MODEL THEORY OF DIFFERENCE FIELDS

MOSHE KAMENSKY

Lecture 1, Aug. 28, 2009

1. Introduction

The purpose of this course is to learn the fundamental results about the model theory of difference fields, as developed in Chatzidakis and Hrushovski [10]. The goal is to get to a level that will allow reading that paper. In this section we give some basic definitions, and mention some of the applications of this theory. A more detailed description of the content of the course can be found in 1.6.

1.1. A *difference field* is a pair (K, σ) where K is a field, and σ is an automorphism of K.

We will study these fields as structures for the language of rings, with an extra symbol σ for the automorphism. Thus, an atomic formula has the form $p(x, \sigma(x), \dots, \sigma^k(x)) = 0$, where x is a tuple of variables, and p is a polynomial (over some base difference field).

The subset of K defined by $\sigma(x) = x$ is a subfield, called the *fixed field*.

1.1.1. *Example.* Any field can be considered a difference field with the identity automorphism.

1.1.2. *Example.* Any (algebraic) automorphism of $\mathbb{P}^1_{\mathbb{C}}$ (or $\mathbb{A}^1_{\mathbb{C}}$) yields an automorphism of the function field $\mathbb{C}(t)$, hence an example of a difference field. For instance, the shift by 1 gives the automorphism $\sigma(f(t)) = f(t+1)$. With this choice, an equation of the form $\sigma(x) - x = p(x)$ becomes a classical difference equation (i.e., an discrete analogue of a differential equation), hence the name.

1.1.3. *Example.* The algebraic closure of any finite field has a canonical structure of a difference field, given by the Frobenius.

We now briefly discuss some applications of the general theory. A more detailed discussion of some of these might follow later.

1.2. **The Manin–Mumford conjecture (Hrushovski [20]).** Let $\mathbf{G}$ be a commutative connected algebraic group (for example, a torus G_m^k). We view $\mathbf{G}$ as scheme or as a definable set, in the sense that we may ask about the points $\mathbf{G}(K)$ in a field K. Nevertheless, the subgroup of n-torsion points (at least when n is prime to the characteristic) is finite (this is obvious in the case of a torus, and easy in general), and therefore, at least in characteristic 0, the subgroup $T(\mathbf{G})$ consisting of all torsion points makes sense without reference to a particular field (as long as it is big enough).

Let $\mathbf{X}$ be a subvariety of $\mathbf{G}$. How many torsion points can $\mathbf{X}$ contain? One way it can contain many of them is when there is a subgroup $\mathbf{H}$ of $\mathbf{G}$ contained in $\mathbf{X}$. In this case, X will contain $T(\mathbf{H})$. More generally, $\mathbf{X}$ may contain the translate of $\mathbf{H}$ by a torsion point. It turns out that (over number fields) these are the only options:

1.2.1. *Theorem.* (Hrushovski [20, Theorem 1.1.2]) Let $\mathbf{X}$ be a subvariety of $\mathbf{G}$, both over a number field K. Then there are finitely many subgroups $\mathbf{G_i}$ and torsion points c_i such that $\mathbf{X}$ contains each $c_i + \mathbf{G_i}$, and $T(\mathbf{G}) \cap \mathbf{X} = \bigcup_i T(\mathbf{G_i}) + c_i$

The theorem was previously proved by Raynaud for Abelian varieties, and extended by McQuillan to the case of semi-Abelian varieties. In the case that $\mathbf{X}$ is a curve, the theorem says that the number of torsion points on $\mathbf{X}$ is finite unless $\mathbf{X}$ itself is a group. This is the Manin–Mumford conjecture. The model theoretic proof uses the theory of difference fields, and provides in addition effective bounds and uniformity in all the data.

1.2.2. We now want to explain how the theory of difference fields enters into the proof. The difficulty with this kind of statements lies in the fact that while the group and the sub-variety are geometric objects, the sub-group of torsion points is not. The idea of the proof is to replace $T(\mathbf{G})$ by a difference equation, which is again a geometric object, though in a different theory. We first explain how to do this for the subgroup T'_p consisting of prime-to-p torsion (for suitably chosen prime number p).

We note that if $\sigma \in \mathrm{Aut}({}^{K^a}\!/_{K})$, where K is a field over which $\mathbf{G}$ is defined, then σ induces a group automorphism on $\mathbf{G}(K^a)$. More generally, given any polynomial $F \in \mathbb{Z}[T]$, $F(\sigma)$ is a well defined endomorphism of $\mathbf{G}(K^a)$ (For example, if $\mathbf{G} = \mathbf{G}_m$ and $F(T) = 2T^3 - 5$, then $F(T)(x) = \sigma^3(x)^2 x^{-5}$.) In particular, the kernel of $F(\sigma)$ is a subgroup of $\mathbf{G}(K^a)$ (definable in the language of difference fields, but not algebraic!)

1.2.3. *Lemma.* (Hrushovski [20, Lemma 5.0.10]) For K a number field, there is a prime number p, an automorphism $\sigma \in \mathrm{Aut}({}^{K^a}\!/_{K})$, and an integer non-zero polynomial F, such that the kernel of $F(\sigma)$ on $\mathbf{G}(K^a)$ contains $T'_p(\mathbf{G})$.

Proof. Let p be a prime of good reduction[1] Then $K \subseteq K_p$, the p-adic completion of K, and $T'_p(\mathbf{G}) \subseteq \mathbf{G}(\mathrm{L})$, where L is the maximal unramified extension of K_p (for the case of a torus, if a torsion element extends the value group by an element x, then x itself is torsion, which is impossible). Therefore, it is enough to prove the claim with K replaced by K_p and K^a by L. We let $\Bbbk$ be the residue field of K_p, and set σ to be any lifting of the Frobenius on $\Bbbk^a$ to L.

Claim. *The residue map induced an injective map on the level of T'_p.*

Proof. For n prime to p, the algebraic map $[n] : \mathbf{G}_\Bbbk \to \mathbf{G}_\Bbbk$ (multiplication by n) is étale, since it induces (usual) multiplication by n on the tangent space at identity. If an n torsion point of $\mathbf{G}_\Bbbk$ is the image of more than one torsion point over K, then $[n]$ is ramified at that point. □

The claim implies that it is enough to find an F as in the lemma with K replaced by a finite field, and σ is the Frobenius. The structure theory of algebraic groups gives an exact sequence $0 \to \mathbf{L} \to \mathbf{G} \to \mathbf{A} \to 0$, with $\mathbf{L}$ an affine group and

[1] The precise definition of good reduction is slightly delicate. For our purposes, it can be defined as a group whose torsion points lie in the maximal unramified extension. This equivalence is shown in Serre and Tate [36]. It follows from the usual definition that $\mathbf{G}$ has good reduction at almost every prime. For a torus it is easy to see that any prime will work.

$\mathbf{A}$ an Abelian variety. Given polynomials F_A and F_L for $\mathbf{A}$ and $\mathbf{L}$ respectively, $F = F_L \cdot F_A$ satisfies the requirement for $\mathbf{G}$. For Abelian varieties, the statement is a result of Weil [40]. For the affine part, we again get a decomposition into a torus and a unipotent part. When $\mathbf{L}$ is a torus, $F(T) = T - q$, where q is the size of K (this is slightly more complicated when the torus is non-split. We ignore this). The unipotent case is an exercise. □

1.2.4. *Remark.* Taking $K = \mathbb{Q}$ and $\mathbf{G} = \mathbf{G}_m$, we see that, for example, the odd torsion points all satisfy $\sigma(x) = x^2$, where σ is a lifting of the 2-Frobenius. Explicitly, this says that if n is odd prime, then $\mathbb{Q}^a$ and F_2^a have "the same" n-th roots of unity.

1.2.5. This result is the starting point for the proof of 1.2.1. The equation $F(\sigma)(x) = 0$ determines a definable set in a suitable theory of difference field, in which it has a finite "dimension". The main results of Chatzidakis and Hrushovski [10] is the study of such definable sets. It is shown that any finite-dimensional set can be decomposed into 1-dimensional sets (in algebraic terms, this is achieved by algebraic change of variable). Then the one-dimensional sets are analysed, and are shown to be of one of three kinds (the "trichotomy"). The theorem follows from a general structure theorem for definable subsets in one of these types. End lecture 1

1.3. **Algebraic dynamics (Medvedev and Scanlon [28]).** An (algebraic) *dynamical system* on a variety $\mathbf{X}$ is simply an (algebraic) morphism $F : \mathbf{X} \to \mathbf{X}$. The dynamical questions are questions about the effects of iterating F on subvarieties and (rational) points; e.g., for which points x are there natural numbers n and m, such that $F^n(x) = F^{n+m}(x)$ (pre-periodic points)? Likewise, for which subvarieties $\mathbf{V}$ of $\mathbf{X}$ are there n and m such that $F^{n+m}(\mathbf{V}) \subseteq F^n(\mathbf{V})$ (I don't know if they are called pre-periodic subvarieties).

For example, if $\mathbf{X}$ is an algebraic group, and $d > 1$ an integer, the map $x \mapsto x^d$ determines a dynamical system. The pre-periodic points are precisely the torsion points, and any sub-group is preserved. Therefore, the Manin–Mumford conjecture implies that if $\mathbf{X}$ is commutative, then the intersection of the pre-periodic points with any subvariety $\mathbf{Y}$ is contained in a pre-periodic subvariety of $\mathbf{Y}$.

In general, given a dynamical system F, one may study it by associating with it the difference equation $\sigma(x) = F(x)$ (this is precisely the equation that captures the prime to d torsion in the above example). One then uses the trichotomy to study the properties of the system.

Medvedev and Scanlon [28] do this in the case that $\mathbf{X}$ is a power of $\mathbb{P}^1$, and $F(x) = (f_1(x_1), \dots, f_n(x_n))$, where each f_i is a rational function on $\mathbb{P}^1$. We note that though the system itself is a direct product, questions such as above do not reduce to the components. However, any periodic subvariety determines a relation on the associated difference equation. The trichotomy implies that the equation decomposes as (essentially) a product of the three types of sets, and there can be no definable relations between them. Thus the question of periodic subvarieties does reduce to each of these parts separately. Within those parts, one was described explicitly in Chatzidakis and Hrushovski [10], and another in Medvedev [27]. The third kind can by definition have only binary non-trivial relations.

1.4. **The asymptotic theory of the Frobenius.** Given an infinite collection of structures (in the same language), its *asymptotic theory* is the set of sentences in that language that hold in all but finitely many of the structure. For example, it is a theorem of Ax (Ax [1]) that the asymptotic theory of the finite fields is the theory of pseudo-finite fields (of characteristic 0). The fixed field of a difference closed field is pseudo-finite, so it is not unreasonable to expect that $ACFA$, the theory of difference closed fields, is the asymptotic theory of the algebraic closures of the finite fields, endowed with the Frobenius.

Hrushovski [19] proved that this is, indeed, the case. Explicitly, it means that any statement that holds in a difference field also holds in all but finitely many of the $\mathbb{F}_p^a$, where the automorphism is the Frobenius. He used this result to prove a generalised version of the Lang–Weil estimates. Poonen also used this result to show that if $F : \mathbf{X} \to \mathbf{X}$ is a dominant map, with $\mathbf{X}$ an (irreducible) algebraic variety over the algebraic closure of a finite field, then the set of periodic points is dense in $\mathbf{X}$.

1.5. **Galois theory of difference equations (defaut).** The Galois theory of linear differential equations is parallel to the usual Galois theory, with finite groups replaced by affine algebraic groups. The construction of the Galois group can be done algebraically (in several ways), but if one views the differential equation as a definable set in a suitable theory of differential fields, this becomes an instance of a general theory of definable automorphism groups in model theory. Hrushovski [17] uses this description to show that the Galois group of a linear differential equation over $\mathbb{Q}(t)$ is always computable.

The theory of linear difference equations $\sigma(x) = Ax$ is parallel to that of differential equations, but the algebraic construction runs into difficulties (which can be overcome) and is valid only over an algebraically closed fixed field. However, viewing the equation as a definable set in $ACFA$, this Galois group again becomes an instance of the general theory (**defaut**). In particular, the group is defined independently of the fixed field, and the methods of Hrushovski [17] are applicable.

1.6. **Plan of the course.** For most of the course, we will concentrate purely on the model theory of difference fields. We will be following Chatzidakis and Hrushovski [10]. They cover the following topics.

1.6.1. *Section 1.*

- Existence and geometric axiomatisation of the model companion of the theory of difference fields ($ACFA$)
- Properties of the fixed field
- Description of the completions of $ACFA$
- Description of definable sets and types in terms of quantifiers
- Algebraic closure of a set
- Elimination of imaginaries and the independence theorem

1.6.2. *Section 2.* Ranks on definable sets and on types

1.6.3. *Section 3.* The concepts of orthogonality, forking, triviality, modularity, etc. in the context of $ACFA$ (and more generally).

1.6.4. *Section 4.* Dichotomy theorem for types of rank 1 (in characteristic 0; this was generalised to arbitrary characteristic in Chatzidakis et al. [12]). Such a type is either a finite cover of the field (with parameters), or is stable, stably-embedded and modular (and more explicitly, either trivial or “group-like”).

1.6.5. *Section 5.* Analysis of finite rank types in term of those of rank 1.

1.6.6. *Section 6.* Examples: description of the results for particular formulas, and counterexamples showing that the results are optimal.

1.6.7. *Section 7.* Groups of finite rank.

Of these, the course covers in detail section 1, in the more general context of a stable theory with a generic automorphism, as developed in Chatzidakis and Pillay [11] and in Hrushovski [18]. This is done in sections 2 through 5. Then sections 2–5 of the paper are covered in much less detail, and the main results are only stated (sections 6 and 7).

In addition to the paper itself, a good reference for this theory is Chatzidakis [9], which mostly overlaps with this course, though from a more algebraic point of view.

2. The model companion of the theory of difference fields

Lecture 2, Aug. 31, 2009

2.1. **Model companions.** We first work in the context of a general first order theory $\mathcal{T}$ (not necessarily complete). Whenever we are given a map of structures $f : M \to N$, we identify an element $m \in M$ with its image $f(m) \in N$. The content of this sub-section can be found in any text on basic model theory, for example Pillay [33] or Marker [25].

Recall that a *definable set* is an equivalence class of formulas with respect to logical equivalence (relatively to $\mathcal{T}$). A formula is *existential* if it has the form $\exists y \phi(x,y)$, where ϕ is quantifier-free (x and y are tuples). It is *universal* if its negation is existential. A definable set if quantifier-free, existential, etc. if it can be represented by a formula with the same property.

Let $f : M \to N$ be a map of structures. By definition, this means that if $\phi(\bar{x})$ is a quantifier-free formula, then $\phi(M)$ consists of the M-points of $\phi(N)$. This need not be the case for arbitrary formulas: first, in $\phi(N)$ might acquire some new elements of M, and second, some elements of $\phi(M)$ might no longer be in $\phi(N)$ (and both can even happen with the same formula ϕ).

2.1.1. *Exercise.* Show that the first happens if and only if the second does (for some other formula). Give an example of this (it is easy to find an example with T complete).

2.1.2. *Elementary embeddings.* Call a formula $\phi(\bar{x})$ f-*good* if $\phi(M) = \phi(N) \cap M$. The map f is called an *elementary embedding* if any formula is f-good.

2.1.3. *Exercise.*

1. Show that the set of f-good formulas forms a boolean sub-algebra.
2. Show that if ϕ is both universal and existential, then it is f-good for all f.
3. Give an example of a definable set that is both universal and existential, but is not quantifier free.
4. If $f : M \to N$ is an elementary embedding, show that M and N have the same theory.

2.1.4. *The Tarski–Vaught test.* $f : M \to N$ is an elementary embedding if and only if any M-definable non-empty subset of N (in one variable) has an M-point.

Proof. Since f-good formulas form a boolean algebra, it is enough to prove that any formula of the form $\psi(\bar{x}) = \exists y \phi(\bar{x}, y)$ is f-good. Assume that $\psi(\bar{m})$ holds in M. Then there is $n \in M$ such that $\phi(\bar{m}, n)$ holds in M. By induction, this also holds in N, hence so does $\psi(\bar{m})$. On the other hand, if $\psi(\bar{m})$ holds in N (but $\bar{m}$ is from M), then $\phi(\bar{m}, y)$ is a non-empty subset of N, and therefore has an M-point n. Again by induction $\phi(\bar{m}, n)$ holds in M, hence so does $\psi(\bar{m})$.

The converse is trivial. □

2.1.5. *Extension of parameters.* Given a theory $\mathcal{T}$ and a subset A of a model of $\mathcal{T}$, we denote by $\mathcal{T}_A$ the theory $\mathcal{T}$ expanded by constant symbols A and by axioms corresponding to all quantifier-free relations on elements of A. Thus, a model of $\mathcal{T}_A$ is the same as a model M of $\mathcal{T}$ together with an embedding of A in M. (In contrast, we denote by $\mathcal{T}(M)$ the theory of the model M, and so $\mathcal{T}_M(M)$ is the theory of M in the language that has a symbol for each of its elements.)

2.1.6. *Exercise.* Let $f : M \to N$ be a map between models of $\mathcal{T}_A$, which is elementary as a map of models of $\mathcal{T}$. Show that it is also elementary with respect to $\mathcal{T}_A$.

2.1.7. *Existentially closed models.* A model M is called *existentially closed* if any quantifier-free formula over M that has points in a model extending M, also has points in M.

2.1.8. *Model completeness.* The following conditions on a theory $\mathcal{T}$ are equivalent:

(1) Any definable set can be represented by an existential formula
(2) Any definable set can be represented by a universal formula
(3) Any existential formula is equivalent to a universal one
(4) Any map of models of $\mathcal{T}$ is elementary
(5) For any model M of $\mathcal{T}$, $\mathcal{T}_M$ is complete.
(6) Any model of $\mathcal{T}$ is existentially closed

The theory $\mathcal{T}$ is called *model complete* if it satisfies these conditions. (The proof of this statement is in 2.1.12.)

End lecture 2
Lecture 3, Sep. 2, 2009

2.1.9. *Existential types.* To prove the equivalence in 2.1.8, we introduce the notion of a type. Given a theory $\mathcal{T}$, a *type* on a given tuple of variables x is a maximal consistent set of formulas on these variables. If a is a tuple in some model of $\mathcal{T}$, then $\text{tp}(a) = \{\phi(x) : \phi(a)\}$ is clearly a type. Conversely, any type is of this form, for some model M and tuple $a \in M$.

An *existential type* is the restriction of a type to existential formulas. Again it can be defined to by the existential type of some element. However, it is not, in general, maximal with respect to inclusion (find an example).

2.1.10. *Exercise.* Let M be a model, $a \in M$ a tuple. Show that if the existential type of a is not maximal, then there is an embedding $f : M \to N$ such that the existential type of $f(a)$ strictly contains that of a. (Hint: realise a proper extension of the existential type of a by an element b in a model N, consider the theory with constants for elements of both M and N, and show that it is consistent with the statement $a = b$).

2.1.11. *Exercise.* Show that the existential type of any tuple in an existentially closed model is maximal.

2.1.12. *proof of 2.1.8.* We prove only that (6) implies (3) since the rest is obvious.

Let $\phi(x)$ be existential. By 2.1.11, any existential type p that does not contain ϕ is, in fact, inconsistent with ϕ. Therefore, it contains an (existential) formula ψ_p that is disjoint from ϕ. Since any element has some existential type, we get that $\phi \vee \bigvee_p \psi_p$ covers the whole sort. By compactness, already $\phi \vee \psi_{p_1} \cdots \vee \psi_{p_k}$ covers. Hence ϕ is equivalent to the universal formula $\neg\psi_{p_1} \wedge \cdots \wedge \neg\psi_{p_k}$.

2.1.13. *Remark.* The above proof can be interpreted as follows. Let S^e be the topological space whose underlying set consists of the existential types, and whose closed sets are determined by the existential formulas. Then the proof shows that (in the case that T is model complete) this space is Hausdorff (2.1.11 asserts that it is T_1).

2.1.14. As explained in the introduction, model theory works by analysing the structure of definable sets. In the context of a particular theory, this is usually possible only if the definable sets are simple enough, i.e., if the theory eliminates quantifiers, or is at least model complete. Moreover, as explained above, the right kind of maps between models of a theory are the *elementary* embeddings. However, checking this condition is, in general, very difficult. The upshot is that we would like to work with model complete theories.

On the other hand, in the application we might be interested in an arbitrary structure. Thus, we would like to know that we can study arbitrary structures in terms of models of a model complete theory.

