^\sigma,\ldots]$, and we thus view it as a difference ring.

Lemma 16.22 *$\widetilde{R}_I^c \simeq B_\infty(R,I)$ as difference ring extensions of R.*

Proof They actually coincide as subrings of $R[t,t^\sigma,\ldots]$. □

$\widetilde{X}_I^c$ can thus be viewed as a limit of the projective system of schemes $X \leftarrow X_1 \leftarrow X_2 \leftarrow \ldots$. The exceptional divisor E on $\widetilde{X}_I^c$ corresponds to the intersection of the pullbacks of all the E_n. Observe that the image of E on X_n will usually have unbounded codimension.

Example 16.23

Let $R = [\sigma]k[X_1,\ldots,X_n]$ be the affine difference ring of affine n-space, with k a difference field. Then $\mathrm{Spec}^\sigma R$ is affine n-space over k. However $\mathrm{Spec}\, R$ is an infinite product of schemes Y_i, each isomorphic to affine n-space over k. Suppose I is an ideal of R generated by an ideal I_0 of $k[X_1,\ldots,X_n]$, corresponding to a subscheme S_0 of affine space, and let Z_i be the result of blowing up Y_i at I_0, EZ_n the exceptional divisor. Then X_n can be identified with

$$Z_1 \times Z_2 \times \ldots \times Z_n \times Y_{n+1} \times Y_{n+2} \times \ldots$$

□

In the above example, we used the following observation concerning ordinary blowing-up of varieties: when S,T are schemes, and C a subscheme of S, $(\widetilde{S \times T})_{C\times T}$ can be identified with $\widetilde{S}_C \times T$.

16.4 Semi-valuative difference rings

Valuative schemes of finite total dimension over $\breve{A}_k$ can be viewed as transformal analogs of smooth curves. But moving a finite scheme over $\mathbb{P}^1$, we encounter singular curves as well. Their ultraproducts lead to considering semi-valuative schemes. We take a look at these in the present subsection; it will not be required in the applications. Lemma 16.31 explains their potential role in a purely schematic treatment of transformal specializations.

An m-semi-valuative A_0-algebra A is by definition a local $k[\breve{t}]_\sigma$-algebra contained in a Boolean -valued valuation ring R over A_0, such that $\sigma^m(t)R \subset A$. We will be interested in this when $A_0 = k[\breve{t}]_\sigma$ or when $A_0 = k[T]$ in characteristic p. An ultraproduct of the latter will be an instance of the former. When R is a transformal valuation domain, we say A is an m-pinched transformal valuation ring.

Definition 16.24 *Let R be a transformal valuation domain, A a difference subring of R, $t \in R$. If $A \cap t^m R \subset A$, we will say that A is an m-pinched valuative domain (with respect to (R,t).)*

Let $X_0 = \mathrm{Spec}^\sigma\, R/tR$; we have a map $f : X_0 \to \mathrm{Spec}^\sigma\, A/tA$. If Y is a component of $\mathrm{Spec}^\sigma\, A/tA$, write $r.dim(X_0/Y)$ for the relative reduced total dimension of $f^{-1}(Y) \subset X_0$ over Y. The valuative multiplicity *of Y is $v.mult(Y) = rk_{ram}(K) + r.dim(X_0/Y)$. Finally, the valuative dimension of Y (viewed as a part of the specialization of A) is $vt.dim(Y) = r.dim(f^{-1}(Y)) + rk_{ram}(K)$.*

Example 16.25 *A plane curve with a point pinched to order m, inside formal ring of normalization.*

More precisely, the local ring of the curve has the required property, inside the local ring of the normalization, over $k[t]$ rather than $k[\check{t}]_\sigma$; an ultraproduct of such rings leads to $k[\check{t}]_\sigma$. We consider this class of examples in more detail.

Consider curves $C \subset \mathbb{P}^n$ over an algebraically closed field k, $P \in C$. C may be reducible and singular. (For simplicity, assume C is reduced.) The normalization $\widetilde{C}$ of C is defined to be the disjoint sum of the normalizations of the irreducible components of C. Let A A be the local ring of C at P, B the product of the local rings of $\widetilde{C}$ at the points above P.

