THE LIMIT THEORY OF THE FROBENIUS BOX 7: CO-ANALYSIS AND INERTIAL DIMENSION

by

James Freitag, William Johnson, Alice Medvedev & Thomas Scanlon

Abstract. —

Résumé **(La théorie limite du Frobenius, Boîte 7: Cöanalyse)**

1. Scattered

Definition 1.1. — Let K be a valued field. A subset $X \subset K$ is said to be *scattered* if $\{|x-y| : x, y \in K\}$ is finite. $X \subset K^n$ is scattered if the projection $\pi_i(X) \subset K$ is scattered for each i.

Remark 1.2. — In point-set topology, a subset X of a Hausdorff topological space is *scattered* if every subset of X has at least one isolated point, or equivalently, no subset of X is dense-in-itself. If a set $X \subset K^n$ is scattered in the above sense, then it is topologically discrete, thus certainly scattered in the topological sense. The above definition seems to be a much stronger condition than being topologically scattered, so the terminology might have a different origin.

Lemma 1.3. — *Assume that X is scattered. Then there are a finite number of equivalence relations $E_0 \subset E_1 \subset \ldots \subset E_n$ on X so that $E_0 = Id$, E_n has a single equivalence class, and for each E_{i+1} class Y, there is a map f_i^Y embedding Y/E_i into the residue field. E_i and f_i^Y are quantifier free definable in the language of valued fields and f_i^Y is defined by a formula with a parameter which depends upon Y. Each of the E_i are defined over any set of parameters over which X is defined.*

Proof. — Let $X \subset K^d$. We proceed by induction on d. First consider the base case where $X \subseteq K^1$. Let $\rho_0 < \rho_1 < \ldots \rho_n$ be the possible values of $|x-y|$ for $x, y \in X$. Let $E_i(x, y)$ hold if $|x-y| \leq \rho_i$. Then $E_0 \subset E_1 \subset \cdots \subset E_n$ has the desired properties. Given an E_i-equivalence class Y, we let $f_{i-1}^Y(x) = res(c_i^{-1}(x-b))$, where b is any

Partially supported by NSF grants FRG DMS-0854998 and DMS-1001550 and an NSF MSPRF.

element of Y and c_i any element of K with absolute value ρ_i. Then f_i^Y induces an injection from Y/E_{i-1} into the residue field.

Now suppose $X \subset K^{d+1}$. The projection of X to the first d coordinates will still be scattered. Applying the inductive hypothesis to this projection, we get a decomposition $E'_0 \subset E'_1 \subset \cdots \subset E'_n$. Applying the base case to $\pi_{d+1}(X)$, we also get a sequence of equivalence relations $F'_0 \subset F'_1 \subset \cdots \subset F'_n$ on $\pi_{d+1}(X)$. For $i \leq n$, let $E_i(x, y)$ assert that the first d coordinates of and x and y satisfy E'_i while $x_{d+1} = y_{d+1}$. For $i \geq n$, let E_i assert that $F'_{i-n}(x_{d+1}, y_{d+1})$. We leave the construction of the f_i^Y as an exercise to the reader.

□

Lemma 1.4. — *Let K be a transformal valued field with Γ_K a torsion-free $\mathbb{Z}[\sigma]$-module. If L is an elementary extension of K and $0 \neq a \in L$ is transformally algebraic over K (i.e., $K(a)_\sigma$ has finite transcendence degree over K), then $v(a)$ is in $\mathbb{Q}(\sigma)\Gamma_K$.*

Proof. — Since a has finite total dimension over K, it is a root of some nonzero transformal polynomial $F(x) = \sum_\nu c_\nu x^\nu$, with $c_\nu \in K$. Because $\sum_\nu c_\nu a^\nu = 0$, there must exist $\mu \neq \nu$ such that $v(c_\nu a^\nu) = v(c_\mu a^\mu) \neq \infty$. Then $(\nu - \mu)v(a) = v(c_\nu) - v(c_\mu)$, so

$$v(a) = \frac{v(c_\nu) - v(c_\mu)}{\nu - \mu} \in \mathbb{Q}(\sigma)\Gamma_K.$$

□

Proposition 1.5. — *Let K be a transformal valued field with Γ_K a torsion-free $\mathbb{Z}[\sigma]$-module. Then any definable set $X \subset K^n$ of finite total dimension is scattered.*