2.1.15. *Definition.* Let $\mathcal{T}_0$ be a theory. A theory $\mathcal{T}$ extending $\mathcal{T}_0$ is called a *model companion* of $\mathcal{T}_0$ if any model of $\mathcal{T}_0$ can be embedded in a model of $\mathcal{T}$, and $\mathcal{T}$ is model complete.

$\mathcal{T}$ is called a *model completion* of $\mathcal{T}_0$ if in addition it eliminates quantifiers.

2.1.16. *Proposition.* The model companion, if it exists, is unique.

Proof. Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two model-companions of a theory $\mathcal{T}$, and let M_1 be a model of $\mathcal{T}_1$. Since $\mathcal{T}_2$ is a model-companion, M_1 embeds into a model N_1 of $\mathcal{T}_2$. Since $\mathcal{T}_1$ is a model-companion, N_1 embeds in a model M_2 of $\mathcal{T}_1$. In particular, we get an embedding of M_1 in M_2, which is elementary since $\mathcal{T}_1$ is model complete. Continuing this way, we get two systems of models and elementary maps. The union is a model of both $\mathcal{T}_1$ and $\mathcal{T}_2$. In particular, any model of $\mathcal{T}_1$ is an elementary substructure of a model of $\mathcal{T}_2$. □

End lecture 3
Lecture 4, Sep. 4, 2009

2.2. **Affine varieties.** We recall some definition results on algebraic varieties. See, e.g., Milne [29] or Hartshorne [16] for details.

Let $\Bbbk$ be an algebraically closed field. An affine variety over $\Bbbk$ is, roughly speaking, a topological space with an algebra of "polynomial functions" on it, whose geometry is determined by these functions. More formally we have:

2.2.1. *Definition.* Let A be a finitely-generated integral domain over an algebraically closed field $\Bbbk$. We set $X = \operatorname{specm}(A)$ (the maximal spectrum of A) to be the set of $\Bbbk$-algebra homomorphisms from A to $\Bbbk$. The *affine variety (over $\Bbbk$) associated to* A is the pair (X, A). The set X is called the set of $\Bbbk$-points of the variety, and A is called the algebra of functions. More generally, if B is a $\Bbbk$-algebra, the set $X(B)$ of B-points of X is defined to be the set of $\Bbbk$-algebra homomorphisms $A \to B$. A map from the variety (X, A) to (Y, B) is by definition an A-point of Y. Note that any such map induces a map of sets $X \to Y$.

The algebra of a variety $\mathbf{X}$ is denoted by $\mathcal{O}(\mathbf{X})$. If L is an algebraically closed field extension of $\Bbbk$, we denote by $\mathbf{X} \otimes L$ the variety over L associated with $\mathcal{O}(\mathbf{X}) \otimes_{\Bbbk} L$.

2.2.2. *Exercise.* Let $\mathbf{X}$ be the variety associated with $A = \Bbbk[x_1, \dots, x_n]$. Show that $\mathbf{X}(B)$ can be canonically identified with the set of n-tuples from B. Show that if $\mathbf{X}$ and $\mathbf{Y}$ are two varieties of this form, then a map between them is given by polynomials (and conversely, any map given by polynomials is a map in the sense above).

2.2.3. *The Zariski topology.* Let $\mathbf{X}$ be the affine variety associated to A. Any element a of A can be viewed as a function $a : X \to \Bbbk$, $x \mapsto x(a)$. Given a subset I of A, we denote by $V(I)$ the subset of X consisting of all points $x \in X$ with $a(x) = 0$ for all $a \in I$. Any subset of this form is called an *algebraic subset* of $\mathbf{X}$. It is clear that a subset I defines the same algebraic subset as the ideal it generates, so we may always assume that I is an ideal. On the other hand, Hilbert's basis theorem implies that any such ideal is finitely generated.

The *Zariski topology* on X is the topology whose closed subsets are the algebraic subsets. The basis theorem implies that this topology is Noetherian: any proper descending chain of closed subsets is finite. With each subset V of X we may associated an ideal $I(V)$ in A, namely the ideal of all elements $a \in A$ that are 0 on all elements of V. If $V = V(I)$ is a closed subset, Hilbert's Nullstelensatz implies that $I(V)$ is the radical ideal of I.[2]

If V is a closed subset of X, then any element of ${}^{A}\!\big/_{I(V)}$ determines a function on V, and distinct elements determine distinct functions. Hence, if $I(V)$ is prime, then V itself has the structure of a variety (and it is called a subvariety of $\mathbf{X}$). This happens precisely if V is irreducible, i.e., if it cannot be written as a union of two proper sub-varieties. The *dimension* of $\mathbf{X}$ is the length of the longest chain of sub-varieties of $\mathbf{X}$.

2.2.4. *Example.* If $\mathbf{X}$ is an affine space $\mathbb{A}^n$, the closed subsets are sets of solutions of a (finite) system of polynomial equations.

End lecture 4
Lecture 5, Sep. 7, 2009

2.2.5. *Generic points.* Let L be a field extension of $\Bbbk$. An L-point of $\mathbf{X}$ is called *generic* if it is not contained in any proper sub-variety of $\mathbf{X}$. It follows that it is not contained in any proper closed subset. Equivalently, it is injective when viewed as an algebra homomorphism. For any L-point x there is a unique subvariety $\mathbf{V}_x$ of $\mathbf{X}$ such that x is generic in $\mathbf{V}_x$ (it is defined by the kernel of x). Conversely, any variety has a canonical generic point, namely, the embedding of the algebra in its fraction field.

[2] The Nullstelensatz can also be stated as saying that any non-zero $\Bbbk$-algebra has a $\Bbbk$-point, or equivalently, that any set of equations that generate a proper ideal has a solution in $\Bbbk$. In this form, it essentially says that the theory of algebraically closed fields is model complete.

A map $f : \mathbf{X} \to \mathbf{Y}$ of varieties (or algebraic subsets) is called *dominant* (or generically onto) if the image is Zariski dense. In other words, it is injective on the level of algebras. This is also equivalent to saying that it maps a generic point of $\mathbf{X}$ to a generic point of $\mathbf{Y}$.

2.2.6. *Families.* If $\mathbf{X}$ and $\mathbf{Y}$ are two affine varieties, their product $\mathbf{X} \times \mathbf{Y}$ is the variety associated with $\mathcal{O}(\mathbf{X}) \otimes_{\Bbbk} \mathcal{O}(\mathbf{Y})$ (this is again an integral domain, since $\Bbbk$ is algebraically closed, see 2.3.8). A *family of algebraic subsets* of $\mathbf{X}$, parametrised by $\mathbf{Y}$, is simply an algebraic subset $\mathbf{Z}$ of $\mathbf{X} \times \mathbf{Y}$. For any L-point $y \in \mathbf{Y}$ (with L an algebraically closed field extension of $\Bbbk$), the *fibre* of $\mathbf{Z}$ over y is by definition the algebraic subset of $\mathbf{X} \otimes L$ corresponding to the radical of the kernel of the map from $\mathcal{O}(\mathbf{X}) \otimes_{\Bbbk} L$ to $\mathcal{O}(\mathbf{Z}) \otimes_{\mathcal{O}(\mathbf{Y})} L$. It is denoted by $\mathbf{Z}_y$ (Exercise: the L-points of $\mathbf{Z}_y$ are precisely the L-points of $\mathbf{Z}$ that map to y.)

2.2.7. *Example.* Assume that $\mathbf{Y} = \mathbb{A}^n$. A family over $\mathbf{Y}$ then corresponds to a finitely generated (reduced) algebra over the polynomial algebra, i.e., polynomials in variables X_i, whose coefficients are polynomials in variables $Y_i \in \mathcal{O}(\mathbf{Y})$. Then giving an L-point of $\mathbf{Y}$ consists of giving values $y_1, \dots, y_n \in L$ to these variables, and the fibre is obtained by substituting these values in the coefficients.

2.2.8. *Example.* Let $\mathbf{X} = \mathbb{A}^1$, $\mathbf{Y} = \mathbb{A}^n$ and $\mathbf{Z}$ the subset of $\mathbf{X} \times \mathbf{Y} = \mathbb{A}^{n+1}$ given by $\sum Y_i X^i = 0$. The fibre over a given point $(y_1, \dots, y_n)$ is the subset of $\mathbf{X}$ given by the polynomial $\sum y_i X^i$. Thus, this family contains all finite subsets of size at most $n-1$. Similarly, for any affine variety, the set of all algebraic subset with "bounded discrete data" is part of a family.

End lecture 5
Lecture 6, Sep. 9, 2009

2.2.9. *Constructible subsets.* A *constructible subset* of a variety $\mathbf{X}$ is a boolean combination of algebraic subsets. Topologically, it is a union of locally-closed subsets, where a locally-closed subset is an open subset of a closed subset.

As mentioned above, an algebraic subset of affine space is given by a finite number of polynomial equations. Thus such a subset can be represented by a certain quantifier-free formula in the language of rings. Hence the same is true for constructible subsets. Conversely, any such formula determines a constructible subset.

It thus follows from the elimination of quantifiers for ACF (and from the fact that constructible sets can be distinguished by points in fields) that constructible subsets of ${\mathbb{A}^n}_{\Bbbk}$ can be identified with definable sets of n-tuples in $ACF_{\Bbbk}$. In other words, to say that a set is constructible is the same as to say that it is definable. It also follows from this remark that the notions of an algebraic subset or a sub-variety of affine space can be used to describe definable subsets in any theory that interprets ACF.

2.2.10. *Fact.* Given a family $\mathbf{Z} \to \mathbf{Y}$, it is natural to ask what can we say about the set of points $\mathbf{Y}$ where the fibre $\mathbf{Z}_y$ has a given property. For example, we may consider the set of points where the fibre has 7 points, or is connected, or has a given dimension. It turns out that all such subsets are constructible subsets of $\mathbf{Y}$. In some cases, this follows easily from the identification with definable sets above, together with elimination of quantifiers. The result we need is not as easy, but also has a model theoretic proof (Dries and Schmidt [13]): The set of points y of $\mathbf{Y}$ where $\mathbf{Z}_y$ is a sub-variety (i.e., is irreducible) is constructible.

2.2.11. *Generic fibres.* Let $\mathbf{Z} \to \mathbf{Y}$ be a family, and let y be a generic point of $\mathbf{Y}$ over an algebraically closed field. The fibre $\mathbf{Z}_y$ is called a *generic fibre* of the family. We make the following observation: let $P(\mathbf{V})$ be a property of algebraic sets $\mathbf{V}$ which is constructible in the above sense (e.g., $P(\mathbf{V})$ says that $\mathbf{V}$ is of size 7). Then the following are equivalent:

(1) The generic fibre $\mathbf{Z}_y$ satisfies P
(2) There is a non-empty open subset $\mathbf{U}$ of $\mathbf{Y}$ such that P holds for all $\mathbf{Z}_u$ where u is a point of $\mathbf{U}$
(3) There is a dense set U of points of $\mathbf{Y}$ such that P holds for all $\mathbf{Z}_u$ for $u \in U$

In particular, in (2) it is enough to check $\Bbbk$-points.

Proof (sketch). Assume (2) does not hold. By assumption, the set of point where P holds is then contained in a proper closed subset. But y is contained in no such subset.

(2) implies (3) since any non-empty open subset of a variety is dense. Finally, assuming (3), we see by constructibility that (2) holds, and since y is also a point of the open set, (1) holds as well. □

End lecture 6
Lecture 7, Sep. 11, 2009

2.3. **The theory** $ACFA$**.** We work in the language of rings, expanded by a unary function symbol σ. Our aim is to describe the model companion of the theory FA of difference domains, i.e., pairs (A, σ) where A is an integral domain, and σ is an injective endomorphism of A.

Let K be an algebraically closed field. From now on, we use the term 'variety over K' to mean an affine sub-variety of affine space over K. If σ is an endomorphism of K, and $\mathbf{U}$ is a variety, we denote by $\mathbf{U}^\sigma$ the variety obtained from $\mathbf{U}$ by applying σ to the coefficients of the defining equations. It can be alternatively described as $\mathbf{U} \otimes_\sigma K$ (i.e., tensor with K viewed as a K algebra via σ).

We note that, setting $A = \mathcal{O}(\mathbf{U})$, we have an injective map of *rings* (rather than K-algebras) from A to $A \otimes_\sigma K = \mathcal{O}(\mathbf{U}^\sigma)$, given by $a \mapsto a \otimes 1$. This map extends the endomorphism σ on K.

If N is a field extension of K, then an N point $x : A \otimes_\sigma K \to N$ of $\mathbf{U}^\sigma$ is the same as an algebra homomorphism f from A to N, where N has the K-algebra structure coming from σ (i.e., $f(xa) = \sigma(x) f(a)$, for $x \in K$ and $a \in N$). In particular, if N itself has an endomorphism σ_1 extending σ, then each point $x : A \to N$ of $\mathbf{U}$ determines a point $\sigma_1(x)$ of $\mathbf{U}^\sigma$ induced by $\sigma_1 \circ x$. In fact, $\mathbf{U}^\sigma$ is the smallest sub-variety containing such points.

2.3.1. *Exercise.* Show that if (N, σ) is a difference field extending (K, σ), and $x : A \to N$ is a generic point of $\mathbf{U}$, then $\sigma(x)$ is a generic point of $\mathbf{U}^\sigma$ (this is easy if σ is an automorphism of K, slightly harder in general).

2.3.2. *Definition.* The theory $ACFA$ is the extension of FA consisting of the following schemes of axioms for a structure (K, σ).

(1) σ is surjective.
(2) K is an algebraically closed field.
(3) Let $\mathbf{U}$ be an affine sub-variety of some affine space, $\mathbf{V} \subseteq \mathbf{U} \times \mathbf{U}^\sigma$ a subvariety projecting dominantly onto both factors, and $\mathbf{W} \subset \mathbf{V}$ a proper algebraic subset. Then there is a point $x \in \mathbf{U}(K)$, such that $(x, \sigma(x)) \in \mathbf{V} - \mathbf{W}$.

We note that the third condition is indeed given by a scheme of first order axioms, by 2.2.10.

In the situation of the third axiom, we will denote by $\mathbf{V}^{\#}$ the set of points of the form $(x, \sigma(x))$ in $\mathbf{V}$. If $f : \mathbf{U} \to \mathbf{U}^{\sigma}$ is a dominant (algebraic) function, we denote by $f^{\#}$ the definable set given by the equation $f(x) = \sigma(x)$ (i.e., $\mathbf{V}^{\#}$ where $\mathbf{V}$ is the graph of f).

2.3.3. *Theorem.* $ACFA$ is the model companion of the theory of difference domains.

Proof. We first show that $ACFA$ is model complete. We use 6. Let M be a model of $ACFA$, and let $\phi(x)$ be a quantifier free formula over M. Then, by inspection, there is a formula $\psi(x_0, \dots, x_n)$ in the language of rings, such that $\phi(x) \iff \psi(x, \sigma(x), \dots, \sigma^n(x))$. Furthermore, there is another formula $\theta(y_0, \dots, y_{n-1}, z_1, \dots, z_n)$ such that $\phi(x) \iff \theta(x, \dots, \sigma^{n-1}(x), \sigma(x), \dots, \sigma^n(x))$. Any formula in ACF is equivalent to a union of locally closed sets, i.e., sets of the form $\mathbf{V} - \mathbf{W}$, with $\mathbf{V}$ a variety, and $\mathbf{W}$ an algebraic subset. Hence we may assume that θ itself is locally closed.

Let a be a solution of ϕ in some model N, and set $b = (a, \sigma(a), \dots, \sigma^{n-1}(a))$. Hence, $(b, \sigma(b))$ is a solution of θ. By 2.2.5, b is generic in some variety $\mathbf{U}$. Therefore, $\sigma(b)$ is generic in $\mathbf{U}^{\sigma}$. Let $\mathbf{V}$ be the sub-variety of $\mathbf{U} \times \mathbf{U}^{\sigma}$ of which $(b, \sigma(b))$ is generic. Since $(b, \sigma(b))$ maps to the generic point b of $\mathbf{U}$ under the first projection, it again follows from 2.2.5 that the projection from $\mathbf{V}$ to $\mathbf{U}$ is dominant (and likewise for the projection to $\mathbf{U}^{\sigma}$). Since $(b, \sigma(b))$ is a solution of θ, it follows that for some proper algebraic subset $\mathbf{W}$ of $\mathbf{V}$, $\mathbf{V} - \mathbf{W}$ is contained in the set defined by θ. By the third axiom, there is an M-point of the form $(c, \sigma(c))$ in $\mathbf{V} - \mathbf{W}$ and therefore in θ. Hence $\phi(c)$ holds in M.

To prove that $ACFA$ is the model companion, it remains to show that any difference domain (A, σ) embeds in a model. Since σ extends uniquely to the fraction field of A, we may assume that A is a field. Furthermore, any endomorphism of A extends to an automorphism of some algebraically closed field containing A (this follows for example from elimination of quantifiers for ACF), so we may assume that A is algebraically closed, and σ is an automorphism.

Assume we are given the data of the third axiom. Let $b : B = \mathcal{O}(\mathbf{V}) \to L = K(\mathcal{O}(\mathbf{V}))$ be the generic point of $\mathbf{V}$. The projections induce maps $p : A = \mathcal{O}(\mathbf{U}) \to B$ and $q : A^{\sigma} \to B$, which are injective, since by assumption, the projections are dominant. Viewing A as a sub-algebra of L, we have the partial endomorphism σ $a \mapsto b(q(a \otimes 1))$ of L, which is injective by 2.3.1. As above, this endomorphism extends to an automorphism of the algebraic closure of L. By definition, $b = (a, \sigma(a))$, where $a = b \circ p$ is the projection of b to $\mathbf{U}$. Finally, since b is a generic point of $\mathbf{V}$, it avoids $\mathbf{W}$ (and any other proper algebraic subset).

Thus we have managed to find an extension that satisfies one instance of the axiom. Taking the union over all instances, we get a model of $ACFA$ containing the original field. □

End lecture 7
Lecture 8, Sep. 14, 2009

2.3.4. *Exercise.* Let (K, σ) be a model of $ACFA$, and let $n \neq 0$ be an integer. Show that (K, σ^n) is also a model of $ACFA$. Show that this is not the case when $n = 0$. Show that there is no model of $ACFA_p$, $p > 0$, where σ is the Frobenius.

2.3.5. *Exercise.* Let (K, σ) be a model of $ACFA$. We denote by C the fixed set of elements $x \in K$ satisfying $\sigma(x) = x$.

1. Show that C is a field
2. Show that C is perfect
3. Show that if $\mathbf{V}$ is a variety that can be defined with parameters from C, then it has a point in C (this says that C is a *pseudo-algebraically closed* (PAC)).
4. Show that C has an extension of degree n, for each n.
5. Show that $\sigma(C^a) = C^a$
6. Show that if $D \subseteq K$ is a finite Galois extension of C, then $\mathrm{Aut}({}^{D}\!/_{C})$ is generated by σ.
7. Conclude that C has a unique extension of each degree (up to isomorphism).

A field that satisfies (2), (3) and (7) is called *pseudo-finite*. It is a result of Ax [1] that the theory axiomatised by these statements is precisely the asymptotic theory of finite fields (cf. 1.4).

2.3.6. *Exercise.* Let (K, σ) be a difference field, with fixed field C, and let A be an $n \times n$ matrix over it. Consider the subset of K^n defined by $\sigma(\bar{x}) = A\bar{x}$.

1. Show that the set of solutions of the above equation forms a vector space over C, of dimension at most n (For the dimension, solve first the case $n = 1$, then consider n-th exterior powers for the general case.)
2. Show that if (K, σ) is a model of $ACFA$, then the dimension is precisely n, and that the set of solutions is dense in K^n.

2.3.7. The theory $ACFA$ is not complete. For instance, the characteristic of the field is not determined. Fixing characteristic complete the theory of the field, but as a difference field, the theory is still not complete: for example, the statement $\forall x(x^2 = 2 \implies \sigma(x) = -x)$ is expressible in the language, but is not determined by $ACFA$ (since $\mathbb{Q}(\sqrt{2})$ with both its automorphisms are embedable in models of $ACFA$). It turns out that the completions are determined by statements of this form. To prove this, we need the following result.

2.3.8. *Lemma.* Let A and B be two integral domains over a perfect field $\Bbbk$. If $K(A) \cap \Bbbk^a = \Bbbk$ (in particular, $\Bbbk$ is algebraically closed), then $A \otimes_{\Bbbk} B$ is a domain.

proof (sketch; cf Eisenbud [15]). By taking the fraction field, we may assume that B is a field extension. Since the problem is finitary, we may assume that this field is finitely generated. Finally, we may assume by induction that $B = \Bbbk(b)$ is generated by one element. The minimal (or 0) polynomial p of b is, by assumption, still irreducible over $K(A)$. Hence $A \otimes_{\Bbbk} B = {}^{A[t]}\!/_{(p(t))}$, where $(p(t))$ is prime. □

End lecture 8
Lecture 9, Sep. 16, 2009

We shall later (3.2.32) explain a vast generalisation of this lemma, in the context of stable theories.