Let $T \in A$ be a non-zero-divisor, corresponding to the restriction of a linear birational map $\mathbb{P}^n \to \mathbb{P}^1$.

Example 16.26 *Let $C \subset \mathbb{P}^n$ be a curve of degree $\underline{d}$. Let $P \in C$, $A, B, \widetilde{C}, T$ as above. Then $T^{\frac{3}{4}\underline{d}^2} B \subset A$.*

(For our purposes here, we could equally well replace A, B by their completions $\hat{A}, \hat{B}$ for the T-adic topology; i.e. the conclusion $T^{\frac{3}{4}\underline{d}^2} \hat{B} \subset \hat{A}$ is what we really need.)

Proof We may assume $P = 0 \in \mathbb{A}^n \subset \mathbb{P}^n$, T a linear map on $\mathbb{A}^n$. Let Y be a generic linear map on $\mathbb{A}^n$. Projecting C via (T, S) to $\mathbb{A}^2$ we obtain a plane curve of degree d, whose local ring $k[S, T] \cap A$ is if anything smaller that A, while the local ring of $\widetilde{C}$ does not change since the map (T, S) is birational on each component of C. So we can assume $C \subset \mathrm{Spec}\, k[S, T]$ is a plane curve.

Assume first that C is irreducible. By [Hartshorne] I ex. 7.2, the arithmetic genus $p_a(C)$ satisfies $p_a(C) = \binom{\underline{d}}{2} \leq \underline{d}^2/2$. By [Hartshorne] IV Ex. 1.8 (a), since the genus of the normalization $\widetilde{C}$ is non-negative, the A-module B/A has length $l \leq p_a(C)$. Thus $A + T^i B = A + T^{i+1} B$ for some $i < l$. In B/A , we have $T(T^i B/A) = (T^i B/A)$, so by Nakayama, $T^i B/A = 0$, and hence $T^l B \subset A$.

If C is reducible (at 0), it is a union of irreducible curves C_i, of degrees $\underline{d}_i$, with $\sum \underline{d}_i = \underline{d}$. More precisely, C is cut out by $\Pi_i f_i$, $f_i \in k[S, T]$ relatively prime polynomials of degrees $\underline{d}_i$; C_i is the curve cut out by f_i.

Let B_i be the local ring of the normalization $\widetilde{C}_i$ of C_i. Then $B = \Pi_i B_i$. Let A_i be the local ring of C_i itself. We can view A as a subring of $\Pi_i A_i$. We will show that $(T^{\underline{d}^2/4} \Pi_i A_i) \subset A$. By the irreducible case, $T^{\underline{d}^2/2} \Pi_i B_i \subset \Pi_i A_i$. So $T^{3\underline{d}^2/4} B \subset A$. It suffices therefore to show that $T^{\underline{d}^2/4} A_i \subset A$ for each i; i.e. that there exists $a_i \in A$ whose image in A_j is $T^{\underline{d}^2/4}$, and whose image in every other A_j is 0. Take for instance $i = 1$.

Let $f_1{}' = f_2 f_3 \cdot \ldots \cdot f_n$, $\underline{d}_1{}' = \deg(f_1{}')$. Note that $\underline{d}_1 + \underline{d}_1{}' = \underline{d}$, so $\underline{d}_1 \underline{d}_1{}' \leq \frac{1}{4}\underline{d}^2$. Now the curves $(f_1), (f_1{}')$ intersect in a subscheme of size $\leq \underline{d}_1 \underline{d}_1{}'$, by Bezout's theorem. So $T^{\underline{d}^2/4} = r f_1 + r' f_1'$, $r, r' \in k[S, T]$. The element $r' f_1'$ of A is the one we looked for: it equals

$T^{\underline{d}^2/4}$ module f_1, and equals 0 modulo f_j for $j \neq 1$. □

Here is a view of some of the combinatorics of the above example.