Proof. — It suffices to show that for $1 \leq i \leq n$, the definable set

$$S_i = \{|x_i - y_i| : x, y \in X(L)\}$$

is bounded in elementary extensions L of K, because this will imply that it is finite for $L = K$, which is the definition of scattered. By the assumption that X has finite total dimension, the x_i, y_i, and $x_i - y_i$ are transformally algebraic over K. By the previous lemma, every element of S_i is in $\{\infty\} \cup \mathbb{Q}(\sigma)\Gamma_K$, which is bounded. □

Let us give a simple example. When $X \subset \mathbb{A}^1$ is given by $\sigma x - x = 0$, then for any solution $v(a) = 0$. Any difference of solutions is a solution, so $val(a - b) = 0$ when a, b are two distinct solutions. Consequently, $res(a) \neq res(b)$. Therefore the residue map gives an injective map from X into $\mathbb{A}^1$. In the proposition, we can take $E_0 \subset E_1$, with E_0 equality and $E_1 = X \times X$.

2. Co-Analysis

Given a structure M and a distinguished sort or definable set V, there is a traditional notion of what it means for a definable set D to be co-analyzable over V. D is co-analyzable over V in zero steps if it is a singleton. D is co-analyzable in $n+1$ steps if there is a definable map $f : D \to V^m$ for some m, such that the fibers of f

are co-analyzable over V in n steps. The parameters used to define f can come from anywhere. This notion is a slight variant of the notion of analyzability, in which one instead requires the fibers of the map to be internal and the base to be analyzable.

In Section 7.2, Hrushovski defines a notion of "co-analyzability" which differs from the traditional version in a few ways. First, he restricts to quantifier free sets. Second, he considers definable families of sets. Roughly speaking, a predicate $P(x;y)$ is going to be co-analyzable in the sense if Hrushovski if $P(L;b)$ is co-analyzable in the traditional sense for every model L and tuple b from L. Finally, he introduces an extra piece of data into the definition, a notion of "inertial dimension," which keeps track of the degrees of freedom from the residue sort that were used in the co-analysis of a set.

Let M be $F(t)^a_\sigma$ for some trivially valued inversive difference field F, with the usual ω-increasing valuation on M such that $v(t) > 1$. Let $\mathcal{L}$ be the language of transformal valued rings over M with a distinguished sort $V = V_{res}$ for the residue field.Let T be the theory of ω-increasing transformal valued fields extending M. We emphasize in what follows that u and v will always denote variables from the residue field sort, or for more general discussion, from the sort V. Variables x, y, z will come from either the home sort (the valued field) or V (the residue sort).

Elements of M are named by constants. In what follows, we will never extend the language by adding constants for elements outside of M.

Let $\Phi(x;y)$ be the set of quantifier-free $\mathcal{L}$-formulas such that if $\phi(x;y) \in \Phi$ then $T \models \phi(x;v) \to \phi_1(x) \wedge \phi_2(v)$ for some quantifier-free ACFA formulas ϕ_1 and ϕ_2 of finite total dimension over M (or M_{res}). Equivalently, a quantifier-free $\mathcal{L}$-formula $\phi(x;y)$ is in $\Phi(x;y)$ when it implies that its arguments are transformally algebraic over M or M_{res}, in the sense that for any $L \models T$ and a, b from L such that $L \models \phi(a,b)$, we have a and b transformally algebraic over M or M_{res}.