2.3.9. *Theorem.* Let (E, σ) be a common sub-structure of two models M_1 and M_2 of $ACFA$, such that E is algebraically closed as a field. Then M_1 and M_2 have the same theory over E.

Proof. By 2.3.8, $M_1 \otimes_E M_2$ is a domain. It has a natural endomorphism σ (coming from the universal property), that extends to the fraction field M. This difference field embeds in some model of $ACFA$, giving an embedding of M_1 and M_2 into the same model N of $ACFA$. Since the theory is model complete, these embeddings are elementary, so all of these models have the same theory by 2.1.3. $\square$

2.3.10. *Corollary.* Let F be a prime field of characteristic p (possibly 0), and let G be its Galois group. There is a one-to-one correspondence between completions of $ACFA_p$ and conjugacy classes in G.

Proof. Let $E = F^a$, and let $\sigma \in G$ be an automorphism of E. By 2.3.3, (E, σ) embeds in a model M of $ACFA$. We let T_σ be the theory of M. By 2.3.9, it does not depend on M.

If $\tau \in G$ is a conjugate of σ, $\tau = \phi^{-1}\sigma\phi$, then (E, τ) is a sub-structure of M as above, with embedding given by ϕ. Hence $T_\tau = T_\sigma$.

If T is any completion of $ACFA_p$, and M a model of T, then $T = T_\sigma$, where σ is the restriction of the automorphism of M to the prime field inside M.

Finally, if $T_\sigma = T_\tau$, let M be a model. By assumption there are embeddings of (E, σ) and (E, τ) in M. Since the image of both embeddings is the same, they determine an automorphism ϕ of E that conjugates between σ and τ. $\square$

2.3.11. *Decidability.* It follows from the last corollary that the theories $ACFA_p$ and $ACFA$ are decidable. For simplicity, we shall explain this only for the case $p = 0$.

We first note that any complete theory $\mathcal{T}$ that has a recursive axiomatisation is decidable: Indeed, to check if a sentence ϕ is in $\mathcal{T}$, we start from the axioms, and try to prove both ϕ and $\neg\phi$. Since $\mathcal{T}$ is complete, we will eventually prove one of them.

As mentioned above, $ACFA_0$ is not complete, but the completions are described in 2.3.10. Given a conjugacy class c of elements of $\mathrm{Aut}({}^{\mathbb{Q}^a}\!/_{\mathbb{Q}})$, the completion T_c is axiomatised by a collection of sentences $\phi_{E,c}$, one for each finite Galois extension E of $\mathbb{Q}$, which describe the action of (an element of) c on E. A given sentence ψ is equivalent in $ACFA_0$ to a finite disjunction $\bigvee_i \phi_{E,c_i}$ (by compactness). These ϕ_{E,c_i} can be found effectively, as before. This disjunction holds in $ACFA_0$ precisely if any conjugacy class of elements of $\mathrm{Aut}({}^{E}\!/_{\mathbb{Q}})$ occurs in it. Hence the same is true for ψ.

2.3.12. *Algebraic closure and definable closure.* Let $\mathcal{T}$ be any theory, and let A be a subset of a model M of $\mathcal{T}$. An element $b \in M$ is *algebraic* over A if there is a formula $\phi(x)$ over A, such that $\phi(b)$ holds, and there are only finitely many $c \in M$ for which $\phi(c)$ holds. It is called *definable* over A if b is the only such c. For example, for $\mathcal{T} = ACF$, b is algebraic over A if and only if it is algebraic over the field generated by A, in the sense of field theory.

The *algebraic closure* $\mathrm{acl}(A)$ of A is the set of all elements algebraic over A, and the *definable closure* $\mathrm{dcl}(A)$ is the set of all elements. So $\mathrm{acl}(A)$ contains $\mathrm{dcl}(A)$, which contains the substructure generated by A. Hence, when computing the algebraic closure, we may first assume that A is a substructure.

We note that $\mathrm{acl}(A)$ and $\mathrm{dcl}(A)$ remain unchanged when passing to elementary extensions, and so for a model-complete theory depend only on A.

2.3.13. *Exercise.* Show that an element a is algebraic over A if and only $\mathrm{tp}(a/A)$ has finitely many solutions in every model, and is definable over A if and only if this type has only one solution in every model. Show that "in every model" cannot be omitted.

End lecture 9
Lecture 10, Sep. 18, 2009

2.3.14. In the context of $ACFA$, if A is a substructure (i.e., a difference field), the field theoretic algebraic closure A^a is definitely contained in $\mathrm{acl}(A)$. We may use the method of 2.3.9 to show that they are in fact the same:

2.3.15. *Proposition.* Let A be a difference sub-field of a model (M, σ) of $ACFA$. Then $\mathrm{acl}(A) = A^a$, the field theoretic algebraic closure of A.

Proof. Let $c \in M$ be an A-algebraic element outside $B = A^a$, with c satisfying an A-formula $\phi(x)$, with n solutions. By 2.3.8, $M \otimes_B M$ is a domain, and as in 2.3.9, the coordinate-wise action determines an endomorphism on it. Thus it embeds in a model N of $ACFA$. Since M is also a model, and $ACFA$ is model complete, the two inclusions of M in N are elementary. In particular, the images c_1 and c_2 under the two inclusions both satisfy ϕ. But since $c \notin B$, these two elements are different, and so we have $2n$ solutions of ϕ. This is a contradiction. □

2.3.16. *Remark.* The same is not true for the definable closure. For example, consider a model of $ACFA_0$ where $\sigma(\imath) = -\imath$ (where $\imath^2 = -1$). Then for each non-zero element x of the fixed field, precisely one of x and $-x$ has a square root in the field. In particular, if σ fixes the square roots of 2, then one of them is definable (over 0) by the formula $x^2 = 2 \wedge \exists y(y^2 = x \wedge \sigma(y) = y)$. See Beyarslan and Hrushovski [5] for an analysis of the definable closure.

2.3.17. *A description of the types.* Theorem 2.3.9 also allows us to describe when two tuples have the same type:

Two tuples b and c in M have the same type over a substructure A if and only if there is an isomorphism (of difference fields) over A from $\mathrm{acl}(A(b))$ to $\mathrm{acl}(A(c))$, sending b to c.

Proof. Assume there is an isomorphism τ, and let $E = \mathrm{acl}(A(b))$. Applying 2.3.9 with E, $M_1 = M$, and $M_2 = M$ with E embedded via τ, we get that M has the same theory over the parameters E, where b is interpreted as b or as $\tau(b) = c$ (and A always interpreted as A). Hence b and c have the same type over A.

Conversely, if b and c have the same type over A, there is, first, a unique map τ (of difference fields) over A from $A(b)$ to $A(c)$ sending b to c. If d is algebraic over $A(b)$, since b and c have the same type, we may find e such that $\phi(b, d)$ holds if and only if $\phi(c, e)$ holds, for all A-formulas ϕ. Thus we may extend τ by setting $\tau(d) = e$. Continuing this way, we get an isomorphism between $\mathrm{acl}(A(b))$ and $\mathrm{acl}(A(c))$. □

2.3.18. *Exercise.* Let $\mathcal{T}$ be an arbitrary complete theory. Show that the following are equivalent:

(1) For each definable set $\mathbf{X}$ the restriction map from the space $S(\mathbf{X})$ of types of elements of X, to the space $S^{qf}(\mathbf{X})$ of quantifier free such types is a bijection (in other words, any quantifier free type has a unique extension).
(2) T admits elimination of quantifiers

2.3.19. *Description of definable sets.* In light of 2.3.18, it is not surprising that 2.3.17 gives information about definable sets in $ACFA$. We already know by model completeness that every definable set is existential. Now we have the following more precise description: Any definable set $\mathbf{X}$ can be represented by a formula of the form $\exists y \phi(x, y)$, where y is just one variable (rather than a tuple), $\phi(x, y)$ is quantifier-free, and the projection from $\phi(x, y)$ to $\mathbf{X}$ has finite fibres.

The proof is similar to the above exercise. The main point is that the complement of a set of the form $\psi(x) = \exists y\phi(x,y)$ for ϕ as above again has the same form. Indeed, taking Zariski closure, we see that ψ is dense in the image of some algebraic map with generically finite fibres. Hence the complement of ψ is the union of some quantifier-free set with the image of the complement of ϕ within its Zariski closure. Using this, one may repeat the above exercise, or the proof of 2.1.8.

We note that using some facts from algebraic geometry, we may require that the projection has additional nice properties (flat, étale).

2.3.20. *Exercise.* Show that $\exists y(y^2 = x \wedge \sigma(y) = -y)$ is an infinite and co-infinite subset of the fixed field. Find an existential representation of the complement. End lecture 10

Lecture 11, Sep. 21, 2009

3. Stable theories

Our aim is now to setup a general model theoretic framework in which the result of $ACFA$ can be interpreted. This will mostly be done in 4, where a "generic" automorphism is attached to models of a theory $\mathcal{T}$. In this section we study a class of theories $\mathcal{T}$ for which this process yields a theory with good properties. Namely, we are interested in stable theories. Before defining them, we need to discuss imaginaries.

3.1. **Elimination of imaginaries.** Let $\mathcal{T}$ be an arbitrary theory, $\mathbf{X}$ a definable set in it. A *definable equivalence relation* on $\mathbf{X}$ is a definable subset $\mathbf{E} \subseteq \mathbf{X} \times \mathbf{X}$ such that $\mathbf{E}(M)$ is an equivalence relation on $\mathbf{X}(M)$ for any model M of $\mathcal{T}$. A *quotient* of such a relation consists of a definable set ${}^{\mathbf{X}}\!/_{\mathbf{E}}$ and a definable surjective function $\pi_{\mathbf{E}} : \mathbf{X} \to {}^{\mathbf{X}}\!/_{\mathbf{E}}$ such that $\mathbf{E}(x,y)$ if and only if $\pi_{\mathbf{E}}(x) = \pi_{\mathbf{E}}(y)$. Clearly, the quotient, if it exists, is unique up to a unique definable bijection. Also, the surjectivity condition can always be attained by taking the image. $\mathcal{T}$ is said to admit *elimination of imaginaries* (EI) if any definable equivalence relation in $\mathcal{T}$ has a quotient.

Below we discuss several equivalent formulations of this condition. I refer the reader to Marker [25] (for example) for details.

3.1.1. *Example.* If G is a group and $\mathbf{X}$ is a definable set in $\mathcal{T}$, an *action* of G on $\mathbf{X}$ is a group homomorphism from G to the group $Aut(\mathbf{X})$ of definable bijections from $\mathbf{X}$ to itself. If G is a finite group acting on $\mathbf{X}$, the relation of being in the same orbit is a definable equivalence relation.

In particular, if $\mathbf{X} = \mathbf{Y}^n$ for some other definable set $\mathbf{Y}$, the symmetric group S_n acts on $\mathbf{X}$ by permuting the coordinates. The equivalence classes can be identified with the set of subsets of $\mathbf{Y}$ of size at most n. If $\mathcal{T}$ eliminates all such imaginaries, we say that $\mathcal{T}$ codes finite sets.

We note that any theory that expands a theory that codes finite sets (without adding new sorts) also codes them. In particular, the theory of fields codes the finite set $\{a_1, \dots, a_n\}$ by the coefficients of the polynomial $(x-a_1)\dots(x-a_n)$ (i.e., the quotient map is given by the elementary symmetric polynomials). Hence any theory of fields (possibly with extra structure, but no new sorts) codes finite sets. In fact, the theory ACF eliminates imaginaries (see, e.g., Marker et al. [26]).

3.1.2. *Homogeneous and saturated models.* Let M be a structure, $A \subseteq M$ a subset, $b, c \in M$ two tuples. If there is an automorphism of M over A that takes b to c, then $\mathrm{tp}(b/A) = \mathrm{tp}(c/A)$. The converse if false, in general. For example, if $f : A \to B$ is a function onto a set B with two elements that has infinite fibres, then the two elements of B clearly have the same type over 0, but if the fibres have different cardinality, there is no automorphism that will take one to the other. End lecture 11

Lecture 12, Sep. 23, 2009 If κ is a cardinal, the model M is called (strongly) *κ-homogeneous* if the converse holds whenever A has cardinality less than κ, i.e., if for any such A, and for any tuples $b, c \in M$ with the same type over A, there is an automorphism of M taking b to c. It is a fact that for any κ, any model is an elementary sub-model of some κ-homogeneous one. We will call a subset of cardinality less than κ *small*.

A κ-homogeneous model is sometimes insufficient to determine all properties of the theory (or the small subset). κ-homogeneity means that the realisation of types over a small set A are the orbits under the action of $\mathrm{Aut}({}^{M}\!/_{A})$. However these orbit might be too small (or even empty) to capture the properties of the type. For example, the real numbers form an ω-homogeneous model of the theory of real-closed fields, but only because any type is realised by at most one element (this model has no automorphisms).

A model M is called *κ-saturated* if any type over a small subset is realised in M. Again it is the case that any model is an elementary sub-model of a κ-saturated and κ-homogeneous model. Working inside this model, we obtain the full information about types over small subsets from the action of the automorphism groups.

The strongest notion is that of a saturated model: A model M is *saturated* if it is saturated in its own cardinality. Such a model is automatically homogeneous in its own cardinality as well. Under additional set-theoretic assumptions (existence of strongly inaccessible cardinals), saturated models of arbitrarily large cardinality exist, and any models is an elementary sub-model of a saturated one. There is a meta-theorem (Shoenfield [38]) that says that "anything" proved using saturated models in model theory does not in fact depend on the set theoretic assumption. Alternatively, in most cases assuming the model is κ-saturated and κ-homogeneous is enough. We will use the term "saturated" to mean one of the two.

See, e.g., Buechler [7, Ch. 2] for more details on this stuff.

3.1.3. *Exercise.* Let M be a saturated model of a complete theory. Show that if X is a definable set, and S is a small set of definable sets (all with parameters) such that $X(M) = \bigcap_{Y \in S} Y(M)$, then X is equivalent to a finite intersection of sets $Y \in S$.

3.1.4. *Exercise.* Assume that M is κ-saturated and κ-homogeneous. For a small $A \subseteq M$, show that $\mathrm{dcl}(A) = \{b \in M : \tau(b) = b \quad \forall \tau \in \mathrm{Aut}({}^{M}\!/_{A})\}$ and that $\mathrm{acl}(A)$ is the set of elements with that have finitely many conjugates over A (a conjugate over A of an element $b \in M$ is an element of the form $\tau(b)$ for some $\tau \in \mathrm{Aut}({}^{M}\!/_{A})$).

3.1.5. *Descent.* Using the above exercise, we get a natural definition of "$\mathrm{dcl}(X)$" and "$X \in \mathrm{dcl}(A)$" for any kind of objects $X \in S$, where S is a set on which $\mathrm{Aut}(M)$ acts. For example, if X is a definable set, defined with parameters in some saturated model M, then any automorphism τ of M moves X to another definable set $\tau(X)$ (by acting on the parameters used to define X; alternatively by applying τ to each of the points $X(M)$). Let $G_X = \{\tau \in \mathrm{Aut}(M) : \tau(X) = X\}$ be the stabiliser of X

(as a set). Then $\mathrm{dcl}(X) = \{a \in M : \tau(a) = a \quad \forall \tau \in G_X\}$. Likewise, $\mathrm{acl}(X) = \{a \in M : \#(G_X \cdot a) < \infty\}$. On the other hand, $X \in \mathrm{dcl}(A)$ for a (small) set of parameters A if $\mathrm{Aut}({}^{M}\!/\!_{A}) \subseteq G_X$, and $X \in \mathrm{acl}(A)$ if ${}^{\mathrm{Aut}({}^{M}\!/\!_{A})}\!/\!_{(G_X \cap \mathrm{Aut}({}^{M}\!/\!_{A}))}$ is finite.

Exercise 3.1.6 will show that $X \in \mathrm{dcl}(A)$ precisely if X can be defined with parameters in A. In this sense saturated models are models in which descent theory works: we may detect parameters via the action of the automorphisms group.

3.1.6. *Exercise.* Show that $X \in \mathrm{dcl}(A)$ precisely if X can be defined by a formula with parameters in A. Show also that if $X \in \mathrm{dcl}(A)$, then $\mathrm{dcl}(X) \subseteq \mathrm{dcl}(A)$.

End lecture 12
Lecture 13, Sep. 25, 2009

3.1.7. *Canonical parameters.* Let X be a definable set (with parameters), and let $B = \mathrm{dcl}(X)$ be the set of elements fixed by G_X. We saw in 3.1.6 that if X can be defined by a formula with parameters in A, then $B \subseteq A$. If X can actually be defined with parameters from B, we say that X can be defined with a *canonical parameter*, and that B is the *canonical base* for X. Thus the canonical base (if it exists) is the smallest subset over which X is defined.

3.1.8. *Example.* Let T be the theory of an equivalence relation with more than one class. Then the class of an element is definable using any element of the class, but any such element is definably closed. Hence there is no smallest subset over which this class is defined.

3.1.9. *Example.* The definable set X given by $x = \sqrt{2} \vee x = -\sqrt{2}$ in a saturated algebraically closed field is fixed by all automorphisms, hence its definable closure is the prime field. It is indeed definable over the prime field, by the formula $x^2 = 2$.

The same definable set also makes sense when the structure is considered in the empty theory. The definable closure of X is now the empty set, over which X is not definable.

3.1.10. *Definable families.* If $Z \subseteq X \times Y$ are definable sets (over 0), we may view, as before, Z as a family of definable subsets of X, parametrised by Y. Given $a \in Y(M)$, we denote by Z_a the a-definable subset $\{x \in X : (x, a) \in Z\}$. We note that any M-definable set has the form Z_a for some Z and a. The family Z is called *canonical* if $Z_a \neq Z_b$ for $a \neq b$.

3.1.11. $\mathcal{T}^{eq}$. Any theory $\mathcal{T}$ has a canonical universal expansion $\mathcal{T}^{eq}$ that eliminates imaginaries (it is universal in the following sense: Any interpretation of $\mathcal{T}$ in a theory that eliminates imaginaries extends to an interpretation of $\mathcal{T}^{eq}$, unique up to a unique isomorphism over $\mathcal{T}$). Any model M of $\mathcal{T}$ has a unique expansion M^{eq} to a model of $\mathcal{T}^{eq}$, and any model of $\mathcal{T}^{eq}$ restricts to a model of $\mathcal{T}$.

$\mathcal{T}^{eq}$ may be constructed as follows: The sorts of $\mathcal{T}^{eq}$ are indexed by pairs (X, E), where X is a definable set in $\mathcal{T}$ and E is a definable equivalence relation on X (all over 0). For any definable function $f : X \to Y$ in $\mathcal{T}$, there is a function symbol $f : (X, =) \to (Y, =)$. In addition, for sort (X, E) there is a function symbol $\pi_E : (X, =) \to (X, E)$. The theory says that the restriction to the sorts of the form $(X, =)$ satisfies $\mathcal{T}$ (where everything is interpreted in the natural way), and that each map π_E is the quotient of the relation $(E, =)$ on $(X, =)$.

3.1.12. *Exercise.* Prove that the theory constructed above indeed satisfies the properties attributed to $\mathcal{T}^{eq}$.

3.1.13. *Proposition.* Let $\mathcal{T}$ be a complete theory[3], and let M be a saturated model. The following are equivalent:

(1) $\mathcal{T}$ eliminates imaginaries
(2) Any M-definable set has the form Z_a for a canonical family Z.
(3) Any M-definable set can be defined with a canonical parameter
(4) For any definable equivalence relation E on X, for any equivalence class C, and for any type p of an element in $C(M)$ over $\mathrm{dcl}(C)$, $p(M)$ is contained in $C(M)$.
(5) Any element of M^{eq} is inter-definable with an element of M.

Proof. We prove only that (4) implies (5). Let e be any element in M^{eq}. Then e is the code of some equivalence class C. Let $A = \mathrm{dcl}(C)$ (in the sense of $\mathcal{T}$). Any automorphism M^{eq} that fixes e fixes $C(M)$ (as a set) and therefore fixes A. Conversely, if τ fixes A pointwise, then it fixes $p(M)$ for any type p over A. By assumption it therefore maps $C(M)$ into itself, so it fixes e. We have shown that an automorphism fixes e if and only if it fixes A pointwise. Hence $\mathrm{dcl}^{eq}(A) = \mathrm{dcl}^{eq}(e)$. In particular, there is a finite tuple $a \in A$ such that $\mathrm{dcl}^{eq}(a) = \mathrm{dcl}^{eq}(e)$. □

3.1.14. *Exercise.* Prove the rest of 3.1.13.

