Notions of difference closures of difference fields

Zoé Chatzidakis
CNRS (UMR 8553) - ENS Paris

SLALM XVIII
Concepcion, 16 - 20 December 2019

Partially supported by the Institut Français du Chili

Setting

Two examples of fields with operators: *differential fields* and *difference fields*.
(All fields are commutative)

A differential field is a field K with an operator D which is additive and satisfies Cauchy's law: $D(ab) = aDb + bDa$.

A difference field is a field K with a distinguished endomorphism σ. Usually, one imposes to σ to be an automorphism.

The algebra of these fields was developed in parallel by Ritt in the 1930's.

Examples: The field $\mathbb{C}(t)$, with the usual derivation d/dt.

The field $\mathbb{C}(t)$, with σ defined as the identity on $\mathbb{C}$ and sending t to $t+1$ (and so $f(t)$ to $f(t+1)$).

Differentially closed fields (of characteristic 0)

The model companion (i.e., a model complete theory axiomatising the class of existentially closed differential fields) of the theory of differential fields of characteristic 0, DCF_0, was first described by A. Robinson, then later studied by L. Blum. They prove to be very interesting model-theoretically: the theory is complete, eliminates quantifiers, eliminates imaginaries, and moreover is ω-stable. The model companion $\mathrm{DCF}_{0,n}$ of fields of characteristic 0 with n commuting derivations enjoys similar properties.

In positive characteristic $p > 0$, the model companion DCF_p of the theory of differential fields of characteristic p also exists, and was studied by C. Wood. It is stable (not superstable), and eliminates quantifiers in the language differential fields augmented by a function symbol for the partial map $x \mapsto x^{1/p}$ (defined on the subfield of constants, 0 outside).

Typical quantifier-free formulas:
Differential fields: language $= \{+, -, \cdot, 0, 1, D\}$; atomic formulas are of the form $P(x_1, \ldots, x_n, Dx_1, \ldots, Dx_n, \ldots, D^m x_1, \ldots) = 0$, where P is a polynomial over $\mathbb{Z}$. If one allows parameters in K, then P has coefficients in K. *Differential equations*.
Difference fields: language $= \{+, -, \cdot, 0, 1, \sigma\}$; atomic formulas are of the form $P(x_1, \ldots, x_n, \sigma x_1, \ldots, \sigma x_n, \ldots, \sigma^m x_1, \ldots) = 0$, where P is as above. *Difference equations*.
Existentially closed: K is e.c. if every finite system of differential/difference equations with coefficients in K which has a solution in an extension of K, already has a solution in K.

Differential closure

Since DCF_0 is ω-stable, results of Shelah imply that if K is a differential field of characteristic 0, then the theory DCF_0 has a prime model over K, and which is unique up to K-isomorphism. This prime model is called the *differential closure of K*. While its existence and uniqueness are fairly easy to show algebraically when K is countable, model-theoretic arguments are needed in the uncountable case.
Recall: a prime model of a (complete) theory T is a model M which elementarily embeds into every model of T. If the theory T eliminates quantifiers, then one can also talk about prime models of T over a subset A (of some large model), since the theory $T \cup \text{qf-Diag}(A)$ is complete.

Difference closed

The model companion of the theory of difference fields exists, is denoted by ACFA. It is not complete, and does not eliminate quantifiers. Its completions are given by describing the isomorphism type of the difference subfield which is the algebraic closure of the prime field, i.e., $\bar{\mathbb{Q}}$ or $\bar{\mathbb{F}}_p$. Definable sets are well understood.

ACFA eliminates imaginaries. Its completions are not stable, but are supersimple. This means that any model has a good notion of independence, and possesses a rank with good properties. The field structure of a model is algebraically closed. But the fixed subfield of a model $\mathcal{U}$, i.e., $\{a \in \mathcal{U} \mid \sigma(a) = a\}$, is not algebraically closed, however it is pseudo-finite.

Difference closure?

In analogy with the differential case, one can call difference fields which are models of ACFA *difference closed*. Then the natural question is:
Do difference fields have a difference closure, and it is unique (up to isomorphism)?

In other words: does the theory ACFA admit prime models over difference fields?

The question has an obvious negative answer, if the underlying field is not algebraically closed: for instance, there are $2^{\aleph_0}$ incompatible ways of defining an automorphism of $\bar{\mathbb{Q}}$, so there can be no prime model of ACFA over $\mathbb{Q}$.
So the first condition we need to impose is that the difference field K be algebraically closed (as a field).

Second obstacle

As we saw before, the fixed subfield F of a difference-closed field $\mathcal{U}$ is *pseudo-finite*, and its theory has the independence property. It is therefore not surprising that:

if the fixed subfield of K is not pseudo-finite, then K does not have a difference-closure.

Indeed, a difference-closed field $\mathcal{U}$ containing K will contain an element $a \notin K$ with $\sigma(a) = a$. We'll do the case where K is algebraically closed, so that a is transcendental over K. But there are $2^{\aleph_0}$ non-isomorphic over K difference fields $(K(a)^{alg}, \sigma)$ containing $K(a)$ and with fixed field pseudo-finite.

New question

So, in order not to run into these trivial obstacles, we need to make further assumptions:

We assume that the difference field K is algebraically closed and that its fixed field F is pseudo-finite.