By Φ_{fn}, we denote a set of *basic functions*; in the setting of transformal valued fields, this will be the set of functions built out of composition from σ-polynomials (on both the residue field sort and the field sort, with coefficients from M_{res} or M, respectively) and maps of the form:

$$(x, y_1, y_2) \mapsto res\left(\frac{x - y_1}{y_1 - y_2}\right),$$

where $y_1 \neq y_2$ and $v(x - y_1) = v(x - y_2)$.

Note that elements of Φ_{fn} and $\Phi(x;y)$ can name parameters from M, but are not allowed to name parameters from bigger models.

We will write $d_V(\phi(v;y)) \leq n$ if for any $L \models T$ and a, b from L, $\phi(a;b)$ implies $t.dim(a/(M(b))_{res}) \leq n$. In other words, $\phi(\cdot;b)$ defines a difference scheme over $M(b)_{(res)}$ with total dimension at most n, for any b in an extension of M. By definition of $\Phi(v;y)$, if $\phi(v;y) \in \Phi(v;y)$, then $d_v(\phi(v;y)) < \infty$. The quantity $d_V(\phi)$ is a measure of the size of the fibers of the relation cut out by ϕ.

We are about to define what it means for $\phi(x;y)$ in $\Phi(x;y)$ to be "co-analyzable" of inertial dimension at most n. This should be thought of as a generalization of d_V to the case where the first argument no longer lives in the residue sort. Specifically, the inertial dimension of $\phi(x;y)$ is a bound on the dimension of the fibers $\phi(L;b)$ in models $L \models T$, in some sense of the word "dimension." The inertial dimension of $\phi(x; y)$ tells us nothing about how much y can vary.

Definition 2.1. — We define the notion of (h, n)-co-analyzability for $P \in \Phi$ by induction, and will sometimes alternatively write co-analyzable in h steps of inertial dimension at most n.

1. $P(x; y)$ is co-analyzable in 0 steps of inertial dimension less than or equal to n if for every $N \models T$ and $b \in N$, $|P(N; b)| \leq 1$.
2. $P(x; y)$ is co-analyzable in $3h + 1$ steps of inertial dimension less than or equal to n if there are
   - $n_1, n_2 \in \mathbb{N}$ such that $n_1 + n_2 \leq n$
   - $Q(x; y, v) \in \Phi(x; y, v)$ which is co-analyzable in at most $3h$ steps of inertial dimension less than or equal to n_1
   - $R(v; y) \in \Phi(v; y)$ with $d_V(R(v; y)) \leq n_2$
   - $g(x; y) = v$, a basic definable function (from Φ_{fn})

   such that for all $N \models T$ and $b \in N$,
   $$P(N; b) \subseteq \{a \in N \mid N \models Q(a; b, g(a, b)) \wedge R(g(a, b); b)\}.$$
   Note that the function g is a map to V^m for some m.
3. $P(x; y)$ is co-analysable in $3h+2$ steps of inertial dimension less than n if there are $Q_1(z, y), \ldots, Q_l(z, y) \in \Phi(z, y)$ such that for $j = 1, \ldots, l$,
   $$\phi_j(x; z, y) := P(x, y) \wedge Q_j(y, z) \in \Phi(x; z, y)$$
   is co-analysable in at most $3h + 1$ steps of inertial dimension less than n and such that
   $$T \vdash P(x, y) \to \bigvee_{i=1}^{l} \exists z : Q_i(z; y),$$
   i.e., for all models $N \models T$ and $b \in N$, if $P(N, b) \neq \emptyset$, then there is $c \in N$ and some j such that $Q_j(c; b)$.
4. $P(x; y)$ is co-analyzable in $3h + 3$ steps of inertial dimension less than or equal to n if there are $P_i(x; y) \in \Phi(x; y)$ for $1 \leq i \leq l$ which are co-analyzable in at most $3h + 2$ steps of inertial dimension less than or equal to n such that $T \vdash P(x; y) \to \bigvee_{i=1}^{l} P_i(x; y)$.

We say that $P(x; y)$ is co-analyzable with inertial dimension $\leq n$ if it is (h, n)-coanalyzable for some h.