3.1.15. *Exercise.* Show that the theory of an infinite set in the empty language does not eliminate imaginaries.

3.1.16. From now on, we assume that we work inside $\mathcal{T}^{eq}$, so we assume elimination of imaginaries. In the context where we would like to prove elimination of imaginaries, this means that we have singled out a subset of the sorts, which we call the real sorts, and would like to show, e.g., that any element is inter-definable with a real one.

End lecture 13

Lecture 14, Sep. 28, 2009

3.1.17. *Weak elimination of imaginaries.* We would like to make use of the fact that a theory codes finite sets. We say that a theory $\mathcal{T}$ has *weak elimination of imaginaries* if it eliminates imaginaries up to finite sets, i.e., any element is inter-definable with the code for a finite set of real elements. Thus, a theory eliminates imaginaries precisely if it weakly eliminates imaginaries and also codes finite sets. In many cases, these two properties are dealt with separately.

If E_1 and E_2 are definable equivalence relations on X_1 and X_2, we say that E_1 and E_2 are *equivalent* if there is a definable equivalence relation E on $X_1 \coprod X_2$ whose restriction to $X_i \times X_i$ is E_i, and such that each element of X_1 is related to some element of X_2, and conversely. Thus, E_1 has a quotient precisely if it is equivalent to the equality on some X_2, in which case the quotient is E.

3.1.18. *Proposition.* Let $\mathcal{T}$ be a complete theory, M a saturated model. The following are equivalent:

(1) Any equivalence relation is equivalent to an equivalence relation with finite classes.
(2) Any M-definable set $\mathbf{X}$ can be defined over $\mathrm{acl}(\mathbf{X})$.
(3) If p is the type of an element of an equivalence class C over $\mathrm{acl}(C)$, then $p(M) \subseteq C(M)$.
(4) Any element e of M^{eq} is definable over the real elements of $\mathrm{acl}(e)$.

[3] the completeness assumption is only for convenience. There are analogous statements for arbitrary theories, which are more cumbersome to state

(5) $\mathcal{T}$ weakly eliminates imaginaries

3.1.19. *Exercise.* Deduce 3.1.18 from 3.1.13.

3.2. **Stability.** There are several ways to define stability. From our point of view, one of the useful properties of fields was the existence of tensor products. Stability generalises a variant of this that we now try to describe.

Let $\Bbbk$ be a subfield of an algebraically closed field M (viewed as a model of ACF), and let A be an extension field of $\Bbbk$ generated (as a field) over $\Bbbk$ be a tuple a. There is then a type p_a over $\Bbbk$, such that $x \in M$ realises p_a precisely if $a \mapsto x$ determines a map of fields from A to M. Conversely, any type over $\Bbbk$ determines, up to isomorphism, a field extension of $\Bbbk$. Thus we may replace a field extension of $\Bbbk$ (together with a fixed generator) by a type over $\Bbbk$.

Now let A and B be two field extensions of $\Bbbk$, generated by a and by b. If we are in the situation of 2.3.8, then $A \otimes_{\Bbbk} B$ is a domain, and the field of fractions $K(A \otimes_{\Bbbk} B)$ is generated by a, b. Thus we get an operation that to any two types p and q over $\Bbbk$ attaches a new type $p \otimes q$. To understand the nature of this type, we note that if we fix a map from B to M (i.e., a realisation the type q), a map over $\Bbbk$ from A to M is the same as a map over B from $A \otimes_{\Bbbk} B$ to M. This map extends to a map on the fraction field precisely if it is injective. Thus, a realisation of $p \otimes q$ is the same as a realisation of p which is "independent" from B.

It turns out that the existence of such an operation on types (over algebraically closed sets) is a characterisation of stable theories (Kim and Pillay [22]). We will define give a definition of stability from which this operation is easy to describe. However, some properties that are obvious in the special case of tensor products, such as symmetry, will require some work to prove.

There are several books available about stability, taking several approaches, among them Buechler [7], Pillay [32], Baldwin [2], Lascar [24] and Shelah [37]. There are also the lecture notes Pillay [34].

End lecture 14
Lecture 15, Oct. 2, 2009

3.2.1. Let $\mathcal{T}$ be a complete and model complete theory with EI. In this section, by a *structure* we mean a definably closed subset of some model. If A is a structure, and $\mathbf{X}$ is an A-definable set, we denote by $\mathcal{D}_A(\mathbf{X})$ the boolean algebra of A-definable subsets of $\mathbf{X}$, and by $\mathcal{S}_A(\mathbf{X})$ the space of types over A of elements of $\mathbf{X}$. We omit the A from the notation if $A = \mathrm{dcl}(0)$.

3.2.2. *Definition.* Let $\mathbf{X}$ be a definable set, A a set of parameters. A *definable type* over A in $\mathbf{X}$ is a collection of boolean algebra homomorphisms $\mathfrak{d} : \mathcal{D}(\mathbf{X} \times \mathbf{Y}) \to \mathcal{D}_A(\mathbf{Y})$, one for every definable set $\mathbf{Y}$, such that for any set B containing A, the collection of definable subsets of $\mathbf{X}$ given by

$$\mathfrak{d}|_B = \{\mathbf{Z}_b : \mathbf{Z} \subseteq \mathbf{X} \times \mathbf{Y}, b \in \mathfrak{d}(\mathbf{Z})(B)\} \tag{1}$$

is consistent. We will call this set the restriction of $\mathfrak{d}$ to B.

3.2.3. *Exercise.* Show that for any set B, $\mathfrak{d}|_B$ is a type over B. Hence, given a type p over B, it makes sense to ask whether it has the form $\mathfrak{d}|_B$ for some B-definable type. We say that the type p itself is definable in this case.

3.2.4. *Definition.* The theory $\mathcal{T}$ is called *stable* if for any algebraically closed set A, any type over A is definable (over A). A definable type whose restriction to A is a given type p is called a *definition scheme* for p.

3.2.5. *Example.* In ACF, a type p over A may be identified (but only if A is algebraically closed!) with the generic point of some A-sub-variety $\mathbf{V}$ of affine space. Thus, a definable set given by polynomial equations will belong to p precisely if the algebraic subset it determines contains $\mathbf{V}$. This is a definable condition (over A), which determines a definable type $\mathfrak{d}_p$ such that $\mathfrak{d}|_A = p$. This shows that ACF is stable.

3.2.6. *Exercise.* Let p be the type in ACF_0 corresponding to the variety given by $y^2 = x^3$. Find a quantifier free presentation of $\mathfrak{d}_p\phi$ for the following formulas $\phi(xy, a, \dots)$:

$$x = ay \tag{2}$$

$$ax^3 + by^2 + cx + dy = 0 \tag{3}$$

$$a(x^2y^2 - x^5) + b(y^3 - yx^3) = 0 \tag{4}$$

$$\begin{split} b^2 = a^3 \wedge \exists z, w(z^2 = w^3 \wedge (z, b) \neq (w, a) \wedge \\ (b - y)(a - w) = (b - z)(a - x)) \end{split} \tag{5}$$

(The last one requires a bit geometry; make a picture)

3.2.7. The definition scheme $\mathfrak{d}$ of a type p over A gives a canonical way of extending p to larger sets of parameters B. This will be our notion of "the generic extension" of p to B. To show that this is indeed a canonical extension, we need to show that each type has a *unique* definition scheme. We would also like to prove that this notion satisfies some natural properties, similar to the ones for tensor products. It seems that there is not direct way to prove these, and all these statements make use of the order property.

3.2.8. *The order property.* A definable set $Z \subseteq X \times Y$ is said to have the *order property* if there are infinite sequences $a_i \in X(M)$ and $b_i \in Y(M)$ (for some model M), such that $(a_i, b_j) \in Z(M)$ if and only if $i \leq j$. A theory $\mathcal{T}$ has the order property is some definable set of $\mathcal{T}$ has it.

3.2.9. *Exercise.* Show that $\mathcal{T}$ has the order property if and only if there is a definable set as above with $X = Y$ and $a_i = b_i$. In other words, the map $I : \mathbb{N} \to X(M)$ extends to an isomorphism of structures $(\mathbb{N}, <) \to (I(\mathbb{N}), Z)$.

End lecture 15

Lecture 16, Oct. 5, 2009

3.2.10. *Lemma.* Let κ be an infinite cardinal. Then there is a linearly ordered set S of cardinality at most κ, whose set of cuts has cardinality bigger than κ.

Proof. Let α be the smallest cardinal such that $2^\alpha > \kappa$, and let S be the subset of 2^α of functions f for which there is a $\beta < \alpha$ such that $f(i) = f(j)$ for $i, j > \beta$ (eventually constant functions). The set S is ordered lexicographically. We claim that S has cardinality at most κ. Indeed, an element of S is given by an element of 2^β for some $\beta < \alpha$ (and an element of 2). The number of such elements is less than κ by the choice of α. However, each element of 2^α realises a different cut. □

3.2.11. *Lemma.* If $\mathcal{T}$ has the order property, then for any cardinal $\kappa \geq \#\mathcal{T}$, there is a set S of cardinality κ such that there are more than κ types over S.

Proof. Let S be a set as in 3.2.10, and add a constant for each element of S. If Z is as in 3.2.9, add the axioms $(c,d) \in Z \wedge (d,c) \notin Z$ for any $c,d \in S$ with $c < d$. Then by the choice of Z these axioms are finitely consistent, hence consistent. But each cut in S extends to at least one type over S (again by the choice of Z). □

3.2.12. *Corollary.* A stable theory does not have the order property

Proof. Exercise:

(1) Show that 3.2.11 is true even if we require that S is algebraically closed.
(2) Show that if $\kappa = \kappa^{\#\mathcal{T}}$ and $\mathcal{T}$ is stable, then there are only κ types over any algebraically closed set of size κ.

□

3.2.13. *Exercise.* Consider a language with two sorts $\mathbf{X}$ and $\mathbf{P}$, and a relation $\mathbf{E}$ on $\mathbf{X} \times \mathbf{P}$. Let S be an infinite set with power set $P(S)$, and let $\mathcal{T}$ be the theory of $S, P(S)$ in this language, where $\mathbf{E}$ is interpreted as membership. Show that $\mathcal{T}$ is not stable.

3.2.14. *Exercise.* The following is from Duret [14]. The following algebraic fact is proved there: Let $\mathbb{k}$ be a field of characteristic not 2, $a_i, b_i \in \mathbb{k}$ elements such that the a_i are distinct and the b_i non-zero. Then the algebraic set (over the algebraic closure of $\mathbb{k}$) given by $y_i^2 = b_i(x - a_i)$ is irreducible. Use this to show that $ACFA$ is not stable. See Beyarslan [4] for a generalisation.

We note that the above actually proves that the fixed field (and therefore $ACFA$) has the *independence property*: There is a definable set $\mathbf{Z} \subseteq \mathbf{X} \times \mathbf{Y}$ and elements $a_i \in \mathbf{X}$ and $b_i \in \mathbf{X}$ such that for any subset B of the b_i, there is an a_i such that $(a_i, b_j) \in \mathbf{Z}$ precisely if $b_j \in B$. Obviously, having the independence property is stronger than having the order property, so the class NIP of theories that do not have the independence property contains the stable theories, and the example above shows that $ACFA$ is not NIP.

3.2.15. We will write a definable type as $\mathfrak{d}^{\mathbf{X}}$ if we wish to specify that it defines a type in $\mathbf{X}$. By definition, an A-definable type consists of homomorphisms $\mathfrak{d} : \mathcal{D}(\mathbf{X} \times \mathbf{Y}) \to \mathcal{D}_A(\mathbf{Y})$. However, $\mathfrak{d}$ canonically extends to a homomorphism $\mathfrak{d} : \mathcal{D}_A(\mathbf{X} \times \mathbf{Y}) \to \mathcal{D}_A(\mathbf{Y})$, via $\mathfrak{d}(\mathbf{Z}_a) = \mathfrak{d}(\mathbf{Z})_a$. As a result, it makes sense to compose two definable types: if $\mathfrak{d}^{\mathbf{X}}$ and $\mathfrak{d}^{\mathbf{Y}}$ are two definable types, we get for each definable set $\mathbf{Z}$ a map $\mathfrak{d}^{\mathbf{X}} \otimes \mathfrak{d}^{\mathbf{Y}} : \mathcal{D}_A(\mathbf{X} \times \mathbf{Y} \times \mathbf{Z}) \to \mathcal{D}_A(\mathbf{Z})$. The following result will imply all properties of stable theories that we need.

3.2.16. *Main lemma.* Any two definable types $\mathfrak{d}^{\mathbf{X}}$ and $\mathfrak{d}^{\mathbf{Y}}$ over A commute: $\mathfrak{d}^{\mathbf{X}} \otimes \mathfrak{d}^{\mathbf{Y}} = \mathfrak{d}^{\mathbf{Y}} \otimes \mathfrak{d}^{\mathbf{X}}$.

Proof. Assume otherwise. By extending parameters we may assume that there is an A-definable subset $\mathbf{Z} \subseteq \mathbf{X} \times \mathbf{Y}$ such that both $(\mathfrak{d}^{\mathbf{X}} \otimes \mathfrak{d}^{\mathbf{Y}})\mathbf{Z}$ holds, but $(\mathfrak{d}^{\mathbf{Y}} \otimes \mathfrak{d}^{\mathbf{X}})\mathbf{Z}$ doesn't. This means, in particular that $\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$ (a subset of $\mathbf{X}$) is consistent, i.e., $x_0 \in \mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$ for some $x_0 \in \mathbf{X}$. This, in turn, means that $(x_0, y_0) \in \mathbf{Z}$ for some $y_0 \in \mathbf{Y}$. Turning now to the fact that $(\mathfrak{d}^{\mathbf{Y}} \otimes \mathfrak{d}^{\mathbf{X}})\mathbf{Z}^c$ also holds, we note that $y_0 \in \mathfrak{d}^{\mathbf{X}}\mathbf{Z}^c$ (since y_0 satisfies the type over A determined by $\mathfrak{d}^{\mathbf{Y}}$). Hence we may find x_1 such that $(x_1, y_0) \notin \mathbf{Z}$. This x_1 satisfies $\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$, hence there is $y_1 \in \mathbf{Y}$ such that $(x_1, y_1) \in \mathbf{Z}$.

Furthermore, since x_0 also satisfies $\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$, we also get that $(x_0, y_1) \in \mathbf{Z}$. So far we have:

$$(x_0, y_0) \in \mathbf{Z} \tag{6}$$

$$(x_1, y_0) \notin \mathbf{Z} \tag{7}$$

$$(x_0, y_1) \in \mathbf{Z} \tag{8}$$

$$(x_1, y_1) \in \mathbf{Z} \tag{9}$$

By induction, assume we have defined, for $i < n$, $x_i \in \mathbf{X}$ and $y_i \in \mathbf{Y}$ such that $(x_i, y_j) \in \mathbf{Z}$ if and only if $j \geq i$, each x_i satisfies $\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$ and each $y_i \in \mathfrak{d}^{\mathbf{X}}\mathbf{Z}^c$. We define x_n to be an element satisfying $\mathfrak{d}^{\mathbf{X}}|_{\{y_i : i<n\}}$ and y_n to be an element satisfying $\mathfrak{d}^{\mathbf{Y}}|_{\{x_i : i \leq n\}}$. Since all $y_i \in \mathfrak{d}^{\mathbf{X}}\mathbf{Z}^c$, we have $(x_n, y_i) \notin \mathbf{Z}$ for $i < n$, and since all x_i satisfy $\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$, we get that $(x_i, y_n) \in \mathbf{Z}$ for $i \leq n$.

This shows that $\mathbf{Z}$ has the order property, contradicting 3.2.12 and the stability of $\mathcal{T}$. □

End lecture 16
Lecture 17, Oct. 7, 2009

3.2.17. *Exercise.* Show that if $\mathfrak{d}^{\mathbf{X}}$ and $\mathfrak{d}^{\mathbf{Y}}$ are A-definable types, then so is $\mathfrak{d} = \mathfrak{d}^{\mathbf{X}} \otimes \mathfrak{d}^{\mathbf{Y}}$. Show that (x, y) satisfies $\mathfrak{d}|_B$ if and only if x satisfies $\mathfrak{d}^{\mathbf{X}}|_B$ and y satisfies $\mathfrak{d}^{\mathbf{Y}}|_{B \cup \{x\}}$.

3.2.18. *Corollary.* If $\mathfrak{d}_1$ and $\mathfrak{d}_2$ are two definable types over A such that $\mathfrak{d}_1|_A = \mathfrak{d}_2|_A$, then $\mathfrak{d}_1 = \mathfrak{d}_2$.

Proof. Let $\mathbf{Z} \subseteq \mathbf{X} \times \mathbf{Y}$ and let $y \in \mathbf{Y}$ be such that $\mathfrak{d}_1\mathbf{Z}(y)$ holds. Let $\mathfrak{d}^{\mathbf{Y}}$ be a defining scheme for the type of y over A. Then $\mathfrak{d}^{\mathbf{Y}}\mathfrak{d}_1\mathbf{Z}$ holds, hence, by 3.2.16, so does $\mathfrak{d}_1\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$. But $\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$ is an A-definable set, so by assumption, $\mathfrak{d}_2\mathfrak{d}^{\mathbf{Y}}\mathbf{Z}$, hence, again by 3.2.16, $\mathfrak{d}_2\mathbf{Z}(y)$ holds as well. □

3.2.19. To summarise, we now know that in a stable theory, any type over an algebraically closed set is definable by a unique definable type. We may thus ignore the distinction between types and definable types. In particular, if p and q are types over an algebraically closed set A, we denote by $p \otimes q$ (the free amalgamation of p and q) the type over A corresponding to the composition of the associated definable types. Likewise, if $B \supseteq A$, we denote by $p|_B$ the B points of the definable type of p. This type is called the (unique) *non-forking extension* of p to B. Thus, 3.2.17 gives that (a, b) realises $p \otimes q$ if and only if a satisfies p, and b satisfies $q|_a$. We now show in several examples that $p \otimes q$ is indeed a "free" amalgamation of p and q.

3.2.20. *Proof of 2.3.8 when* $\Bbbk = \mathrm{acl}(\Bbbk)$. The two domains A and B correspond to types p and q over $\Bbbk$ in ACF. Since $\Bbbk$ is algebraically closed and ACF is stable, we may form $p \otimes q$. Let C be the field generated by a realisation. Then we have maps over $\Bbbk$ from A and B into C, hence a map $A \otimes B \to C$. We claim that the map is injective. Indeed, we have an identity $\phi(a, b) : \sum a_i b_i = 0$, where the a_i are, without loss of generality, linearly independent over $\Bbbk$, then b satisfies $\mathfrak{d}_p\phi(y)$, a formula over $\Bbbk$. Hence if b is non-zero, there is a non-zero tuple $c \in \Bbbk$ satisfying $\mathfrak{d}_p\phi$ as well, contradicting the linear independence of a.

3.2.21. *Proposition.* Let $p(x)$ and $q(y)$ be two types of $A = \mathrm{acl}(A)$, and let $f(x)$ and $g(y)$ be two A-definable functions. Then the formula $\phi(x, y)$ given by $f(x) = g(y)$ belongs to $p \otimes q$ if and only if there is $a \in A$ such that $f(x) = a$ is in p and $g(y) = a$ is in q.

Proof. Assume that $f(x) = g(y)$ is in $p \otimes q$, let x_0 realise p, and let $a = f(x_0)$. Then x_0 satisfies $d_q\phi$, and so the formula $g(y) = a$ is in $q|_{x_0}$. But a can be the only z such that $g(y) = z$ is in q. Therefore, a is definable over A, and hence belongs to A. □

End lecture 17
Lecture 18, Oct. 9, 2009

3.2.22. *Exercise.* Let $A = \mathrm{acl}(A)$, let $\mathbf{X}$ and $\mathbf{Y}$ be two definable sets, let $f : \mathbf{X} \to \mathbf{Y}$ be an A-definable function, and let, for each $\mathbf{V}$, $f^{-1} : \mathcal{D}_A(\mathbf{Y} \times \mathbf{V}) \to \mathcal{D}_A(\mathbf{X} \times \mathbf{V})$ be the inverse image map. If $\mathfrak{d}^{\mathbf{X}}$ is an A-definable type on $\mathbf{X}$, show that $f_*(\mathfrak{d}^{\mathbf{X}})$ given by $f_*(\mathfrak{d}^{\mathbf{X}})(\mathbf{Z}) = \mathfrak{d}^{\mathbf{X}}(f^{-1}(\mathbf{Z}))$ is an A-definable type on $\mathbf{Y}$. Show that if x realises $\mathfrak{d}^{\mathbf{X}}|_B$, then $f(x)$ realises $f_*(\mathfrak{d}^{\mathbf{X}})|_B$. Show that if $\mathfrak{d}_1$ and $\mathfrak{d}_2$ are A-definable types, and f, g are A-definable functions on the corresponding domains, then $f_*(\mathfrak{d}_1) \otimes g_*(\mathfrak{d}_2) = (f, g)_*(\mathfrak{d}_1 \otimes \mathfrak{d}_2)$.