Is it enough to guarantee that the field K has a difference-closure?

NO . . . but the examples are harder.

Some comments about the examples

So, one needs to find formulas over K which contain no isolated types.
In characteristic 0, there is a particular example, using the j-function (a modular function); there are more examples in the same vein, but no infinite family is known.
In positive characteristic $p > 0$, there are many families of examples.

However these examples do not answer the question of which difference fields admit a difference closure: Are there some which are not already difference closed?

Recall that a *type* (in the tuple of variables x) over a subset A of a model M of T is a set $\Sigma(x)$ of formulas with parameters in A, which is finitely satisfiable in M; we will usually work with *complete* types, i.e., such that given a formula $\varphi(x)$, either $\Sigma(x) \vdash \varphi(x)$ or $\Sigma(x) \vdash \neg\varphi(x)$.
Example: $A \subset M \models T$, a an n tuple in M, x an n-tuple of variables. Then the *type of a over A*, $tp(a/A)$, is the set of formulas with parameters in A satisfied by a in M.

If κ is an infinite cardinal, a model M of T is κ*-saturated* if every type over a subset A of M of size $< \kappa$ is realised in M.

Other notions of prime models

($T = T^{eq}$) A model M of a theory T is $\aleph_\varepsilon$-*saturated* if whenever $A \subset M$ is finite, then every type over $\mathrm{acl}(A)$ is realised in M.

Let κ be an infinite cardinal or $\aleph_\varepsilon$, $B \subset M$. One says that M is κ-*prime over* B if M is κ-saturated, contains B, and B-embeds (elementarily) into every κ-saturated model of T containing B.

Shelah showed that if T is superstable and complete, and κ is either $> |T|$ or $\aleph_\varepsilon$ if $|T| = \aleph_0$, then κ-prime models exist and are unique up to isomorphism.
This solves the problem of isolated types not being dense in a non totally transcendental theory.

These notions have a direct translation into algebraic terms. For instance, with $\kappa = \aleph_1$: A model $\mathcal{U}$ is $\aleph_1$-saturated if and only if every countable set of difference equations (with coefficients in $\mathcal{U}$) which has a solution in a difference field extending $\mathcal{U}$, already has a solution in $\mathcal{U}$.
The notion of κ-prime then corresponds to a natural notion of closure for this property.

The result

Theorem. *Let K be an algebraically closed difference field of characteristic* 0, *κ an uncountable cardinal or $\aleph_\varepsilon$, and assume that the fixed field F of K is pseudo-finite and κ-saturated. Then K has a κ-difference closure $\mathcal{U}$, and it is unique up to isomorphism over K.*

Comments

Does not work in characteristic $p > 0$. In fact, I am pretty sure that κ-difference closure only exist when K is already κ-difference closed.

Why does it work in characteristic 0?

Because of the dichotomy stable/fixed field - more later on that.

Some ingredients of the proof

(1) The first thing we notice is the following:

Let K be a difference field, with fixed field F pseudo-finite and κ-saturated. Then there is a κ-saturated model $\mathcal{U}$ (of ACFA) containing K and with fixed field F.

The fixed field is also *stably embedded*, and its induced structure (by the ambiant $\mathcal{U}$) is the pure field structure.
The proof then follows the usual strategy, and to do that, we need to take a closer look at 1-types.

(2) The generic 1-type: it says that the element does not satisfy any non-trivial difference equation. This type is stationary.
In algebraic terms, it is called a transformally transcendental element, and gives rise to a notion of transformal transcendence basis. No problem: a κ-saturated model must have a transformal transcendence basis of cardinality $\geq \kappa$, and if needed, we realise the type κ many times.

(3) The non-generic types have finite rank and are analysable in terms of types of rank 1. What makes this work in characteristic 0, is that types of rank 1 are either
stable stably embedded (and locally modular);
or almost internal to the fixed field.

More precisely: given a and $K = \mathrm{acl}(K)$, with $\mathrm{SU}(a/K) < \omega$, there are $a_1, \ldots, a_n \in \mathrm{acl}(Ka)$, with $a \in \mathrm{acl}(Ka_1, \ldots, a_n)$, and for each i, $tp(a_i/\mathrm{acl}(Ka_1, \ldots, a_{i-1}))$ is either stable stably embedded, or internal to $\mathrm{Fix}(\sigma)$.

(3 - contd) The stable stably embedded types of rank 1 behave very much like usual strongly minimal types - notions of dimension, etc.

(4) The "internal to $\mathrm{Fix}(\sigma)$" part is more delicate, since those types are very much unstable. What one can show is that given such a type $tp(a/L)$, there is a finite subset A of L such that $tp(a/\mathrm{acl}(A)) \vdash tp(a/L)$. This is what makes things work. One can then redo Shelah's proof.

κ-prime models

Under the assumptions of the theorem, show that κ-prime models exist, and are characterized in the usual fashion:
every 1-type over K realised in the model is κ-isolated (it is implied by its restriction to some subset of size $< \kappa$); and every sequence of K-indiscernibles has length $\leq \kappa$.

Then show that any two κ-saturated models satisfying these two properties are isomorphic over K.

Thank you