Lemma 16.27 *Let E be a sub-semi-group of $\mathbb{N}$, generated by a finite X with greatest element b. Then for some $Y \subset \mathbb{Z}/b\mathbb{Z}$, for all $n \geq b(b-1)+1$, $n \in E$ iff $n \operatorname{mod} b \in Y$.*

Proof Let $Y(i)$ be the image in $\mathbb{Z}/b\mathbb{Z}$ of $E \cap [ib+1, (i+1)b]$. Clearly $\emptyset \subset Y(0) \subseteq Y(1) \subseteq \ldots$, so for some $i < b$ we have $Y(k) = Y(k+1)$. Now $X + Y(k) \subset Y(k+1) (\operatorname{mod} b)$, so $X + Y(k) = Y(k) (\operatorname{mod} b)$, and it follows that $E + Y(k) = Y(k) (\operatorname{mod} b)$. The claim follows, with $Y = Y(k)$. □

Remark 16.28 *In 8.22, let A be an m-pinched valuative subring of R, with respect to t, $t \in A$. Then A/tA has total dimension $\leq d$ over F; if equality holds, then $tA \cap K[0] = (0)$ (in fact $tR \cap K[0] = (0)$.)*

Proof The inequality is immediate from 9.3; so we will concentrate on the case: $tA \cap K[0] \neq (0)$, and show that A/tA has total dimension $< d$. (But the argument also serves in general to show the weak inequality.)

Assume $tR \cap K[0] \neq (0)$. The proof of 8.20-8.22 shows that R/t_mR has total dimension $< m+d$. (In general, it has total dimension $\leq m+d$.)

Since $t_mR \subset A$, we have $(t_mR \cap A)^2 \subset t_mA$: if $t_mr_i \in A$, then $(t_mr_1)(t_mr_2) = t_m(t_mr_1r_2) \in t_m(t_mR \cap A)$. So the difference rings A/t_mA, $A/(t_mR \cap A)$ differ only by nilpotents, and thus have the same total dimension. Now $A/(t_mR \cap A)$ embeds into R/t_mR, and so is well-mixed, and has total dimension $< m+d$ (resp. $\leq m+d$.)

Let J_n be the smallest well-mixed algebraically radical ideal of A containing t_n. So $J_m = t_mR \cap A$, by the above argument; thus ${}^{A}\sqrt{J}_m = (\sqrt{J}_m) \cap A$ is an algebraically prime ideal.

Let A'' be a finitely generated F-subalgebra of A/tA. Lift the generators to A/t_mA, and let A' be the $F[\check{t}]_\sigma[m-1]$-subalgebra of A/t_mA generated by them. We have seen that A/t_mA has total dimension $< m+d$, hence so does A', and thus also $A'/\sqrt{(0)}_{A'} \leq A/J_m$. In this ring, t_m is not a 0-divisor. Thus the proof of 9.3 takes over, and shows that A'/tA' has total dimension $< d$. Hence so does A''. □

The usefulness of 16.28 is limited by the fact that it does not apply to an arbitrary weak component of A/tA. We can use valuation-theoretic rather than schematic multiplicities:

Lemma 16.29 *In 8.22, let A be an m-pinched valuative subring of R, with respect to t, $t \in A$. Let Y be a component of $\operatorname{Spec}^\sigma A/t$. Then*

$$v.mult(Y) + r.dim(Y) \leq d$$

If equality holds, then $P = 0$. If every weakly Zariski dense component of R/tR is Zariski dense, then Y is Zariski dense in $\operatorname{Spec} R[0]$.

Proof By the argument in 16.28, A/t_mA, $A/(t_mR \cap A)$ differ only by nilpotents; they thus have the same reduced total dimension. Similarly, the maps $A/t_mA \to A/tA$, $A/(t_mR \cap A) \to$

$A/(tR \cap A)$ are isomorphisms on points, and also respect the reduced total dimension. Thus $r.dim(A/tA) = r.dim(A/(tR \cap A)) \leq r.dim(R/tR)$; $r.dim(Y) + r.dim(X_0/Y) \leq r.dim(f^{-1}(Y))$. So $vt.dim(Y) \leq vt.dim(f^{-1}(Y)) \leq vt.dim(R/tR)$. The lemma follows from 8.22. □

By an m- *semi-valuative* difference scheme (over $\breve{A}_k$), we will mean one one of the form $\mathrm{Spec}^\sigma A$, where A is an m-semi-valuative $k[\breve{t}]_\sigma$-algebra. We will also say that the singularity of $\mathrm{Spec}^\sigma A$ is of order $\leq m$.