3.2.23. *Exercise.* Let $A = \mathrm{acl}(A)$, and let σ be an automorphism of A. For any type p over A, define $\sigma(p)$ to be the type $\{Z_{\sigma(a)} : Z_a \in p\}$. Show that if p and q are types over A, then $\sigma(p \otimes q) = \sigma(p) \otimes \sigma(q)$ (Hint: use uniqueness).

3.2.24. *Definition.* Let A be an algebraically closed set in a stable theory, and let b and c be two elements in a model containing A. We say that b is *free* (or *independent*) from c over A, denoted by $b \mathop{\smile\!\!\!\!\!\!|}_A c$, if $\mathrm{tp}((b, c)/A) = \mathrm{tp}(b/A) \otimes \mathrm{tp}(c/A)$. If B and C are (possibly infinite) subsets of a model, we say that B is free from C over A if $b \mathop{\smile\!\!\!\!\!\!|}_A c$ for any finite tuples $b \in B$ and $c \in C$.

3.2.25. *Exercise.* Prove the following properties of the freeness relation

(1) If $b \mathop{\smile\!\!\!\!\!\!|}_A c$, b_1 is definable over b and c_1 is definable over c, then $b_1 \mathop{\smile\!\!\!\!\!\!|}_A c_1$. Hence the definition above for (possibly) infinite sets makes sense.
(2) If $B \mathop{\smile\!\!\!\!\!\!|}_A C$ then $C \mathop{\smile\!\!\!\!\!\!|}_A B$.
(3) $B \mathop{\smile\!\!\!\!\!\!|}_A D$ if and only if $B \mathop{\smile\!\!\!\!\!\!|}_A C$ and $B \mathop{\smile\!\!\!\!\!\!|}_C D$ for any (algebraically closed) $C \subseteq D$.
(4) For any A and B there is a C-isomorphic copy A' of A over C such that $A' \mathop{\smile\!\!\!\!\!\!|}_C B$.

3.2.26. *Exercise.* In terms of the freeness relation, 3.2.21 says that if b and c are free over $A = \mathrm{acl}(A)$, then $\mathrm{dcl}(A, b) \cap \mathrm{dcl}(A, c) = A$. Generalise this statement by replacing dcl with acl.

3.2.27. *Canonical bases.* As with M-definable sets, we may ask what is the minimal set over which a type $p(x)$ over A is definable. We would like to say, as in 3.1.7, that $\mathrm{dcl}(p)$ is the set of elements fixed by automorphisms fixing p. However, note that in 3.1.7 we considered automorphisms fixing the M-points of a definable set. Likewise, we define the *canonical base* $\mathcal{C}b(p)$ of the type p (over $A = \mathrm{acl}(A)$) to be the set of all elements of M fixed by all automorphisms of M fixing $p|_M$ (as a type). Clearly, $\mathcal{C}b(p) \subseteq A$.

3.2.28. *Proposition.* Let p be a type over $A = \text{acl}(A)$, and let q be the restriction of p to $\mathcal{C}b(p)$. Then:

1. $\mathcal{C}b(p)$ is the smallest definably closed subset of A over which p is defined
2. For any extension r of q to an algebraically closed set D, $\mathcal{C}b(r) \supseteq \mathcal{C}b(p)$, and $\mathcal{C}b(r) = \mathcal{C}b(p)$ precisely if $r = p|_D$.

Proof.
1. We already noticed that $\mathcal{C}b(p)$ is contained in any set over which p is defined. On the other hand, if $\sigma(p|_M) = p|_M$, then for any ϕ, $\sigma(\mathfrak{d}_p\phi(M)) = \mathfrak{d}_p\phi(M)$, so by 3.1.13, $\mathfrak{d}_p\phi$ is definable with a parameter in $\mathcal{C}b(p)$.
2. Assume that σ fixes $\mathcal{C}b(r)$ pointwise. Then it fixes $r|_M$, so $r \cup \sigma(r)$ is consistent. Therefore, also $p \cup \sigma(p)$ is consistent, so σ fixes $\mathcal{C}b(p)$ pointwise.

   For the second statement, r and $p|_D$ are two types which are definable over A and agree on the restriction to A, so by uniqueness they are equal. Then converse is obvious.

□

End lecture 18
Lecture 19, Oct. 12, 2009

3.2.29. *Non algebraically closed sets.* So far we have concentrated on types over algebraically closed sets. We now know that any such type (in a stable theory) has a unique definition scheme. Over arbitrary sets, either uniqueness or existence fails (depending on the point of view):

3.2.30. *Example.* Let p be the type determined by $x^2 = 2$ in ACF_0. Any extension of this type to $\mathbb{Q}^a$ will have to include $x = a$ for some square root a of 2. Since there are two of them which are conjugate, there is no definition scheme over $\mathbb{Q}$, and over $\mathbb{Q}^a$ there are two definition schemes that agree with p when restricted to $\mathbb{Q}$. We note that all problems appear already when we extend to $\mathbb{Q}(\sqrt{2})$, which happens to be the canonical base of each extension (this is an instance of the above proposition). On the other hand, if we extend to any field that does not contain a square root of 2, no problems occur.

3.2.31. *Canonical bases (general case).* The above example is typical. Given a type p over an arbitrary set A (which we may safely assume to be at least definably closed), there is a subset $\mathcal{C}b(p)$ of A, such that p can be canonically extended from $\mathcal{C}b(p)$ to any "non-problematic" set. More precisely, we have:

3.2.32. *Proposition.* Let $\mathcal{T}$ be a stable theory. For each definably closed set A there is a map $\mathcal{C}b$ from the space of types over A to the set of subsets of A, with the following properties:

1. For any automorphism σ, $\sigma(\mathcal{C}b(p)) = \mathcal{C}b(\sigma(p))$.
2. (Existence) Let $B \supseteq A$ be such that $B \cap \text{acl}(\mathcal{C}b(p)) \subseteq A$. Then there is an extension q of p to B, such that $\mathcal{C}b(q) = \mathcal{C}b(p)$.
3. (Uniqueness) For any $B \subseteq A$ containing $\mathcal{C}b(p)$, p is the unique extension of $p|_B$ with canonical base contained in B.

Proof. When A is algebraically closed, the canonical base was already defined. Let p be a type over A, and let p_1 be any extension of p to $\text{acl}(A)$. Any element b of $\mathcal{C}b(p_1)$ is algebraic over A, and we denote by $[b]$ the (element coding the) finite set of conjugates of b over A (thus $[b] \in A$). We define $\mathcal{C}b(p) = \{[b] : b \in \mathcal{C}b(p_1)\}$. Hence $\mathcal{C}b(p) \subseteq A$. Note that this does not depend on the choice of p_1, since all such types (and therefore their canonical bases) are conjugate over A.

(1) Exercise (the case when A is algebraically closed is in 3.2.23).

(2) Let p_1 be an extension of p to $\operatorname{acl}(A)$, let q_1 be the non-forking extension of p_1 to $\operatorname{acl}(B)$, and let q be the restriction of q_1 to B. Since, by definition, $\mathcal{C}b(p_1) = \mathcal{C}b(q_1)$, each element c of $\mathcal{C}b(q)$ is a finite set of elements, each of which is algebraic over $\mathcal{C}b(p)$. Therefore c itself is algebraic over $\mathcal{C}b(p)$. Thus, $\mathcal{C}b(q) \subseteq \operatorname{acl}(\mathcal{C}b(p))$. On the other hand, $\mathcal{C}b(q) \subseteq B$, so by assumption, $\mathcal{C}b(q) \subseteq A$. This means that the code for the set of conjugates of an element b of $\mathcal{C}b(p_1) = \mathcal{C}b(q_1)$ over B is in A. Therefore, the sets of conjugates over B and over A are equal, so $\mathcal{C}b(p) = \mathcal{C}b(q)$.

(3) Assume that q is another such extension. Extend p and q to p_1 and q_1 over $\operatorname{acl}(A)$. and let p_2 and q_2 be the restrictions of these types to $\operatorname{acl}(B)$. p_2 and q_2 restrict to the same type over B, so there is an automorphism τ over B taking p_2 to q_2.

By the assumption, the canonical base of p_1 is contained in $\operatorname{acl}(B)$, so p_1 is the unique non-forking extension of p_2 to $\operatorname{acl}(A)$. Therefore, τ takes $\mathcal{C}b(p_2)$ to a conjugate over A. By composing with an automorphism over A, we may thus assume that τ fixes $\mathcal{C}b(p_2)$. But then $p_1 = q_1$ and therefore $q = p$.

□

3.2.33. *Exercise.* Show that in the existence part above, it was enough to assume that $B \cap \operatorname{dcl}(\bigcup_r \mathcal{C}b(r)) \subseteq A$, where the union is over all extensions r of p to $\operatorname{acl}(A)$.

3.2.34. *Proof of 2.3.8, general case.* Similar to the algebraically closed case. View A as a substructure extending $\Bbbk$, and B as a type p over $\Bbbk$. By the proposition above and the assumption, there is a non-forking extension of q of p to A. As before, if $\sum a_i x_i = 0$ is in q for some tuple a in A, then , since p and q have the same canonical base, there is a tuple c_i in $\Bbbk^a$ with the same property. Furthermore, the space of all such tuples is defined over $\Bbbk$. Therefore, it has a non-zero $\Bbbk$-point (this is Hilbert 90).

3.2.35. *Facts.* The following basic facts about stability will not be proved (or used) here. The converse of 3.2.12 is true: If a theory does not have the order property, then it is stable. It follows as before that $\mathcal{T}$ is stable if for some cardinal κ, there are only κ types over any set of cardinality κ ($\mathcal{T}$ is said to be *κ-stable* in this case). It also follows that any complete theory interpretable in a stable theory is itself stable.

Furthermore, for any definable type $\mathfrak{d}$ and any definable set $\phi(x, y)$, $\mathfrak{d}\phi$ is in fact a boolean combination of set $\phi(c_i, y)$, where the c_i are independent realisations of $\mathfrak{d}$. Cf. Pillay [32, ch. 1].

End lecture 19
Lecture 20, Oct. 16, 2009

4. Theories of a generic automorphism

In this section we present an abstract version of the situation in $ACFA$. The main reference is Chatzidakis and Pillay [11]. The idea is that $ACFA$ is obtained from the theory ACF be adding an automorphism, and then taking the model companion. This makes sense with any theory: Starting with a theory $\mathcal{T}$, we may add a function symbol σ, and the axioms stating that σ is an automorphism of the $\mathcal{T}$ structure, and then ask for the model companion $\mathcal{T}_\sigma$. It turns out that in many cases it exists. One may then ask what are the consequences of good properties of

$\mathcal{T}$ for $\mathcal{T}_\sigma$. The main good property $\mathcal{T}$ may have in this context is stability, and in this context one may deduce most of the statements of 2.3 (and more). One of the main goals we have here is elimination of imaginaries for $ACFA$.

4.1. **Basic properties.**

4.1.1. *Definition.* Let $\mathcal{T}$ be a theory with elimination of quantifiers and elimination of imaginaries. We let $\mathcal{T}_\sigma^0$ be the theory in the language of $\mathcal{T}$ expanded by a unary function symbol σ, whose restriction to the language of $\mathcal{T}$ is the universal part of $\mathcal{T}$, and that says that σ is an injective endomorphism of the structure of $\mathcal{T}$.

The theory $\mathcal{T}_\sigma$ of a *generic automorphism of* $\mathcal{T}$ is the model companion of $\mathcal{T}_\sigma^0$.

4.1.2. *Example.* For $\mathcal{T} = ACF$ we have $\mathcal{T}_\sigma = ACFA$ (2.3.3)

4.1.3. *Remark.* The theory $\mathcal{T}_\sigma$ need not exist, but by 2.1.16 if it exists, it is unique. The question of existence is not easy. Chatzidakis and Pillay [11] show that if definable sets have a reasonable notion of dimension (finite Morley rank, definability of the Morley degree, "dimension theorem"), then $\mathcal{T}_\sigma$ exists, and can be axiomatised similarly to $ACFA$. Baldwin and Shelah [3] describe a combinatorial necessary and sufficient condition on $\mathcal{T}$, when $\mathcal{T}$ is stable. We will address this question, since we are interested in $ACFA$, which already exists. *from now on, we assume that $\mathcal{T}_\sigma$ exists.*

4.1.4. *Exercise.* Show that $\mathcal{T}_\sigma$ implies that σ is an automorphism (i.e., surjective).

4.1.5. *Exercise.* Describe $\mathcal{T}_\sigma$ when $\mathcal{T}$ is the theory of equality (in an infinite set). More precisely, show that $\mathcal{T}_\sigma$ admits elimination of quantifiers, describe the types, and show that it is ω-stable.

We now assume that $\mathcal{T}$ is stable, and revisit some of the results we proved for $ACFA$.

4.1.6. *Theorem (compare 2.3.9).* Assume $\mathcal{T}$ is stable. Let (M_1, σ_1) and $(M_2.\sigma_2)$ be two models of $\mathcal{T}_\sigma$, and let (E, σ) be a substructure of both, algebraically closed as a $\mathcal{T}$-structure. Then M_1 and M_2 have the same theory over E.

Proof. Let p_1 and p_2 be the types of M_1 and M_2 over E, in the sense of $\mathcal{T}$. Let M be the definable closure of a realisation of $p_1 \otimes p_2$. We view M_1 and M_2 as substructures of M. We claim that the union of σ_1 and σ_2 defines an automorphism of the union $M_1 \cup M_2$. Indeed, if $x_i \in M_i$, then $\sigma_i(x_i)$ satisfies $\sigma(\mathrm{tp}(x_i/E))$, so $(\sigma_1(x_1), \sigma_2(x_2))$ satisfies $\sigma(\mathrm{tp}(x_1/E)) \otimes \sigma(\mathrm{tp}(x_2/E))$ (since both of these types are restrictions of the p_i, cf 3.2.25). By 3.2.23, the last type is $\sigma(\mathrm{tp}(x_1/E) \otimes \mathrm{tp}(x_2/E))$. This shows that $\sigma = \sigma_1 \cup \sigma_2$ is an automorphism of M.

By the definition of $\mathcal{T}_\sigma$, (M, σ) embeds in a model $\tilde{M}$ of $\mathcal{T}_\sigma$. Furthermore, since $\mathcal{T}_\sigma$ is model-complete, the inclusions of M_i in $\tilde{M}$ are elementary, so they all have the same theory. □

4.1.7. *Corollary (compare 2.3.17).* Assume that M is a model of $\mathcal{T}_\sigma$, where $\mathcal{T}$ is stable. The type of $a \in M$ over A is determined by the quantifier free type of $\mathrm{acl}_\sigma(A(a))$.

Proof. The same as in 2.3.17 (exercise). □

4.1.8. *Proposition (compare 2.3.15).* Let A be a $\mathcal{T}_\sigma$ sub-structure of a model M, algebraically closed as a $\mathcal{T}$ structure ($\mathcal{T}$ stable). Then A is algebraically closed in $\mathcal{T}_\sigma$.

Proof. As in the proof of 4.1.6, let p be the type of M over A in the sense of $\mathcal{T}$, and let N be a model of $\mathcal{T}_\sigma$ in which $p \otimes p$ is realised. Thus, we get two embeddings $f_i : M \to N$.

Now let a be the tuple of conjugates over A of some $\mathcal{T}_\sigma$-algebraic element a_0, and let $b^i = f_i(a)$. Since the embeddings are elementary, all the components b^i_j are conjugate. Hence for each j there is some k such that $b^1_j = b^2_k$. By 3.2.21, each $b^i_j \in A$. □

End lecture 20
Lecture 21, Oct. 26, 2009

4.1.9. *Exercise (generalising 4.1.5).* Assume that $\mathcal{T}$ is stable and has the property that for any subset A of a model, $\operatorname{acl}(A) = \operatorname{dcl}(A)$. Show that $\mathcal{T}_\sigma$ eliminates quantifiers and is again stable.

In contrast, we have:

4.1.10. *Proposition.* Assume that there are algebraically closed $\mathcal{T}$ structures A and B such that $\operatorname{acl}(A \cup B) \neq \operatorname{dcl}(A \cup B)$. Then $\mathcal{T}_\sigma$ has the independence property.

Proof. As in 3.2.14, we will show that already the fixed set $\mathbf{C}$ has the independence property. Let r be an element algebraic but not definable over $A \cup B$. So there is a formula $\phi(z, x, y)$ (of $\mathcal{T}$) and tuples $a \in A$ and $b \in B$, such that any element s satisfying $\phi(s, a, b)$ satisfies the whole type of r over $A \cup B$.

Let p be the type of b over A. Given a natural number n, let $b_1, \dots, b_n$ be a realisation of $p^{\otimes n}$. We get, for each i, an element r_i algebraic but not definable over $A \cup b_i$. Hence we may find an automorphism of $B_i = \operatorname{acl}(A \cup b_i)$ fixing $A \cup b_i$ but not r_i. Given $I \subseteq \{1, \dots, n\}$, let σ_i be such an automorphism of B_i if $i \in I$, and the identity on B_i otherwise. As in the proof of 4.1.6, the automorphisms combine to give an automorphism σ on the definable closure D of the B_i, which is the identity on A and the b_i. This σ-structure embeds into a model M of $\mathcal{T}_\sigma$, so in M we may find b_i and for each subset I of $\{1, \dots, n\}$ an a_I, such that $\psi(a_I, b_j)$ precisely if $j \in I$, where $\psi(x, y)$ is the subset of $\mathbf{C} \times \mathbf{C}$ defined by $\exists z \in \mathbf{C}\phi(z, x, y)$. □

4.2. **Generic automorphisms.** In general, we have no explicit axiomatisation of $\mathcal{T}_\sigma$, so even though we assume it exists, we have no description of its models. In this section we give such a description: a model of $\mathcal{T}_\sigma$ is a model of $\mathcal{T}$ together with a *generic* automorphism.

For $ACFA$, the explicit axiomatisation is usually more useful than this description, so we will not use any of this for $ACFA$.

We begin with some general preliminaries.

4.2.1. *Proposition.* A theory $\mathcal{T}$ has a $\forall\exists$ axiomatisation if and only if its class of models is closed under unions of chains.

Proof. We may assume $\mathcal{T}$ is complete. Let M_0 be a model of $\mathcal{T}_{\forall\exists}$, and let $\mathcal{T}_1$ be the extension of $\mathcal{T}_{M_0}$ by all universal sentences over M_0. We claim that $\mathcal{T}_1$ is consistent: otherwise, $\mathcal{T}$ is inconsistent with a sentence of the form $\forall y \phi(y, a)$, where ϕ is a quantifier free formula over 0, and $a \in M_0$. Therefore, it is inconsistent with $\exists z \forall y \phi(y, z)$. But this is an $\exists\forall$ sentence that holds in M_0, hence is in $\mathcal{T}$.

Let N_1 be a model of $\mathcal{T}_1$. We view M_0 as a substructure of N_1. By the definition of $\mathcal{T}_1$, any universal sentence with parameters in M_0 that holds in M_0, holds in N_1. Now let $\mathcal{T}_2 = \mathcal{T}_{N_1} \cup \mathcal{T}_{M_0}(M_0)$ (i.e., the quantifier-free theory of N_1 in its language, extended by the full theory of M_0 in its language). We claim that $\mathcal{T}_2$ is consistent: If $\phi(x)$ is a quantifier free formula over M_0 such that $\phi(a)$ holds in N_1, then $\exists x \phi(x)$ holds there, whence it holds in M_0.

Let M_1 be a model of $\mathcal{T}_2$. Again we think of N_1 as a sub-model of M_1. So we have $M_0 \subseteq N_1 \subseteq M_1$, and the inclusion of M_0 in M_1 is elementary (since M_1 is a model of the full theory of M_0 in its language), and where N_1 is a model of $\mathcal{T}$. Repeating this construction, we get a chain, whose union is a model of $\mathcal{T}$, but is also an elementary extension of M_0. Therefore, M_0 is a model of $\mathcal{T}$. □

End lecture 21

4.2.2. *Corollary.* Any model-complete theory has a $\forall\exists$ axiomatisation (exercise).