The number of steps in a co-analysis has no significance and is merely a bookkeeping device for managing inductive proofs.

Remark 2.2. — Note that if $T \vdash P(x; y) \to Q(x; y)$, and $Q(x; y)$ has inertial dimension at most n, then so does $P(x; y)$. This can be seen by using step $3h + 1$. For large enough h, $Q(x; y)$ will be $(3h, n)$-co-analyzable. Taking $R(; y)$ to be a finite-total dimension formula implied by $Q(x; y)$ and letting $g(x, y)$ be the unique function with range V^0, we have
$$T \vdash P(x; y) \implies Q(x; y, g(x, y)) \wedge R(g(x, y); y),$$

so that $P(x; y)$ is $(3h + 1, n)$-co-analyzable.

Remark 2.3. — By virtue of step $3h + 3$, the union of two predicates of inertial dimension n has inertial dimension n.

Remark 2.4. — There are a couple differences between our definition and Hrushovski's original definition of co-analyzability. First, we have allowed the second parameter of a co-analyzable formula $P(x; y)$ to range over things other than V. This seemed in the spirit of the name "co-analyzable." Moreover, without making this change, all co-analyzable sets (fibers of co-analyzable relations) would be internal to the residue sort, which was probably not the desired intent.

We have also added a third step to allow for finite unions of co-analyzable sets. Explicitly, we have forced it to be true that the union of two sets of inertial dimension at most n has inertial dimension at most n. This property is used by Hrushovski in the subsequent proofs, and didn't seem to be implied by the original definition.

In the most general setting, step $3h + 3$ cannot be obtained from the other two steps. Suppose we omit step $3h + 3$, and only use Hrushovski's original two steps. Suppose Φ_{fn} only contains constant functions. One can easily verify that in this setting, only singletons are co-analyzable. Specifically, if $P(x; y)$ is co-analyzable, then in any model $N \models T$, $|P(N; b)| \leq 1$ for $b \in N$. On the other hand, the union of two such relations will probably not have this property, showing that a union of two co-analyzable relations need not be co-analyzable. To avoid this problem, we have added in step $3h + 3$ explicitly.

Definition 2.5. — Let L be a model of T and let D be a quantifier-free definable set in L. We say that D is co-analyzable of inertial dimension at most n if there is some predicate $P(x; y)$ co-analyzable in some number of steps of inertial dimension at most n and some b in L such that D is $P(L; b)$.

Definition 2.6. — Let L be a model of T and let K be a substructure of L, such that the home sort and residue sort of K are both fields, but the residue map need not be surjective. Let $a \in L$. Then we define $\dim_V(a/K) \leq n$ if there is some predicate $P(x; y)$ co-analyzable with inertial dimension at most n and some $b \in K$ such that $L \models P(a; b)$.

Remark 2.7. — This doesn't depend on the choice of the ambient model L, i.e., it would not change if we replaced L by a larger model. This follows from our choice to include only quantifier-free formulae in $\Phi(x; y)$. Is this the only place we used the quantifier-free assumption?

Definition 2.8. — Let X and Y be $\emptyset$-definable sets. Then X is *internal* to Y if there is a definable surjective map $f : Y^m \to X$ for some m. Sometimes we ask for an injective map $X \to Y^m$ rather than a surjective map $Y^m \to X$. This is slightly stricter in general, but equivalent when Y eliminates imaginaries. The definition of co-analyzability seems to be following this stricter definition, with g playing the role of the injective map.

Part of the point of the previous definition of internality is that even though X and Y are B-definable (adding constants, we may assume $B = \emptyset$), any map f which "witnesses internality" might require additional parameters.