(We think of $\mathrm{Spec}^\sigma A$ as a disjoint union of generalized curves over $\breve{A}_k$, with a pinching of order m over $t = 0$.)

$X_{\to 0}{}'$ is the variant of $X_{\to 0}$ defined using semi-valuative difference schemes, rather than valuative ones. ($2d+1$ -semi-valuative difference schemes will do; any larger number will lead to substantially the same scheme. cf. 16.26.) $X_{\to 0}{}'$ has the same points as $X_{\to 0}$, but may be thicker at the edges. $X_{\to 0}{}'$ can probably also be defined as the intersection of the closed projections of of $(\hat{X})_0$ as $\hat{X}$ ranges over all open blow-ups of X, with vertical components removed.

Lemma 16.30 *Let Z be a weak component of $X_{\to 0}{}'$. Then there exists a $T = \mathrm{Spec}^\sigma \bar{A}$, $\bar{A}$ a pinched valuative domain over $k[\breve{t}]_\sigma$, and a morphism $T \to X$ over $\breve{A}_k$, such that Z is contained in the image of T_0 in X_0.*

Proof Same as the proof of 9.9, except that we consider $h_j : R \to B_j \subset A_j \subset L_j$, with $T^{\sigma^{2d+1}} A_j \subset B_j$. The condition is that p contains $\cap_{j \in J} {h_j}^{-1}(tB_j)$, and in the conclusion p contains $h^{-1}(t\bar{B})$. □

Lemma 16.31 *Let $D \subset k$ be a difference domain such that X descends to D ($X' \times_D k = X$, X' a difference scheme over D.) Let Y^0 (resp. Y) be a finitely presented well-mixed difference scheme with $X_{\to 0} \subset Y \times_D k$. Then for some $a_0 \in K$, with $F = D[1/a_0]$,*

(i) *For all difference fields L and all homomorphisms $h : F \to L$, if $X_h = X' \otimes_{F[\breve{t}]_\sigma} \breve{L}_t$, $(X_h)_{\to 0}{}' \subset Y_h$ (as schemes.)*

(ii) *For all sufficiently large q, all $h : F \to L = K_q$, if $\bar{X} = X_{\breve{h_t}}$, for any reduced (but possibly reducible) curve C over L and map $C \to \bar{X}$, $C \cap \bar{X}_0 \subset Y_h^0$ (as varieties.)*

(iii) *For all sufficiently large q, all $h : F \to L = K_q$, if $\bar{X} = X_{\breve{h_t}}$, for any reduced (but possibly reducible) curve C over L and map $C \to \bar{X}$, $C \cap \bar{X}_0 \subset Y_h$ (as schemes.)*

Proof The assumption that Y^0 is finitely presented over D means that locally, on $\mathrm{Spec}^\sigma A \subset X$, Y^0 is defined by a finitely generated ideal I (among well-mixed ideals.) Let $r_1, \ldots, r_j$ be generators for I. Then (i) and (ii) amount to showing that each r_i lies in a certain ideal. In (i), the condition is that r_i should vanish on the image $f(T_0)$ for any valuative T and $f : T \to (Y^0)_h$. In other words, given a map $g : A \to R$ with R an $2d+1$-semi-valuative ring over $\breve{L}_t$, where d is the total dimension of X_t, that $g(r_i) \in tR$. If (i) fails, there are F_ν approaching K and $h_\nu : F \to L_\nu$, and $g_\nu : A \to R_\nu$, with $g(r_i) \notin tR_\nu$. Taking an ultraproduct, we obtain $g_* : S \to R_*$ with $g(r_i) \notin tR_*$. But this contradicts the definition of $X_{\to 0}{}'$, and the assumption $X_{\to 0}{}' \subset Y^0$.