Lecture 22, Oct. 28, 2009

4.2.3. *Proposition.* Let $\mathcal{T}$ be model complete. Then it is the theory of the existentially closed models of $\mathcal{T}_\forall$.

Proof. Let M be an existentially closed model of $\mathcal{T}_\forall$. By 4.2.2, we must prove M satisfies any axiom of the form $\forall x \exists y \phi(x,y)$ in $\mathcal{T}$. Let N be a model of $\mathcal{T}$ containing M, and let $a \in M$. Then in N there is an element b such that $\phi(a,b)$, hence, since M is existentially closed, such an element exists in M. □

We now go back to our situation.

4.2.4. *Definition.* Let M be a model of a theory $\mathcal{T}$, and let κ be a cardinal. We say that an automorphism σ of M is κ-generic if for any algebraically closed σ-substructure $A \subseteq M$, any σ-embedding $f : (A,\sigma) \to (B,\tau)$, where B is another algebraically closed σ-structure of cardinality less than κ, there is an embedding over (A,σ) of (B,τ) in (M,σ).

We will abbreviate "$\#\mathcal{T}^+$-generic" by "generic".

4.2.5. *Exercise.* Show that if σ is a κ-generic automorphism of M, then M is κ-saturated.

4.2.6. *Theorem.* Let $\mathcal{T}$ be stable, and assume $\mathcal{T}_\sigma$ exists. Then it is precisely the theory of the pairs (M,σ), where M is a model of $\mathcal{T}$, and σ is generic.

The proof is given in 4.2.9. We start by showing that the class of pairs (M,σ) with σ generic satisfies some of the properties of $\mathcal{T}_\sigma$. We assume that $\mathcal{T}$ is stable.

4.2.7. *Lemma.* Any σ-structure (M,σ) can be extended to a model with a generic automorphism

Proof. Embed M in any model of $\mathcal{T}$, and extend σ in any way. Thus we may assume that M is a model. Let A and B be as in the definition. Since A is algebraically closed, we may form $A_1 = M \otimes_A B$, with an automorphism on it. Embed A_1 in a model; continue. □

4.2.8. *Lemma.* Let (E, σ) be a common σ-substructure of (M_1, σ_1) and (M_2, σ_2), where E is algebraically closed, and σ_i are generic. Then the (M_i, σ_i) have the same theory over (E, σ).

Proof. This is a general method known as "back-and -forth". We will prove, by induction on the complexity of a formula over E, the following statement: Let A be an algebraically closed σ-substructure of M_1 of cardinality at most $\#\mathcal{T}$, and let $f : (A, \sigma_1) \to (M_2, \sigma_2)$ be an embedding over E, as guaranteed by the genericity of σ_2. Then $\mathrm{tp}_{M_1}(A/E) = \mathrm{tp}_{M_2}(f(A)/E)$. This will prove the claim since the type contains the whole theory.

The claim is obviously true for quantifier free formulas. Thus we need to show that if $\exists x \phi(x, a_1, \dots, a_n)$ holds in (M_1, σ_1) (where ϕ is a formula over E, and $a_i \in A$), then $\exists x \phi(x, f(a_1), \dots, f(a_n))$ holds in (M_2, σ_2). Let b be an element of M_1 such that $\phi(b, a_1, \dots, a_n)$ holds. Since A is algebraically closed, there is, by the genericity of σ_2, an extension of f to $\mathrm{acl}(b, A)$. Now by the induction hypothesis, $\phi(f(b), f(a_1), \dots, f(a_n))$. □

4.2.9. *Proof of 4.2.6.* Let M be a model of $\mathcal{T}$. By 4.2.3, we need to show that (M, σ) is existentially closed if and only if it is elementarily equivalent to (N, τ) with τ generic.

Assume that σ is generic. If $f : (M, \sigma) \to (N, \sigma)$ is any embedding (where (N, σ) is any σ-structure), by 4.2.7 we may embed (N, σ) in some σ-structure where σ is generic, and by 4.2.8 this embedding is elementary. Hence (M, σ) is existentially closed.

Conversely, if (M, σ) is a model of $\mathcal{T}_\sigma$, we may, again by 4.2.7, embed it in a model (N, τ) with τ generic. Since $\mathcal{T}_\sigma$ is the model completion, we may, in turn, embed N in a model of $\mathcal{T}_\sigma$. Continuing this way, we get a chain, where the composition of any two embeddings is elementary (by model completeness of $\mathcal{T}_\sigma$ and by 4.2.8). Therefore, the union is elementarily equivalent to both (M, σ) and (N, τ).

4.2.10. *Remark.* Hence if $\mathcal{T}_\sigma$ exists, then the theory of all generic automorphisms is model-complete. It is possible to prove that conversely, if this theory is model-complete, then $\mathcal{T}_\sigma$ exists (and therefore equals to it). The proof is an exercise, using 4.2.7 and 4.2.2.

4.2.11. *Exercise (compare 2.3.4).* Show that if (M, σ) is a model of $\mathcal{T}_\sigma$, then so is (M, σ^n) for any $n \neq 0$ (Hint: Assume for simplicity that $n = 2$. Let (A, τ) and (B, τ) be as in the definition, where on A, $\tau = \sigma^2$. Consider the type of B over A, and use stability. You may assume for simplicity that τ is the identity on A).

End lecture 22
Lecture 23, Oct. 30, 2009

4.2.12. *Exercise.* In this exercise we explain that generic automorphisms are indeed generic. This is taken from Lascar [23]. We will use the following definitions. Let X be a topological space, κ a cardinal. A subset A is *nowhere dense* if the interior of its closure is empty. It is called *κ-meagre* if it is the union of κ many nowhere dense sets, and *κ-comeagre* if its complement is κ-meagre. Finally, X is *κ-Baire* if any κ-meagre subset has empty interior. Thus, a Baire space is one in which "meagre" gives a reasonable notion of smallness. (The usual notions of meagre subset and Baire space are obtained when $\kappa = \aleph_0$.)

Fix a cardinal κ, and let G be a group acting faithfully on a set M. Call a subset of M small if its cardinality is less than κ. For an element $g \in G$ and a small subset

A, let

$$U_{g,A} = \{h \in G : h\restriction_A = g\restriction_A\}$$

(1) Show that the $U_{g,A}$ form a basis for a Hausdorff topology on G, so that G becomes a topological group.
(2) Show that a subset U of G is dense if and only if for any $g \in G$ and any small subset A, there is $h \in U$ with $h\restriction_A = g\restriction_A$.
(3) Show that if the cardinality of M is at most κ, then G is a κ-Baire space. (Hint: Show that the intersection of κ many dense open subsets V_α intersects any basic open subset $U_{g,A}$. You may assume $M = \kappa$, and define by induction a chain of partial maps g_α, $\alpha < \kappa$, such that $g_0 = g$, and g_α and its inverse are defined on α, each extends to an element of G, and all such extensions belong to V_α.)
(4) Let $A \subseteq B$ be small subsets, $g \in G$. Show that each of the following sets (hence their intersection $V_{g,A,B}$) is open:

$$\{h \in G : h\restriction_A \neq g\restriction_A\} \tag{10}$$

$$\{h \in G : u^{-1}hu\restriction_B = g\restriction_B \text{ for some } u \in G_A\} \tag{11}$$

Here, G_A is the subgroup of G of elements fixing A pointwise.
(5) Assume now that small subsets $A \subseteq B$, and $g \in G$ have the following property: given any small C containing A, and an element $h \in G$ such that $g\restriction_A = h\restriction_A$, there are elements $f \in G$ and $u \in G_A$, such that $u^{-1}fu\restriction_B = g\restriction_B$ and $f\restriction_C = h\restriction_C$.

Show that $V_{g,A,B}$ is dense. (Hint: it is enough to show that intersects any basic open set $U_{g,C}$ where $g\restriction_A = h\restriction_A$).
(6) Show that if the continuum hypothesis holds for infinite cardinals smaller than κ, then the number of subsets $V_{g,A,B}$ is κ. Conclude that in this case, the intersection of all $V_{g,A,B}$ as above is comeagre.
(7) Finally, assume that the continuum hypothesis holds, and let M be a saturated model of a countable stable theory, of cardinality $\aleph_1$. Let G be its automorphism group. Show that $V_{g,A,B}$ is dense in G whenever A is an algebraically closed countable subset, and conclude that the set of $\aleph_1$-generic automorphisms of M is comeagre in G.

End lecture 23

5. The independence theorem

5.1. In this section we continue to assume that $\mathcal{T}$ is a stable theory, and $\mathcal{T}_\sigma$ exists. We have seen that for any two types p and q of $\mathcal{T}$ over an algebraically closed set, we may canonically define a new type $p \otimes q$. In terms of subsets of a model, this can be phrased as follows: Given two sets A and B extending the algebraically closed set C, there is a copy A_1 of A over C, such that A_1 and B are independent over C, and if A_2 is another copy with the same property, then the isomorphism from A_1 to A_2 can be extended to an isomorphism over $\mathrm{acl}(B)$ from $\mathrm{acl}(A_1 \cup B)$ to $\mathrm{acl}(A_2 \cup B)$ (cf. 3.2.25).

If we are to have a notion of "free amalgamation" on types in $\mathcal{T}_\sigma$, it is reasonable to expect that it should be compatible with the one in $\mathcal{T}$. However, we have seen (3.2.14 and 4.1.10) that there are several ways to extend $p \otimes q$ to a type of $\mathcal{T}_\sigma$, and there is no good way of deciding which is the free one. Instead, we just use freeness as a property of subsets inherited from $\mathcal{T}$, and drop the uniqueness requirement:

5.1.1. *Definition.* Let E be an algebraically closed σ-structure. We say that the T_σ type $p(x, y)$ over E is a *free amalgam* of $q(x) = p\restriction_x$ and $r(y) = p\restriction_y$ if $\bar{p} = \bar{q} \otimes \bar{r}$, where for each type s, $\bar{s}$ is its restriction to $\mathcal{T}$.

Likewise, the σ-substructures A and B (of some model) are *free* over E if they are free in the sense of $\mathcal{T}$.

Thus, any two types (over an algebraically closed subset) have a free amalgam, but it is not unique. The uniqueness of the amalgamation in a stable theory implies the existence of a *higher amalgam*: Any three free amalgams $p(x, y)$, $q(y, z)$ and $r(x, z)$ can be amalgamated to a 3-type $s(x, y, z)$, provided it agrees on the "intersections" (see 5.2.5 below). Thus, this property can be viewed as a weakening of the uniqueness. The independence theorem, proved below, states that this property holds in $\mathcal{T}_\sigma$ (this independence theorem has nothing to do with the independence property discussed above). It turns out that to prove this property it is convenient to analyse higher amalgamation in the stable theory $\mathcal{T}$. The presentation below is taken from Hrushovski [18]. As an application, we will finally prove elimination of imaginaries for $ACFA$.

5.2. **Higher amalgamation.** We denote by $\mathbb{P}(n)$ the set of subsets of $[n] = \{1, \ldots, n\}$, and by $\mathbb{P}_-(n)$ the set of proper subsets.

5.2.1. *Definition.* Let $\mathcal{T}$ be a stable theory (as usual, with QE and EI). An *n-amalgamation problem* $\mathfrak{a}$ for $\mathcal{T}$ consists of the following data:

(1) An algebraically closed $\mathcal{T}$-structure $\mathfrak{a}(s)$ for all $s \in \mathbb{P}_-(n)$.
(2) An elementary map $\mathfrak{a}(f) : \mathfrak{a}(s) \to \mathfrak{a}(t)$ whenever $f : s \to t$ is an inclusion

This data should satisfy the following conditions:

(1) $\mathfrak{a}$ commutes with composition, and sends the identity on s to the identity on $\mathfrak{a}(s)$ (the definition so far says that $\mathfrak{a}$ is a functor from $\mathbb{P}_-(n)$ to the category of algebraically closed $\mathcal{T}$-sets.) It follows that we may view elements of $\mathfrak{a}(s)$ as elements of $\mathfrak{a}(t)$ whenever $s \subseteq t$.
(2) For $s_0 \subseteq s_1, s_2 \subseteq s_3$, $\mathfrak{a}(s_1)$ and $\mathfrak{a}(s_2)$ are free over $\mathfrak{a}(s_0)$ inside $\mathfrak{a}(s_3)$
(3) Any element of $\mathfrak{a}(s)$ is algebraic over $\cup_{i \in s}\mathfrak{a}(i)$.

A map from one amalgamation problem $\mathfrak{a}_1$ to another $\mathfrak{a}_2$ is a map of functors, i.e., a collection of elementary maps $f_s : \mathfrak{a}_1(s) \to \mathfrak{a}_2(s)$ such that $f_t = f_s\restriction_{\mathfrak{a}(t)}$ whenever $t \subseteq s$.

A *solution* to an amalgamation problem $\mathfrak{a}$ is an algebraically closed set $\mathfrak{a}([n])$ together with maps $\mathfrak{a}(f) : \mathfrak{a}(s) \to \mathfrak{a}([n])$ for any inclusion $f : s \to [n]$, such that the conditions above hold with $\mathbb{P}_-$ replaced by $\mathbb{P}$. A map between two solutions $\mathfrak{a}_1$ and $\mathfrak{a}_2$ of the same problem $\mathfrak{a}$ is a functorial map as above, whose restriction to $\mathfrak{a}$ is the identity.

We say that *n-existence* holds in $\mathcal{T}$ if any n-amalgamation problem has a solution, and that *n-uniqueness* holds if there is at most one solution to each such problem, up to isomorphism.

5.2.2. Thus, by definition, 2-existence and uniqueness holds in any stable theory. We will see below that this implies 3-existence as well. It is important to note that all structures are required to be algebraically closed. This is the cause of the possible failure of higher existence and uniqueness:

5.2.3. *Exercise.* Assume that for any algebraically closed set A and B, $\operatorname{acl}(A\cup B) = \operatorname{dcl}(A\cup B)$. Show that all amalgamation problems have a unique solution.

5.2.4. *Example.* The following example from Piro et al. [35] (and due to Hrushovski) shows the failure of 3-uniqueness (and 4-existence). Let A be an infinite set, and let B be the set of pairs (d, i), where $d \in A^{(2)}$ is a subset of size 2 of A, and $i \in \mathbb{Z}/2\mathbb{Z}$. Let $P \subseteq B^3$ be the ternary relation such that $P((d_1, i_1), (d_2, i_2), (d_3, i_3))$ holds if and only if $i_1 + i_2 + i_3 = 0$, and each d_i, d_j have precisely one element in common (for $i \neq j$).

We consider the theory $\mathcal{T}$ of (A, B, P, f), where we let $f : B \to A^{(2)}$ be the projection. This theory is a reduct of the theory of $(A, \mathbb{Z}/2\mathbb{Z})$, which is just a pure set with some independent finite bit, so is stable. Hence $\mathcal{T}$ itself is also stable, but it is important that $\mathbb{Z}/2\mathbb{Z}$ is not part of the structure in $\mathcal{T}$. We also note that if $a \in A$ is a singleton, then $\operatorname{acl}(a) = a$.

Let $a_1, a_2, a_3 \in A$. We consider the 3-amalgamation problem with $\mathfrak{a}(i) = a_i$ for $i \in [3]$. Then $\mathfrak{a}(\{i, j\}) = \operatorname{acl}(a_i, a_j)$ contains the two elements $b_k = (\{a_i, a_j\}, 0) \in B$ and $c_k = (\{a_i, a_j\}, 1) \in B$ (where $\{i, j, k\} = \{1, 2, 3\}$), which are conjugate over a_i, a_j. This amalgamation problem now has the following two solutions: in each case, $\mathfrak{a}(\{a_1, a_2, a_3\})$ consists of $\operatorname{acl}(\{a_1, a_2, a_3\}$, and in particular, contains the b_i and c_i. In both the maps $f_3 : \mathfrak{a}(\{1, 2\}) \to \mathfrak{a}(\{1, 2, 3\})$ and $f_2 : \mathfrak{a}(\{1, 3\}) \to \mathfrak{a}(\{1, 2, 3\})$ are the inclusion, but the map $f_1 : \mathfrak{a}(\{2, 3\}) \to \mathfrak{a}(\{1, 2, 3\})$ is the inclusion in the first solution, and exchanges b_1 and c_1 in the second one. Then both are solutions, but they are not isomorphic over the problem, since $P(c_1, b_2, b_3)$ holds, but $P(b_1, b_2, b_3)$ doesn't. Thus, 3-uniqueness fails. Using the two solutions as components in a 4-amalgamation problem, it is easy to show that 4-existence fails as well. This example was generalised to any $n > 2$ in Pastori and Spiga [30].

We note that in this example, the failure of uniqueness occurred because $c_1 \in \mathfrak{a}(2, 3)$ is definable over $b_2, b_3 \in \operatorname{dcl}(\mathfrak{a}(1, 3) \cup \mathfrak{a}(1, 2))$, but not over $\mathfrak{a}(2) \cup \mathfrak{a}(3)$. We will see below that this happens in general.

5.2.5. *Proposition.* Assume $\mathcal{T}$ admits n-existence and n-uniqueness. Then it admits $n + 1$-existence.

Proof. Let $\mathfrak{a}$ be an $n+1$ problem. We define an n-problem $\mathfrak{b}$ by $\mathfrak{b}(u) = \mathfrak{a}(u \cup \{n+1\})$. By n-existence, $\mathfrak{b}$ has a solution $(\mathfrak{b}([n]), g_u))$. We set $\mathfrak{a}([n + 1]) = \mathfrak{b}([n])$ and $g_{u\cup\{n+1\}} = f_u$. Since $\mathfrak{b}$ is a solution, it is clear that the g_v satisfy the conditions of a solution for $\mathfrak{a}$. We need to define $g_{[n]} : \mathfrak{a}([n]) \to \mathfrak{b}([n])$. For any v properly contained in $[n]$, $\mathfrak{a}(u)$ is already included in $\mathfrak{b}([n])$. Hence $\mathfrak{b}([n])$ contains a solution of the restriction of $\mathfrak{a}$ to $\mathbb{P}_-(n)$. Since $\mathfrak{a}([n])$ is also a solution, we get by uniqueness an embedding of $\mathfrak{a}([n])$ in $\mathfrak{b}([n])$. □

5.2.6. We defined amalgamation problems (and solutions) only in the context of stable theories. However, we have only used stability to have some notion of independence. Hence the definitions make sense in the context of any reasonable notion of independence. In particular, it makes sense for $\mathcal{T}_\sigma$, with independence defined as in 5.1.1. As mentioned there, in the more general case we need not have 2-uniqueness. However, in light of 5.2.5, it makes sense to consider 3-existence as a weak substitute. We now wish to characterise 3-existence for $\mathcal{T}_\sigma$ in terms of $\mathcal{T}$.

Given an n-amalgamation problem $\mathfrak{a}$ and a subset $u \subseteq [n]$, we denote by $\mathfrak{a}(<u)$ the definable closure of all $\mathfrak{a}(v)$ where v is strictly contained in u. We denote by $\mathfrak{a}(\not\supseteq$

u) the definable closure of all $\mathfrak{a}(v)$ where v does not contain u. For example, if $n = 3$ and $u = \{1, 2\}$, then $\mathfrak{a}(<u) = \mathrm{dcl}(\mathfrak{a}(1) \cup \mathfrak{a}(2))$, and $\mathfrak{a}(\ngeq u) = \mathrm{dcl}(\mathfrak{a}(1, 3) \cup \mathfrak{a}(2, 3))$. So in the example above, there was an element of $\mathfrak{a}(u) \cap \mathfrak{a}(\ngeq u)$ that was not in $\mathfrak{a}(<u)$.

5.2.7. *Theorem.* Assume that $\mathcal{T}$ is stable and $\mathcal{T}_\sigma$ exists. Let $\mathfrak{a}$ be an n-amalgamation problem for $\mathcal{T}$ that has a solution. The following are equivalent:

(1) The solution of $\mathfrak{a}$ is unique.
(2) Any expansion of $\mathfrak{a}$ to a problem in $\mathcal{T}_\sigma$ has a solution.
(3) In $\mathfrak{a}([n])$, $\mathfrak{a}(v) \cap \mathfrak{a}(\ngeq v) = \mathfrak{a}(<v)$ for all subsets v of size $n - 1$.

In particular, if $\mathcal{T}$ admits n-existence, then n-uniqueness in $\mathcal{T}$ is equivalent to n-existence in $\mathcal{T}_\sigma$.