Remark 2.9. — Let us analyze the definition of co-analyzability in several simple instances. First, assume that $P(x; y)$ is $(1, n)$-co-analyzable. By hypothesis, $P(x; y) \Rightarrow Q(x; y, g(x, y))$, where $Q(x; y, u)$ is $(0, n)$-co-analyzable. The $(0, n)$-co-analyzability of $Q(x; y, u)$ means that x is uniquely determined by y and u. Consequently, $g(x, b)$ must be an injective map from each fiber $P(x; b)$ into the residue sort. Further, we track the dimension of the fibers via R. In the case of $(1, n)$-co-analyzability, n_1 as in the definition is 0, so $n = n_2$, the dimension of $g(x, y)$ over y. So, we can see that the image of $g(x, b)$ as x varies in the fiber above b is bounded by the inertial dimension n. Thus, $(1, n)$-co-analyzability corresponds to uniform internality to V with the injective map on the fiber above b requiring no parameters except for b.

Now suppose that $P(x; y)$ is $(2, n)$-co-analyzable. Then by definition, there are finitely many $Q_j \in \Phi(z; y)$ so that for each j, $P(x; y) \wedge Q_j(z; y) \in \Phi(x; yz)$ is $(1, n)$-co-analyzable. The z's here serve as the extra parameters needed to define the maps witnessing internality, making $(2, n)$-co-analyzability the same thing as internality (of the fibers). There is no requirement on Q_j besides that it is finite total dimension. So, in the notion of co-analysis, the inertial dimension is essentially measuring the sum of the dimensions of the fibrations we are taking; we don't care about the additional parameters, except to say that we can find them in some finite total dimension difference varieties.

Using the analysis in Section 1, one can show that scattered sets are co-analyzable. So all sets of finite total dimension are co-analyzable.

Example 2.10. — Consider the system of difference equations given by

$$x^\sigma - x = t$$
$$ty^\sigma = xy$$

Let $P_3(x, y; -) \in \Phi(x, y; -)$ be the formula asserting x and y satisfy the above difference equations. One can show that P_3 is $(4, 2)$-co-analyzable.

Let K be a model of T (i.e., an extension of M) and let c be an element from some model of T extending K. Let $Rk(c/K)$ be a $\mathbb{N} \cup \{\infty\}$-valued function depending only on the quantifier-free Φ-type of c over K. Specifically, in the case of ω-increasing transformal valued fields, let $Rk(c/K)$ be $rk_{val}(K(c)/K)$. In general, we will assume that $Rk(c/K)$ depends only on $K(c)$, so that $Rk(c_1/K) = Rk(c_2/K)$ whenever c_1 and c_2 generate the same Φ_{fn} substructure over K. We will write $K(c)$ for the Φ_{fn}-structure generated by c over K. We will write $Rk(L/K)$ for $Rk(c/K)$ if $L = K(c)$. By assumption, this is well-defined. In the case we care about, $Rk(L/K)$ is $rk_{val}(L/K)$.

We remind the reader of our conventions regarding the residue sort: variables v and u always come from the residue sort. For the remainder of this section, we will assume that Rk has the next three properties; from these, we will prove that rank is a bound for inertial dimension.

The first is an initial bound on the dimension of tuples in the V-sort of an extension of Φ_{fn}-structures of bounded rank:

IB : If $Rk(L/K) \leq n$ and $a \in V(L)$, then for some $\phi(v;y) \in \Phi(v;y)$ with $d_V(\phi) \leq n$ and $b \in K$, $L \models \phi(a,b)$.

Remark 2.11. — In the case of differential fields, take $Rk(L/K)$ ot be the supremum of the transcendence degree of tuples $\bar{a} \in C(L)$ over $C(K)$. Of course, this property holds almost by definition.

Remark 2.12. — In the case of transformal valued fields, this is saying that if $rk_{val}(L/K) \leq n$, then every element of $res(L)$ has transcendence degree at most n over $res(K)$, which is clear from the definition.

FD : If $K \leq K' \leq L$, then $Rk(L/K') + Rk(K'/K) = Rk(L/K)$.

Remark 2.13. — In the differential case, this again holds essentially by definition.

Remark 2.14. — In the case of transformal valued fields, this is clear.