As for (ii) and (iii), let A be the product of the local rings at 0 of C, and let R be the corresponding product of the local rings of the the normalizations of the components of C. Each such local ring is a discrete valuation domain, and can be viewed as a transformal valuation domain using the Frobenius map $x \mapsto x^q$. By Lemma 16.26, if $\underline{d}$ is a projective degree of C, then $t^{\frac{3}{4}\underline{d}^2} R \subset A$. Now by 11.22, $\underline{d} \leq O(1)q^d$. So $\underline{d}^2 \leq O(1)^2 q^{2d} \leq q^{2d+1}$ (for large enough q.) so $t^{\sigma^{2d+1}} R \subset A$. Continue as in (i).

□

References

[Ax] J. Ax., The elementary theory of finite fields, *Annals of Math.* 88 (1968), pp. 239-271

[Bombieri] Enrico Bombieri, Thompson's Problem ($\sigma^2 = 3$), *Invent. Math.* **58** (1980), no. 1., pp. 77-100

[Chatzidakis-Hrushovski] Z. Chatzidakis, E. Hrushovski, Model theory of difference fields, Trans. Amer. Math. Soc. 351 (1999), pp. 2997-3071

[Chatzidakis-Hrushovski-Peterzil] , Model Theory of difference fields II: the virtual structure and the trichotomy in all characteristics, Proc. London. Math. Soc. (3) 85 (2002) pp. 257-311

[Cohn] Richard M. Cohn, *Difference Algebra*, Interscience tracts in pure and applied mathematics 17, Wiley and Sons New York - London - Sydney 1965

[Cohn79] Cohn, R.M., *The Greenspan bound*, Proc. AMS 79 (1980), pp. 523-526

[Cohn87] Cohn, R.M., *Order and Dimension*, Proc. AMS 87 (1983), pp. 1-6.

[deJong] A.J. de Jong, Smoothness, semi-stability and alterations, preprint obtained from `ftp://ftp.math.harvard.edu/pub/AJdeJong/alterations.dvi`

[Deligne 74] P. Deligne, La conjecture de Weil, I., *Publ. Math. Inst. Hautes Etude. Sci. Publ. Math. 43* (1974) pp. 273-307

[Deligne 81] P. Deligne, La conjecture de Weil, II., *Publ. Math. Inst. Hautes Etude. Sci. Publ. Math. 52* (1981) pp. 313-428

[Deligne 77] P. Deligne, SGA $4\frac{1}{2}$, Springer-Verlag Lecture notes in mathematics 569, Berlin-Heidelberg 1977

[Fujiwara97] Fujiwara, Kazuhiro, Rigid geometry, Lefschetz-Verdier trace formula and Deligne's conjecture. *Invent. Math.* 127 (1997), no. 3, 489–533

[Fulton] Fulton, W., *Intersection Theory*, Springer-Verlag Berlin-Tokyo 1984

[Freitag-Kiehl] E. Freitag and R. Kiehl, *Étale Cohomology and the Weil Conjecture*, Springer-Verlag, Berlin-Heidelberg 1988

[Fried-Jarden86] Fried, M and Jarden, M, Field Arithmetic, Springer-Verlag, Berlin 1986

[Gromov] M. Gromov, Endomorphisms of symbolic algebraic varieties, *J. Eur. Math. Soc. (JEMS)* 1 (1999), no. 2, 109–197

[Hartshorne] Robin Hartshorne, *Algebraic Geometry*, Springer-Verlag New York - Berlin 1977

[Haskell-Hrushovski-Macpherson] D. Haskell, E. Hrushovski, D. Macpherson, *Definable sets in algebraically closed valued fields*, Part I (preprint.)