Proof. (1) Let $\mathfrak{a}_\sigma$ be an expansion of $\mathfrak{a}$ to a problem in $\mathcal{T}_\sigma$. By assumption, $\mathfrak{a}$ has a solution $A = \mathfrak{a}([n])$, with embeddings $f_u : \mathfrak{a}(u) \to A$ for each proper subset u of $[n]$. Let σ_u be the automorphism of $\mathfrak{a}(u)$ given with $\mathfrak{a}_\sigma(u)$. Then $f_u \circ \sigma_u$ is another solution, so by uniqueness, there is an automorphism σ of A such that $f_u \circ \sigma_u = \sigma \circ f_u$. Then (A, σ) is a solution of $\mathfrak{a}_\sigma$ (note that (A, σ) is algebraically closed by 4.1.8).

(2) By symmetry we may assume $v = [n-1]$. Assume that there is an element $b \in \mathfrak{a}([n-1]) \cap \mathfrak{a}(\ngeq [n-1])$ that is not in $\mathfrak{a}(<[n-1])$. Then there is an automorphism σ fixing $\mathfrak{a}(u)$ pointwise for each proper subset u of $[n-1]$, but that does not fix b. Consider the problem $\mathfrak{a}_\sigma$ where $\sigma_{[n-1]} = \sigma$, and σ_u is the identity for any other u. This is an n-problem for $\mathcal{T}_\sigma$, so by assumption has a solution. But such a solution is the identity on $\mathfrak{a}(\ngeq [n-1])$, hence on b, contradicting the choice of $\sigma_{[n-1]}$.

(3) Let (B, b_u) be a solution of $\mathfrak{a}$. The assumption means that for any u of size $n-1$, any automorphism of $\mathfrak{a}(u)$ fixing $\mathfrak{a}(<u)$, extends to an automorphism of B fixing the rest of the problem. Hence (by composing them), any automorphism of the problem extends to an automorphism of the solution. If (C, c_u) is another solution, we may (by 2-uniqueness) fix an isomorphism $T : B \to C$ over all the $\mathfrak{a}(i)$. Then $c_u^{-1} \circ T \circ b_u$ is an automorphism of the problem, hence it extends to an automorphism of the solutions.

□

5.2.8. *Corollary.* Assume that $\mathcal{T}$ has n-existence, and let $\mathfrak{a}$ be an n-amalgamation problem for $\mathcal{T}_\sigma$, such that $\mathfrak{a}(0)$ is a model of $\mathcal{T}$. Then $\mathfrak{a}$ has a solution (and the restriction of $\mathfrak{a}$ to $\mathcal{T}$ has a unique solution).

Proof. It is enough to prove that the third condition of 5.2.7 holds. Let b_n be an element of $\mathfrak{a}([n-1]) \cap \mathfrak{a}(\ngeq [n-1])$. Thus, there is a definable function f and elements $b_i \in \mathfrak{a}([n] - i)$ for $i < n$, such that $b_n = f(b_1, \dots, b_{n-1})$. In turn, there are formulas $\phi_i(y, x_1, \dots, \check{x}_i, \dots, x_n)$ (x_i omitted) for $i \leq n$, and $c_j \in \mathfrak{a}(j)$ such that $\phi_i(y, c_1, \dots, c_n)$ has a finite (and minimal) number of solutions, b_i being one of them.

We may assume that $\mathfrak{a}([n-1])$ is independent of $\mathfrak{a}(n)$ over the model $M = \mathfrak{a}(0)$. In particular, b_n and the c_i ($i < n$) are independent from c_n. Hence (see 5.2.9) there is an element c'_n in M, and elements b'_i ($i < n$) such that $\phi_i(b'_i, c_1, \dots, \check{c}_i, \dots, c_{n-1}, c'_n)$ and $f(b'_1, \dots, b'_{n-1}) = b_n$ both hold (and shows b'_i to be algebraic over the c_i, c'_n). But this means that $b'_i \in \mathfrak{a}([n-1] - i)$, so $b_n \in \mathfrak{a}(<[n-1])$. □

5.2.9. *Exercise.* Let $\mathcal{T}$ be a stable theory, and let a and b be independent over a model M. Show that if $\phi(a, b)$ holds, where ϕ is a formula over M, then there is $b' \in M$, such that $\phi(a, b')$ holds as well.

5.2.10. *Corollary (the Independence Theorem).* Let $p(x, y)$, $q(x, z)$ and $r(y, z)$ be types in $\mathcal{T}_\sigma$ over a model of M, such that $p\restriction_x = q\restriction_x$, $p\restriction_y = r\restriction_y$ and $q\restriction_z = r\restriction_z$, and such that each is an independent amalgam of its restrictions. Then there is a type $s(x, y, z)$ whose restrictions to any two variables is the given one, and such that each variable is independent from the other two.

Proof. Exercise □

5.2.11. *Definition.* A theory $\mathcal{T}$ is called *simple* if there is a ternary relation $A \mathop{\smile\!\!\!\!\!\!|}_C B$ on definably closed sets, satisfying the following conditions:

1. If σ is an automorphism, $A \mathop{\smile\!\!\!\!\!\!|}_C B$ if and only if $\sigma(A) \mathop{\smile\!\!\!\!\!\!|}_{\sigma(C)} \sigma(B)$.
2. For any A, B and C, there is an automorphism σ over C, such that $A \mathop{\smile\!\!\!\!\!\!|}_C \sigma(B)$.
3. $A \mathop{\smile\!\!\!\!\!\!|}_C B$ if and only if $B \mathop{\smile\!\!\!\!\!\!|}_C A$.
4. For $A \subseteq B \subseteq C$, $D \mathop{\smile\!\!\!\!\!\!|}_A C$ if and only if $D \mathop{\smile\!\!\!\!\!\!|}_B C$ and $D \mathop{\smile\!\!\!\!\!\!|}_A B$.
5. 3-existence holds over a model, where in the definition of amalgamation, independence is taken in the given sense.
6. For any A and B, there is a subset C of B, of cardinality at most $\#\mathcal{T}$, such that $A \mathop{\smile\!\!\!\!\!\!|}_C B$.

Given such a ternary relation, we say that a type p over B *does not fork* over $C \subseteq B$, if for some realisation A of p, $A \mathop{\smile\!\!\!\!\!\!|}_C B$. By 1, this will holds for any other realisation as well.

5.2.12. *Remarks.*

1. It thus follows from 3.2.23, 3.2.25 and 5.2.5 that any stable theory is simple, and p does not fork over A precisely if it is definable over $\operatorname{acl}(A)$. It can be shown a simple theory is stable precisely if 3-existence is replaced by 2-uniqueness (over arbitrary algebraically closed subsets).

   On the other hand, by the same facts and 5.2.10 in place of 5.2.5, $\mathcal{T}_\sigma$ is simple (but in general unstable). In fact, we showed something a bit stronger, since 3-existence holds even over models of $\mathcal{T}$.
2. The original definition of simple theories was given by Shelah in terms of an explicit construction of an independence relation. This relation can be defined in an arbitrary theory, but a theory was called simple if that particular relation has the properties above. Then Kim and Pillay [22] proved that if a theory has a freeness relation as above it must be the original non-forking relation. In particular, a simple theory has a unique such relation.

5.2.13. A definable set $\mathbf{X}$ is said to be *strongly minimal* if it is infinite, and any definable subset of $\mathbf{X}$, even with parameters, is finite or co-finite. The theory $\mathcal{T}$ is strongly minimal if there is a strongly minimal set $\mathbf{X}$ of $\mathcal{T}$, such that any other definable set can be embedded into a power of $\mathbf{X}$.

Examples include the theory of equality, the theory of algebraically closed fields, and the theory of a vector space (over a fixed field).

5.2.14. *Exercise.* Show that $\mathbf{X}$ has a unique non-algebraic type. In particular, a strongly minimal theory is ω-stable (hence stable).

5.2.15. *Exercise.* If $\mathcal{T}$ is strongly minimal, show that any infinite algebraically closed $\mathcal{T}$ structure is a model.

5.2.16. *Corollary.* If $\mathcal{T}$ is a strongly minimal theory where $\operatorname{acl}(0)$ is infinite, then $\mathcal{T}_\sigma$ admits n-existence for all n.

Proof. Exercise □

5.3. **Indiscernible sequences.** We make a small detour to introduce and study indiscernible sequences. These will be used to prove independence in 5.4.3. Though indiscernible sequences can be defined (and used) in arbitrary theories, we restrict to the stable context immediately after the definition.

5.3.1. *Definition.* Let A be a sub-structure (in an arbitrary theory). A sequence (a_i) of elements (in some model containing A) is *indiscernible* over A if for any n, all sub-sequences of a_i of length n have the same type over A.

We will mostly consider indiscernibles over 0, adding A to the language if necessary.

5.3.2. *Exercise.* Let p be a type in a stable theory. Define a sequence recursively by taking a_α to be a realisation of $p|_{\{a_i : i<\alpha\}}$ (of course, the sequence is not well defined, but we can always produce such a sequence). Show that (a_i) is an indiscernible sequence. Such a sequence is called a *Morley sequence*.

On the other hand, note that any constant sequence is indiscernible, so not any indiscernible sequence is a Morley sequence.

5.3.3. *Exercise.* Let $I = (a_i)$ be an indiscernible sequence in a stable theory. Show that for every formula $\phi(x, y)$ and any element b, $\phi(a_i, b)$ either holds for almost all a_i, or for finitely many (use compactness to show that otherwise ϕ has the order property). Conclude that given a set B of parameters, the set

$$Av(I/B) = \{\phi(x, b) : b \in B, \phi(a_i, b) \text{ holds for almost all } i\} \tag{12}$$

is a complete type over B. This type is called the *average type* of I (over B).

5.3.4. *Proposition.* Let (a_i) be an indiscernible sequence in a stable theory, such that for any two disjoint finite subsets B and C of the $\{a_i\}$, $\operatorname{acl}(B) \cap \operatorname{acl}(C) \subseteq \operatorname{acl}(0)$. Then any two such B and C are independent over A (in particular, a_1 is independent from a_2).

Proof. Let $p = Av(I/I)$. Then p is definable over $\operatorname{acl}(I)$. Given a formula $\phi(x, y)$, let $\mathfrak{d}_p\phi = \psi(y, a_{i_1}, \dots, a_{i_n})$ be the p-definition of ϕ. Then we have that $\psi(a_j, a_{i_1}, \dots, a_{i_n})$ holds iff $\phi(x, a_j) \in p$, iff $\phi(a_k, a_j)$ holds for infinitely many k, iff $\phi(a_k, a_j)$ holds for all $k \neq j$, since the sequence is indiscernible. Hence $\psi(a_j, a_{i_1}, \dots, a_{i_n})$ has the same truth value for all j. Again by indiscernability, we conclude that $\psi(a_j, a_{j_1}, \dots, a_{j_n})$ has the same truth value for any other sequence $a_{j_1}, \dots, a_{j_n})$ (of the same order type). In other words, $\psi(y, a_{j_1}, \dots, a_{j_n})$ defines the same subset of I for all tuples $(a_{j_1}, \dots, a_{j_n})$. A similar argument shows that the all define the same subset of $\operatorname{acl}(I)$. Hence they are all equivalent.

It follows that the code of $\mathfrak{d}_p\phi$ is definable over any finite subsequence of length n. It now follows from the assumption that $\mathfrak{d}_p\phi$ is definable without parameters.

This is true for all formulas, so the canonical base of p is in $\mathrm{acl}(0)$, and p is the non-forking extension over I of its restriction to $\mathrm{acl}(0)$. □

We next consider the question of producing an indiscernible sequence. We will need extract an indiscernible sub-sequence from a sequence that we can make arbitrarily large. For that, the following generalisation of Ramsey's theorem is useful.

5.3.5. *The Erdös–Rado theorem.* For any cardinal λ, there is another cardinal κ such that for any subset X of $\kappa^{(2)}$, there is a subset Y of κ, $\#Y > \lambda$, such that $Y^{(2)}$ is contained in either X or its complement. Explicitly, κ can be taken to be the successor of 2^{λ}. (And more generally, this is true also with λ colours rather than 2, and with n-element subsets in place of 2 if we take $\kappa = \beth_{n-1}(\lambda)^{+}$.)

See Chang and Keisler [8, p. 7.2] for details and a proof.

5.3.6. *Exercise.* Show that for any cardinal λ there is a cardinal κ, such that for any theory $\mathcal{T}$ of cardinality at most λ, any sequence of elements of length at least κ contains an indiscernible sequence over A.

5.4. **Imaginaries in $\mathcal{T}_\sigma$.** Our main application of 3-existence will be to prove elimination of imaginaries for $\mathcal{T}_\sigma$. In this sub-section, we denote by acl_σ the algebraic closure of $\mathcal{T}_\sigma$ in its original sorts, by acl^{eq} the full algebraic closure, and by acl the algebraic closure in $\mathcal{T}$.

We need the following group theoretic result.

5.4.1. *Neumann's Lemma.* Let G be a group acting on a set X, such that the orbit of each element is infinite. Then for any two finite subsets A, B of X, there is an element $g \in G$, such that gA and B are disjoint.

5.4.2. *Exercise.* Let $\mathcal{T}$ be a theory, let M be a saturated model, and let $A, B \subseteq M$ be small subsets, disjoint from $\mathrm{dcl}(0) = \mathrm{acl}(0)$. Show that there is an automorphism σ of M such that $\sigma(A)$ and B are disjoint.

5.4.3. *Lemma.* Assume we have an equivalence relation in $\mathcal{T}_\sigma$ on a definable set $\mathbf{X}$. Let e be the equivalence class of an element $a \in \mathbf{X}$, let $E = \mathrm{acl}_\sigma(e)$, and let $p = \mathrm{tp}_\sigma(a/E)$. Then there is a realisation b of p, such that b is independent from a over E and is equivalent to it.

If there is a realisation of p that is not equivalent to a, then there is such a realisation which is independent from a over E.

Proof. We assume, to simplify notation, that $E = \mathrm{dcl}(0)$. Let $A = \mathrm{acl}_\sigma(a) - E$. Then A is disjoint from $\mathrm{acl}^{eq}(e)$, so by 5.4.2, there is an automorphism τ, fixing e, such that $\tau(A)$ and A are disjoint. In particular, $a_1 = a$ and $a_2 = \tau(a)$ have the same equivalence class e, and satisfy $\mathrm{acl}(a_1) \cap \mathrm{acl}(a_2) \subseteq E$.

We may construct, by induction on n, a sequence (a_n) extending a_1, a_2, such that $\mathrm{tp}_\sigma(a_n, a_{n+1}) = \mathrm{tp}_\sigma(a_1, a_2)$, and a_{n+1} is independent of $a_1, \dots, a_{n-1}$ over a_n. Indeed, assume a_n was constructed. Let $q = \mathrm{tp}_\sigma(a_2/a_1)$. Since a_1 and a_2 have the same type, the assumption implies that all the a_i have the same type as well. Hence, there is an automorphism τ that takes a_1 to a_n. Let a_{n+1} be a realisation of a non-forking extension of $\tau(q)$ to $a_1, \dots, a_{n-1}$.

We next claim that if $n < m$, then $\mathrm{acl}(a_1, \dots, a_n) \cap \mathrm{acl}(a_{n+1}, \dots, a_m) = E$. We prove this by induction on m. For $m = 2$ this was shown above. Assume

$b \in \mathrm{acl}(a_1, \dots, a_n) \cap \mathrm{acl}(a_{n+1}, \dots, a_m)$. Since a_m is independent from the others over a_{m-1}, we have that $a_{n+1}, \dots, a_m$ is independent from $a_1, \dots, a_n$ over $C = \{a_{n+1}, \dots, a_{m-1}\}$ (where this set includes at least a_{m-1}). By 3.2.26, $b \in \mathrm{acl}(C)$. Now, if $n < m-1$, we may use the induction hypothesis. If $n = m-1$, we get that $b \in \mathrm{acl}(a_n) \cap \mathrm{acl}(a_{n+1})$, in which case the claim again follows from the above (since (a_n, a_{n+1}) has the same type as (a_1, a_2)).

We thus found a sequence of elements (a_i), all belonging to the same equivalence class, all satisfying p, and such that $\mathrm{acl}(a_1, \dots, a_n) \cap \mathrm{acl}(a_{n+1}, \dots, a_m) = E$ for $n < m$. By compactness, we may find an arbitrary long sequence with these properties. By 5.3.6, we may further assume that the sequence is indiscernible. By 5.3.4, any two elements of this sequence are equivalent independent realisations of p.

If p is satisfied by a non-equivalent element, then in particular the class e itself is not algebraic. We may thus find a sequence as above, all of whose elements are pairwise inequivalent. The rest is similar. □

5.4.4. *Theorem.* Assume that $\mathcal{T}_\sigma$ admits 3-existence. Then $\mathcal{T}_\sigma$ weakly eliminates imaginaries. In particular, if $\mathcal{T}$ codes finite sets, then $\mathcal{T}_\sigma$ eliminates imaginaries.

Proof. We will the third criterion of 3.1.18. So assume that we are given an equivalence relation in $\mathcal{T}_\sigma$, and a is an element whose equivalence class is e (in $\mathcal{T}_\sigma$). We need to show that any element realising the type p of a over $E = \mathrm{acl}_\sigma(e)$ is equivalent to a.

Assume that this is not the case. Then, by 5.4.3, there is a counter-example b which is independent from a. Also by 5.4.3, there is an element c realising p, free from a, and equivalent to a. Let $q = \mathrm{tp}(a, b)$, $r = \mathrm{tp}(a, c)$. Then q, q and r form a 3-amalgamation problem, which by assumption has a solution $s(x, y, z)$. By then s says that x and y are equivalent, as are y and z, but not x and z. This is a contradiction. □

5.4.5. *Remark.* Tracing through the proofs of 5.4.4 and 5.4.3, we see that the automorphism was not used. In other words, the following statement was proved: Suppose that a theory $\mathcal{T}$ is an expansion of $\mathcal{T}_0$ with no additional sorts, and suppose that $\mathcal{T}$ satisfies 3-existence with respect to the independence notion that comes from $\mathcal{T}_0$. Then $\mathcal{T}$ has weak EI.

From now on, we assume that $\mathcal{T}_\sigma$ eliminates imaginaries (as we saw, this holds in $ACFA$)

5.4.6. *Definition.* Let $\mathbf{X}$ be a definable set in a theory $\mathcal{T}$. $\mathbf{X}$ is *stably embedded* if any subset of $\mathbf{X}^n$ definable with parameters in $\mathcal{T}$ is definable with parameters from $\mathbf{X}$.

In other words, the collection of definable subsets of $\mathbf{X}$ with parameters is the same in $\mathcal{T}$ and in the induced structure on $\mathbf{X}$.

5.4.7. *Exercise.*

(1) Show that if $\mathcal{T}$ is stable, then any definable set is stably embedded (if $\mathbf{Y}$ is an a-definable subset of $\mathbf{X}^n$, where a is a canonical parameter, consider $\mathrm{tp}(a/\mathbf{X}(M))$)
(2) Show that in the theory of 3.2.13, the set $\mathbf{X}$ is not stably embedded.

5.4.8. *Corollary.* The fixed set $\mathbf{C}$ in $\mathcal{T}_\sigma$ is stably embedded. If there is a small subset $\mathbf{C}_0$ of $\mathbf{C}$, such that $\operatorname{acl}(\mathbf{C}) = \operatorname{dcl}(\operatorname{acl}(\mathbf{C}_0) \cup \mathbf{C})$, then any definable subset of $\mathbf{C}^n$ is given (within $\mathbf{C}$) by a formula not involving σ (but possibly with parameters from $\mathbf{C}_0$). We note that this condition is satisfied in $ACFA$.

Proof. Assume $\mathbf{X} \subseteq \mathbf{C}^n$ is definable in $\mathcal{T}_\sigma$ over a. Since $\mathcal{T}_\sigma$ eliminates imaginaries, we may assume that a is a canonical parameter. Hence, it is fixed by any automorphism that fixes Y pointwise. Since σ is one such automorphism, $a \in \mathbf{C}$.