M : If $K \leq K' \leq N \models T$ and $c \in N$, then $Rk(K'(c)/K') \leq Rk(K(c)/K)$.

The three properties given here are sufficient to ensure that any such abstract rank function Rk which possesses them will serve as an upper bound for inertial dimension.

Remark 2.15. — In the differential case, this again holds essentially by definition. This might not be a particularly interesting case to consider if all of the properties follow so easily, but we will see.

In the case of transformal valued fields, this follows from the very technical result Proposition 6.35 of section 6.6... except only in the setting where all the fields have transformal dimension 0 over the base field M. This won't actually hold, meaning that we may need to make a slight adjustment to the setup. In particular, it may make more sense to work in the theory having a sort for each set defined over M of finite total dimension. This would be similar to the technique of studying compact complex manifolds model-theoretically by having one sort for each compact complex manifold.

As an extra complication, strictly speaking we can't apply Proposition 6.35 to get **M** because the K and L appearing in the above three axioms are not actually transformal valued fields (models of T), but instead models of $T_\forall$. We can more or less pretend [1] that $T_\forall$ implies the two sorts are both fields, but $T_\forall$ is missing the requirement that the residue map be surjective. We are currently looking into ways around this problem.

Remark 2.16. — From these three properties **IB**, **FD**, and **M**, we will draw conclusions about dim_V and Rk. Logically speaking, this happens on an abstract level, for any rank satisfying the properties of Rk.

1. This would hold if we had added the symbol $\div$ to the language.

Lemma 2.17. — *Let $K \leq L \models T_\forall$. Assume $Rk(L/K) \leq n$. Let $a \in L$ so that a/K is co-analyzable.*[2] *Then $dim_V(a/K) \leq n$.*

Proof. — The assumption that a/K is co-analyzable means that $\dim_V(a/K) < \infty$, so that there is some co-analyzable predicate $P(x;y) \in \Phi(x;y)$ and some $b \in K$ such that $L \models P(a;b)$. The proof proceeds via induction on the number of steps in the co-analysis of $P(x;y)$.

Suppose that $P(x;y)$ is co-analyzable in 0 steps. Then $P(x;y)$ has inertial dimension zero, so $\dim_V(a/K)$ is zero, hence certainly less than whatever n. So there is nothing to show in this case.

Suppose that $P(x;y)$ is co-analyzable in $3h+1$ steps. Then there is some $Q(x;y,v)$ which is co-analyzable in $3h$ steps with $g \in \Phi_{fn}(x,y;v)$ and $d = g(a,b)$ so that $L \models Q(a;b,d)$. Now we are going to apply the induction hypothesis to $a/K(d)$; note that d is in the V-sort. Forget about Q; we only used it to find d.

By induction, we know that there is some $Q_1 \in \Phi(x;y_1,v)$ and $b_1 \in K$ so that $\mathcal{M} \models Q_1(a;b_1,d)$, with $dim_V(Q_1) \leq Rk(L/K(d))$. Now, by **IB**, there is some $R_1(v;y_2) \in \Phi(v;y_2)$ so that $d_V(R_1) \leq Rk(K(d)/K)$ and some $b_2 \in K$ with $R_1(d,b_2)$.

Now, consider $P_1(x;y,y_1,y_2) := Q_1(x;y_1,g(x,y)) \wedge R_1(g(x,y);y_2)$. Note that we know $L \models P(a;b,b_1,b_2)$ and this formula shows that

$$dim_V(a/K) \leq dim_V(Q_1) + d_V(R_1) \leq Rk(L/K(d)) + Rk(K(d)/K).$$

By **FD**, the right ride of the inequality is $Rk(L/K)$.