[Haskell-Hrushovski-Macpherson] *Definable sets in algebraically closed valued fields*, Part II, preprint

[Herwig-Hrushovski-Macpherson] Bernhard Herwig, Ehud Hrushovski, Dugald Macpherson, *Interpretable groups, stably embedded sets, and Vaughtian pairs J. London Math. Soc.* (2) 68 (2003) pp. 1-11

[Hyot] William L. Hoyt, On the moving lemma for rational equivalence, J. of the Indian Math. Soc. 35 (1971) p. 47-66

[Hrushovski01] Hrushovski, E., The Manin-Mumford Conjecture and the Model Theory of Difference Fields, Annals of Pure and Applied Logic 112 (2001) 43-115

[Hrushovski02] Hrushovski, E., Computing the Galois group of a linear differential equation, in *Differential Galois Theory*, Banach Center Publications 58, Institute of Mathematics, Polish Academy of Sciences, Warszawa 2002

[Hrushovski-Pillay] Hrushovski, E. and Pillay, A., Definable subgroups of algebraic groups over finite fields, Crelle's journal 462 (1995), 6991; Groups definable in local fields and pseudo-finite fields, Israel J. of Math., 85 (1994) pp. 203-262

[Jacobi] Jacobi, C.G.J *De investigando ordine systematis aequationum differentialium vulgarium cujuscunque,* (Ex. Ill. C. G. J. Jocbi maunscriptis posthumis in medium protulit C. W. Borchardt), Journal für die riene und angewandte Mathematik, Bd. 64, p. 297-320

[Kleiman68] S.L. Kleiman, Algebraic cycles and the Weil conjectures, in *Dix exposes sur la cohomologie des schemas*, A. Grothendieck and N. H. Kuiper (eds.), Mason et Cie., Paris, North Holland, Amsterdam 1968

[KMP] Kondrat'eva, M. V.; Mikhalv, A. V.; Pankrat'ev, E. V. *Jacobi's bound for systems of differential polynomials.* (Russian) Algebra, 79–85, Moskov. Gos. Univ., Moscow, 1982.

[Lando] Lando, B.A., Jacobi's bound for first order difference equations, Proc. AMS **52** (1972), pp. 8-12

[Lang59] Serge Lang, Abelian varieties, Springer -Verlag New-York Tokyo 1983

[Lang-Weil] Lang, S. and A. Weil, Number of points of varieties in finite fields, Amer. Jour. Math., **76** (1954), 17-20

[Lipshitz-Saracino73] Lipshitz L., Saracino D., The model companion of the theory of commutative rings without nilpotent elements, Proc. Amer. Math. Soc. 38, 1973, 381-387.

[Macintyre95] A. Macintyre, Generic automorphisms of fields, *Annals of Pure and Applied Logic* **88** Nr 2-3 (1997), 165 – 180

[Macintyre1973] Macintyre A.J., Model-completeness for sheaves of structures, Fund. Math. 81, 1973/74, no. 1, 73-89.

[Milne1980] J.S. Milne, *Étale Cohomology*, Princeton University Press, Princeton, NJ 1980

[Pink92] Richard Pink, On the calculation of local terms in the Lefschetz-Verdier trace formula and its applications to a conjecture of Deligne, *Annals of Mathematics* **135** (1992) pp. 483-525

[Ritt] Ritt, J. F., *Jacobi's problem on the order of a system of differential equations,* Ann. of Math., (2) 36 (1935), 303-312

[Scanlon] Quantifier elimination for the relative Frobenius, in *Valuation Theory and Its Applications Volume II* , conference proceedings of the International Conference on Valuation Theory (Saskatoon, 1999), Franz-Viktor Kuhlmann, Salma Kuhlmann, and Murray Marshall, eds., Fields Institute Communications Series, (AMS, Providence), 2003, 323 - 352

[Seidenberg1980] A. Seidenberg, editor, *Studies in Algebraic Geometry*, Studies in Mathematics, vol. 20, Mathematical Association of America 1980

[Serre] J. -P. Serre, Local fields, Springer Verlag 1979.

[shoenfield] Shoenfield, J. R., *Mathematical Logic.* Reprint of the 1973 second printing. Association for Symbolic Logic, Urbana, IL; A K Peters, Ltd., Natick, MA, 2001 ; Shoenfield, J. R., Unramified forcing, *in Axiomatic Set Theory*, Proc. Sympos. Pure Math., Vol. XIII, Part I, Univ. California, Los Angeles, Calif., 1967, pp. 357–381 ; Amer. Math. Soc., Providence, R.I 1971

[Weil48] André Weil, *Variétés abeliennes et courbes algébriques*, Hermann, Paris, 1948