Now, consider the case that $P(x;y)$ is co-analyzable in $3h+2$ steps. In this case, we know $L \models P(a,b)$ with $b \in K$ and $T \models P(x;y) \Rightarrow \bigvee_{j=1}^{m} \exists z Q_j(z;y)$ and for each j, $P(x;y) \wedge Q_j(z;y) \in \Phi(x;z,y)$ is co-analyzable in $3h+1$ steps. Set $r := Rk(L/K)$ and let

$$\Phi_r := \{\phi \in \Phi(x;y,y',z) \,|\, dim_V(\phi) \leq r\}.$$

We remind the reader that $b \in K$ was fixed above. Let

$$\Phi'_r = \{\neg\phi(x;b,b',z) \,|\, \phi \in \Phi_r,\, b' \in K\}.$$

Claim. — We will show that for any $j = 1, \ldots, m$,

$$T_\forall \cup tp_\Phi(a/K) \cup \Phi'_r \cup \{Q_j(z,b)\}$$

is inconsistent.

Proof. — Suppose not. Then there is some $N \models T_\forall$ with $K \leq N$ such that there are

(1) $$a_1 \in N \text{ with } a_1 \models tp_\Phi(a/K)$$

(2) $$d \in N \text{ with } N \models Q_j(d,b)$$

(3) $$\text{for any } \phi \in \Phi_r,\ N \models \neg\phi(a,b,b',d).$$

Let $K' := K(d)$ be the Φ_{fn} substructure generated by d over K; similarly, set $L' := K(a',d) = K'(a')$. By [**M**], $Rk(L'/K') \leq r$. By the fact that $P(x;y) \wedge Q_j(z;y)$

2. This means that there is some predicate $P(x;y)$ which is co-analyzable, and some $b \in K$ such that $L \models P(a,b)$.

is co-analyzable in $3h+1$ steps, a'/K' is co-analyzable in $3h+1$ steps. So, applying the induction hypothesis, $dim_V(a'/K') \leq Rk(L'/K') \leq r$. So, there is some $\phi(x; b, b', d) \in tp_\Phi(a'/K')$ with $dim_V(\phi) \leq r$. Now, we can see that $N \models \phi(a', b, b', d)$. This contradicts (3). □

Now we know that for each $j \in \{1, \ldots, m\}$,

$$T_\forall \cup tp_\Phi(a/K) \cup \Phi'_r \cup \{Q_j(z, b)\}$$

is inconsistent.

By compactness, there is some $P'_j(x; b, b') \in tp_\Phi(a/K)$ and some finite disjunction $\bigvee_{j'} \phi_{jj'}(x; y, y', z)$ of elements of Φ_r such that

$$T_\forall \vdash P'_j(x; y, y') \wedge Q_j(z; y) \rightarrow \bigvee_{j'} \phi_{jj'}(x; y, y', z).$$

A priori, $P'_j(x; b, b')$ could depend on j, but replacing it with $\bigwedge_j P'_j(x; b, b')$, we may assume that it does not. Replacing it with $P'(x; b, b') \cap P(x; b)$, we may assume that $P'(x; y, y') \rightarrow P(x; y)$.

Because the $\phi_{jj'}$ have inertial dimension at most r, so does their disjunction $\bigvee_{j'} \phi_{jj'}(x; y, y', z)$. Consequently, $P'(x; y, y') \wedge Q_j(z; y)$, as an element of $\Phi(x; y, y', z)$ has inertial dimension at most r. Now, since

$$T \vdash P'(x; y, y') \rightarrow P(x; y) \rightarrow \bigvee_j (\exists z) Q_j(z, y)$$

it follows that $P'(x; y, y')$ is co-analyzable in some $3k+2$ steps with inertial dimension at most r. Since $P(a; b, b')$ holds, it follows that $\dim_V(a/K) \leq r$ as claimed.

Finally, consider the case where $P(x; y)$ is co-analyzable in $3h+3$ steps. Then $P(x; y)$ implies some finite disjunction $\bigvee_i P_i(x; y)$. Since $P(a; b)$ holds, $P_i(a; b)$ must hold for some i. Applying the inductive hypothesis to $P_i(x; y)$, we are done. □

Again, we remind the reader that the assumption that $Rk(-)$ satisfies the above three properties (**IB, FD, M**) is in effect.

Proposition 2.18. — *Let $P \in \Phi(x; y)$ be co-analyzable. Assume that if $K \leq \mathcal{M} \models T$, $a \in \mathcal{M}$, and $b \in K$ with $P(a, b)$. Then $dim_V(P) \leq Rk(K(a)/K)$.*

Proof. — The result follows by compactness and the previous two lemmas. (We again use the fact that the inertial dimension of a finite union is the maximum of the inertial dimension of the things being unioned.) □

Finally, when we specialize to the case of ω-increasing transformal valued fields:

Proposition 2.19. — *Let $\mathcal{M}$ be an algebraically closed ω-increasing transformal valued field of transformal dimension one over an inversive difference field. Assume the value group of $\mathcal{M}$ is $\mathbb{Q}_\sigma$. Let $\phi \in \Phi(x)$ be a quantifier-free formula in the language of transformal valued fields over $\mathcal{M}$. Assume that $\phi(x)$ is V_{res}-analyzable and that for any ω-increasing transformal valued field extension $L = M(c)$ with $\phi(c)$, $rk_{val}(L/\mathcal{M})$. Then $dim_V(phi(x)) \leq n$.*

Proof. — Once we note that the three properties (**IB, FD, M**) are satisfied by $Rk(-)$, this proposition is a special case of 2.18. So, let K' be a finitely generated extension of K'' over K of transformal dimension zero.

Property **IB** follows since $tr.deg.(K'_{res}/K''_{res} \leq rk_{val}(K'/K'')$. Property **FD** follows because vector space dimension is additive in extensions. **M** follows from 6.35. □

Proposition 2.20. — *Suppose that $P(x;y)$ is (h,n)-co-analyzable with the theory T, the set of formulas Φ, and the set of basic functions Φ_{fn} fixed.*

1. *For some finite $T_0 \subset T$, $\Phi_0 \subset \Phi$, $\Phi_{fn,0} \subset \Phi_{fn}$, P has inertial dimension less than or equal to n with respect to $T_0, \Phi_0, \Phi_{fn,0}$.*
2. *Let $\mathcal{M}_q \models T_q$, where $q \in S \subset \mathbb{N}$ with S infinite, be a family of models of T_0. Suppose that for any $P(v;y) \in \Phi_0$ (note that the first variable is in the residue sort), there is β so that for all q and all $b \in \mathcal{M}_q$,*

$$|P(\mathcal{M}_q, b)| \leq \beta q^{d_V(P)}.$$

Then for any $Q(x;y) \in \Phi_0$ with $dim_V(Q) \leq n$ with respect to $T_0, \Phi_0, \Phi_{fn,0}$, there is some β_1 such that for all q and for all $b \in M_q$,

$$|P(\mathcal{M}_q, b)| \leq \beta_1 q^n.$$

Proof. — (1) follows by book keeping during the co-analysis. (2) uses induction on the number of steps. At $3h+2$ and $3h+3$ steps, we may increase β. At $3h+1$ steps, the exponent increases, and we use the assumption to count the number of points in fibers, since they live in the residue sort. □

Remark 2.21. — The situation in which we will be interested will be when the theory T_0 is that of k-increasing transformal valued fields (for some large enough k). Then we will take $\mathcal{M}_q := K_q(t)^{alg}$ with some nontrivial valuation and with the q-Frobenius. So, the previous proposition takes us from counting points in the residue sort to counting points in definable sets co-analyzable in the residue sort.

James Freitag & William Johnson & Alice Medvedev & Thomas Scanlon
E-mail : `freitag@math.berkeley.edu` • *E-mail :* `willij@math.berkeley.edu`
E-mail : `alice@math.berkeley.edu` • *E-mail :* `scanlon@math.berkeley.edu`, Department of Mathematics, University of California, Berkeley, Evans Hall, Berkeley, CA 94720-3840, USA
Url : `http://www.math.berkeley.edu/~scanlon`