For the second statement, we may assume that $\mathbf{C}_0$ is an elementary substructure of $\mathbf{C}$ (in the induced theory). We show that any automorphism τ of the induced structure on $\mathbf{C}$ over $\mathbf{C}_0$ can be extended to an automorphism of the model. Since $\mathbf{C}_0$ is an elementary substructure, any subset of $\mathbf{C}^n$ definable over $\operatorname{acl}(\mathbf{C}_0)$ is fixed by τ. Hence, τ can be extended to an automorphism of $\operatorname{dcl}(\operatorname{acl}(\mathbf{C}_0) \cup \mathbf{C}) = \operatorname{acl}(\mathbf{C})$ whose restriction to $\operatorname{acl}(\mathbf{C}_0)$ is the identity. Since σ preserves $\operatorname{acl}(\mathbf{C}_0)$ (as a set), τ commutes with σ. In other words, it is an automorphism of the σ-structure $(\operatorname{acl}(\mathbf{C}), \sigma)$. By 4.1.7, τ is an automorphism of $\mathbf{C}$ (in the sense of the full theory).

For the case of fields, we know that the fixed field has a unique extension of each degree (2.3.5), so the condition holds with any elementary subfield (the same argument works whenever the "absolute Galois group" $\operatorname{Aut}(\mathbf{C}^a/\mathbf{C})$ depends only on the theory). □

Lecture 34, Nov 18, 2009

5.4.9. *Exercise.* Let ϕ be a definable automorphism in $\mathcal{T}$. Show that (M, σ) is a model of $\mathcal{T}_\sigma$ if and only if $(M, \sigma \circ \phi)$ is. Show that the fixed set of $\sigma \circ \phi$ is stably embedded in (M, σ). In the context of $ACFA$, it easy to see that the only definable automorphisms in ACF are ϕ^n, where ϕ is the Frobenius if the characteristic is positive, the identity otherwise.

6. The dichotomy theorem

In this section, we will mostly specialise again to the context of $ACFA$. The dichotomy theorem is one of the main results of the paper Chatzidakis and Hrushovski [10] (in characteristic 0), as well as of Chatzidakis et al. [12], which provides a more conceptual proof, valid in all characteristics. Our aim in this section is to explain the statement of the theorem.

The dichotomy theorem gives an analysis of the structure of "finite rank" types in $ACFA$. An example of a type of rank 1 is the type of any non-algebraic element of the fixed field. We have seen that the fixed field is unstable, and so are the twisted fixed fields of 5.4.9. The dichotomy implies that any unstable type of rank 1 is strongly related to one of these fields. If the type is not related to one of the fields in this way, then it is *modular*: any two algebraically closed sets are independent over their intersection. This property has important applications: for instance, a definable subset of a modular group is a boolean combination of cosets of subgroups (compare this to the statement of the Manin–Mumford conjecture, 1.2.1).

For higher (finite) rank types, it is shown that they can be analysed in terms of the rank 1 types. Again it follows that if a type is unrelated to the (twisted) fixed fields, then it is modular. In particular, the only proper subfields are the twisted fixed fields.

We mention that classically, these notion come from the ω-stable context. A similar (and simpler) analysis is available in the analogous context of differentially closed fields (which are ω-stable; cf. Pillay [31]). In fact, at least the rank 1 case is true in almost complete generality (assuming the existence of a "Zariski geometry", Hrushovski and Zilber [21]. The differential fields case follows from this, and the difference field case from a variant of this. See also Bouscaren [6].)

To state all of that precisely, we need to define the ranks, explain how to relate types of rank 1, and how can one analyse types of higher rank in terms of rank 1 types. We start with the ranks.

6.1. **Ranks.** We define three notions of rank. The S_1-rank is defined for an arbitrary theory, and the SU-rank is defined for arbitrary simple theories. They coincide when either of them is finite. The third rank is specific to $\mathcal{T}_\sigma$ where $\mathcal{T}$ is ω-stable. It is less geometric, but has the advantage that it can be easily computed in examples.

6.1.1. *Definition.* Let $\mathcal{T}$ be an arbitrary theory. The S_1*-rank* on definable sets of $\mathcal{T}$ is defined inductively as follows:

(1) $S_1(\mathbf{X}) \geq 0$ if (and only if) $\mathbf{X}$ is non-empty.
(2) $S_1(\mathbf{X}) \geq n+1$ if there is an indiscernible sequence $I = (a_i)$, and a family $\mathbf{Z}_{a_i}$ such that:
    (a) For $i \neq j$, $S_1(\mathbf{Z}_{a_i} \cap \mathbf{Z}_{a_j}) < n$ (i.e., not $\geq n$)
    (b) For each i, $S_1(\mathbf{Z}_{a_i} \cap \mathbf{X}) \geq n$.

    We note that rank conditions above refer to sets definable with parameters, so are computed in $\mathcal{T}_A$ for some A containing the parameters (see 6.1.4).

We say that $S_1(\mathbf{X}) = n$ if $n+1 > S_1(\mathbf{X}) \geq n$. If there is no such n, then we say $S_1(\mathbf{X}) = \infty$. For a partial type p we set $S_1(p) = \min\{S_1(\mathbf{X}) : p \subseteq \mathbf{X}\}$.

6.1.2. *Exercise.* Show that $S_1(\mathbf{X}) > 0$ if and only if $\mathbf{X}$ is infinite, and that if $\mathbf{X} \subseteq \mathbf{Y}$ then $S_1(\mathbf{X}) \leq S_1(\mathbf{Y})$.

6.1.3. *Definition.* Let $\mathcal{T}$ be a simple theory. The rank SU (simple U-rank) on types is defined as follows: $SU(p) \geq \alpha$ (for an ordinal α) if for any $\beta < \alpha$, p has a forking extension q, such that $SU(q) \geq \beta$.

Again, $SU(p) = \alpha$ if $SU(p) \geq \alpha$, but not $SU(p) \geq \alpha + 1$, and $SU(p) = \infty$ if $SU(p) = \alpha$ fails for all α.

If $\mathbf{X}$ is a definable set, we define $SU(\mathbf{X}) = \sup\{SU(p) : p \in \mathbf{X}\}$.

The theory $\mathcal{T}$ is called *supersimple* (or *superstable* if $\mathcal{T}$ was stable) if any type has some SU rank.

For any of the ranks R, we set $R(a/A) = R(\mathrm{tp}(a/A))$. End lecture 34

Lecture 35, Nov 20, 2009

6.1.4. *Exercise.* Let $\mathbf{X}$ be a definable set, A a set of parameters. Show that $SU(\mathbf{X})$ is the same when computed in $\mathcal{T}$ or in $\mathcal{T}_A$. Show that the same is true for S_1.

6.1.5. *Proposition.* Let $\mathcal{T}$ be a simple theory, and let $p \subseteq q$ be types. If q is a non-forking extension, then $SU(p) = SU(q)$. Conversely, if $SU(p) = SU(q) < \infty$, then q is non-forking.

Proof. It is clear that $SU(q) \leq SU(p)$. Assume q is a non-forking extension, let $SU(p) \geq \alpha$, and for $\beta < \alpha$, let r be an extension of p with $SU(r) \geq \beta$. By 2-existence, we may extend q to a non-forking extension of r. By induction, their SU ranks are equal.

The converse is obvious. □

6.1.6. *Exercise.* Show that a simple theory is supersimple if and only if for each type p over A there is a finite subset A_0 of A, such that p does not fork over A_0 (i.e., p is a non-forking extension of its restriction to A_0).

6.1.7. *Proposition.* Assume that $SU(b/E)$ and $SU(a/Eb)$ are both finite. Then $SU(ab/E) = SU(b/E) + SU(a/Eb)$ (there is also a statement in general, which we will not need).

Proof. We show by induction that if $SU(b/E) > n$ then $SU(a/Eb)+n < SU(ab/E)$. Assume $SU(b/E) > n+1$, and let $F \supset E$ be such that $SU(b/F) > n$. By induction, $SU(a/Fb) + n < SU(ab/F)$. We may take F to be independent from a over E, b, so by 6.1.5, $SU(a/Fb) = SU(a/Eb)$. On the other hand, $SU(ab/F)$ is definitely a forking extension of $SU(ab/E)$, so $SU(ab/E) \geq SU(ab/F)+1 > n+1+SU(a/Eb)$.

Conversely, assume $SU(ab/E) > n + 1$. Let F be such that $SU(ab/F) > n$. Since ab is not independent from F over E, transitivity implies that either b is not independent from F over E, or a is not independent from Fb over Eb. In the first case, $SU(b/F) < SU(b/E)$, and in the second, $SU(a/Fb) < SU(a/Eb)$. Either way $n < SU(b/F) + SU(a/Fb) < SU(b/E) + SU(a/Eb)$ □

End lecture 35

Lecture 36, Nov. 23, 2009

6.1.8. *Corollary.* For any type $p(x, y)$, $SU(p) \leq SU(p\restriction_x) + SU(p\restriction_y)$ (assuming everything is finite). Equality holds if and only if p is a free amalgam of its restrictions.

6.1.9. *Exercise.* Let $\mathbf{X}$ be a strongly minimal set. Show that $SU(\mathbf{X}^n) = n$. Conclude that for ACF, $SU(\mathbf{Y})$ is the (Zariski) dimension of the Zariski closure of $\mathbf{Y}$ (In fact, it is possible to prove, using a similar method, that any ω-stable theory is superstable).

6.1.10. *Proposition.* Let p be a type (over 0) in a simple theory $\mathcal{T}$, A a set of parameters. Then the set of non-forking extensions of p to A is a closed subset of the corresponding space of types.

Proof. The general statement follows easily from the explicit definition of forking (see Wagner [39]). In the case of $\mathcal{T}_\sigma$ (where $\mathcal{T}$ is stable), the non-forking extensions are precisely the expansions of the unique non-forking extension of a type $p(x_0, \dots)$ in $\mathcal{T}$, together with the formulas $\sigma^i(x_0) = x_i$. So the closed set is defined by the type and these formulas. □

6.1.11. *Exercise.* Let q be a forking extension of p to A. Show that there is an A-definable set $\mathbf{X}$ containing q, such that no point of $\mathbf{X}$ is independent from A (thinking about ACF, p is the generic point of a variety, and q is the generic point of a proper sub-variety, whose Zariski closure over the original field is the bigger variety. Then $\mathbf{X}$ will be a dense subset of this proper sub-variety.)

6.1.12. *Proposition.* Let $\mathbf{X}$ be a definable set such that $SU(\mathbf{X})$ is finite. Then $SU(\mathbf{X}) = S_1(\mathbf{X})$.

Proof. We show by induction on n that $SU(\mathbf{X}) > n$ if and only if $S_1(\mathbf{X}) > n$. For $n = -1$, this follows from 6.1.1.

Assume that $S_1(\mathbf{X}) > n$. Then there is an indiscernible sequence $\mathbf{Z}_i$ of definable subsets of $\mathbf{X}$, with $S_1(\mathbf{Z}_i) > n - 1$ and $S_1(\mathbf{Z}_i \cap \mathbf{Z}_j) < n$. By enlarging the base, we may assume that the (codes for the) $\mathbf{Z}_i$ are independent. By enlarging it even more, we may also assume that the base is a model (6.1.4).

By induction, we get a type $p_i \in \mathbf{Z}_i$ with $SU(p_i) > n-1$, but for any $p \in \mathbf{Z}_i \cap \mathbf{Z}_j$, $SU(p) < n$. We claim that p_1 is a forking extension of its restriction p to the base. Otherwise, p_1, p_2 and the type of $\mathbf{Z}_1$ and $\mathbf{Z}_2$ give define a 3-amalgamation problem, which by simplicity has a solution. This solution gives a type q that belongs to $\mathbf{Z}_1 \cap \mathbf{Z}_2$, but which is a non-forking extension of p_1. This is a contradiction.

Now assume that $SU(\mathbf{X}) > n$. Then $\mathbf{X}$ contain a type p with $SU(p) > n$. Let q be a forking extension with $SU(q) > n-1$. Let $\mathbf{Y}$ be as in 6.1.11. Then $SU(\mathbf{Y}) > n-1$.

Let $\mathbf{Y}_i$ be an independent sequence of conjugates of $\mathbf{Y}$, with canonical parameters b_i. It might not be the case that $SU(\mathbf{Y}_1 \cap \mathbf{Y}_2) < n$. However, we claim that $SU(\mathbf{Y}_1 \cap \mathbf{Y}_2) < SU(\mathbf{Y}_1)$. To show that, it is enough to show that for any $a \in \mathbf{Y}_1$, $\mathrm{tp}(a/b_1, b_2)$ forks over b_1. Otherwise, we have $b_2 \mathop{\smile\!\!\!\!\!\!|}\ b_1$ and $b_2 \mathop{\smile\!\!\!\!\!\!|}_{b_1} a$, so by transitivity, we get that $b_2 \mathop{\smile\!\!\!\!\!\!|}\ a$, contradicting the choice of $\mathbf{Y}_2$. Similarly, $SU(\bigcap_{i=1}^{k} \mathbf{Y}_i) < SU(\bigcap_{i=1}^{k-1} \mathbf{Y}_i)$ for any $k > 1$. Hence for some k, $SU(\bigcap_{i=1}^{k} \mathbf{Y}_i) < n$. We take k to be the least with this property, and define for $j > 0$, $\mathbf{Z}_j = \bigcap_{i=j}^{k+j-2} \mathbf{Y}_i$. Then $SU(\mathbf{Z}_j) > n$, but for $i \neq j$, $\mathbf{Z}_i \cap \mathbf{Z}_j$ is contained in the intersection of k of the $\mathbf{Y}_i$, hence has SU rank less than n. By induction, we may replace SU by S_1 and get the result. □

6.1.13. *Corollary.*

(1) If $S_1(\mathbf{X}) = n$ then $S_1(a/E) = n$ for some $a \in \mathbf{X}$.
(2) For any type p, $S_1(p) \geq SU(p)$ (See Buechler [7, example 6.1.2] for an example when the inequality can be strict).
(3) If there is an indiscernible sequence of definable subsets $\mathbf{Z}_i$ of $\mathbf{X}$, such that $SU(\mathbf{Z}_i \cap \mathbf{Z}_j) < n$ for $i \neq j$, and $SU(\mathbf{Z}_i) \geq n$ for all i, then $SU(\mathbf{X} > n$. (In fact, this is true even when n is an ordinal.)

Proof. Exercise □

End lecture 36
Lecture 37, Nov 30, 2009

6.1.14. *Definition.* Let M be a model of $\mathcal{T}_\sigma$. The *σ-degree* $\deg_\sigma(b/A)$ of an element b over a substructure A is defined to be the SU-rank, in $\mathcal{T}$, of the $\mathcal{T}$-type of $\mathrm{acl}_\sigma(A, b)$ over A. Obviously, this only depends on the type $p = \mathrm{tp}_\sigma(b/A)$, so we define $\deg_\sigma(p) = \deg_\sigma(b/A)$.

As before, for a definable set $\mathbf{X}$ we define $\deg_\sigma(\mathbf{X}) = \sup\{\deg_\sigma(p) : p \in \mathbf{X}\}$.

6.1.15. *Remark.* We recall that in $\mathcal{T}_\sigma$, forking is defined in terms of the underlying stable theory $\mathcal{T}$. Thus, if $p = \mathrm{tp}_\sigma(b/A)$ (where A is a σ-structure), then $SU_{\mathcal{T}_\sigma}(p) \leq \deg_\sigma(p)$. In particular, if $\mathcal{T}$ is superstable, then $\mathcal{T}_\sigma$ is supersimple. There are examples of difference equations of σ-degree bigger than 1, but SU-rank 1 (Chatzidakis and Hrushovski [10, example 6.6]).

6.1.16. *Exercise.* Show that $\deg_\sigma(a/A)$ is the transcendence degree of the difference field generated by a and A over A.

6.1.17. *Exercise.* Let p be a type in $\mathcal{T}_\sigma$ such that $\deg_\sigma(p) = d$ is finite. Recall that if (M, σ) is a model of $\mathcal{T}_\sigma$, then so is (M, σ^n) for $n > 0$ (4.2.11). Hence p can be interpreted in (M, σ^n). On the other hand, (M, σ) is a reduct of (M, σ), so this interpretation determines a (partial) type definable set p_n in (M, σ) (example: if p was given by $\sigma(x) = f(x)$, then p_n is given by $\sigma^n(x) = p(x)$). Show that $\deg_\sigma(p_n) = nd$.

6.1.18. *Exercise.* Let $\mathbf{U}$ be a variety of dimension d, $f : \mathbf{U} \to \mathbf{U}^\sigma$ a surjective rational function. Show that $\deg_\sigma(\sigma^n(x) = f(x))$ is nd ($n > 0$). In particular, for the fixed fields $\mathbf{C}_n$ given by $\sigma^n(x) = x$, we have $SU(\mathbf{C}_n) = \deg_\sigma(\mathbf{C}_n) = n$.

6.1.19. *Remark.* Let $\mathbf{Z} \subseteq \mathbf{X} \times \mathbf{Y}$ be a family. The subset of $\mathbf{Y}$ consisting of those points b such that $SU(\mathbf{Z}_b) = n$ (for a given integer n) is not, in general, definable (see Chatzidakis and Hrushovski [10, example 6.3]. At least in $ACFA$, it is a union of definable sets, see lemma 2.15 there). In this sense, SU-rank is not definable, in general. However, if SU-rank is definable in $\mathcal{T}$, then $\deg_\sigma$ is definable in $\mathcal{T}_\sigma$. For $\mathcal{T}$ strongly minimal, it is easy to show that the rank is indeed definable, so for such $\mathcal{T}$, $\deg_\sigma$ is definable in $\mathcal{T}_\sigma$ (in particular, it is definable in $ACFA$). This is another advantage of the σ-degree.

6.2. **Sets of rank 1.** We now wish to state the main dichotomy (or trichotomy) theorem, describing sets of rank 1 in $ACFA$. This description is given up to non-orthogonality, which we define next.

If p is a partial type, q a partial type extending p, we say that q is a non-forking extension of p if some complete type extending q is a non-forking extension of its restriction to the parameters of p.

6.2.1. *Definition.* Two partial types p_0 and q_0 in a simple theory are *orthogonal* if for any non-forking extensions p and q to the same set, any realisation of p is independent from any realisation of q (in other words, any common extension of p and q is a free amalgam). We write $p_0 \perp q_0$ for this condition.

6.2.2. *Exercise.* Show that p is orthogonal to both q and r if and only if it is orthogonal to any free amalgam of q and r.

6.2.3. *Exercise.* Show that if p and q are orthogonal types of ranks n and m, and realisations a and b satisfy a relation $(a, b) \in \mathbf{Z}$, then $SU(\mathbf{Z}) \geq n + m$, and the image of the projections of $\mathbf{Z}$ on p and q has ranks n and m, respectively. Thus, if p and q are orthogonal, there are almost no relations between their realisations.

6.2.4. *Exercise.* Show that in a stable theory, $\not\perp$ is an equivalence relation on types of SU-rank 1. Show that in $\mathcal{T}_\sigma$ (where $\mathcal{T}$ is stable), if p, q and r are rank 1 types over A, B and C, respectively, and $A \mathop{\smile\hskip-0.9em ^| \ }_B C$, then $p \not\perp q$ and $q \not\perp r$ imply $p \not\perp r$.

6.2.5. *Exercise.* Show that any type of rank ω is orthogonal to any type of finite rank.

End lecture 37

Lecture 38, Dec. 2, 2009

6.2.6. *Definition.* Let L be a subset of $\mathbf{X}(M)$ (for some model M of an arbitrary theory $\mathcal{T}$), which is invariant (as a set) under automorphisms (for example, L can $\mathbf{Y}(M)$ for some partial type $\mathbf{Y}$). L is said to be *modular* if any $A, B \subseteq L$ are independent over $\operatorname{acl}(A) \cap \operatorname{acl}(B)$. We say that L is modular over a set E if it is modular when considered in the theory $\mathcal{T}_E$.

We note that it is enough to check this condition for finite subsets A and B. In particular, if $\mathbf{Y}$ is a partial type, then the set of its realisations in one model is modular if and only if it is such in any other. Therefore, we may talk about $\mathbf{Y}$ being modular.

6.2.7. *Exercise.* Let R be an infinite field, M a non-zero R-vector space. Let $\mathcal{T}$ be the theory of M in a language with symbols for the abelian group structure on M, and with a function symbol $r : M \to M$ for any $r \in R$.

(1) Show that for any subset $A \subseteq M$, $\operatorname{acl}(A)$ is the vector-space generated by A.
(2) For subsets $A, B, C \subseteq M$, define $A \mathop{\smile\hskip-0.9em^| \ }_C B$ to hold if the intersection of the spaces generated by A and B is contained in the space generated by C. Show that this relation satisfies all the axioms of 5.2.11, as well as 2-uniqueness. Hence $\mathcal{T}$ is stable.
(3) By definition, M is modular.

6.2.8. *Example.* $ACF_{\Bbbk}$ is not modular: for example, if $\mathbf{Z}$ is a generic 3-dimensional subvariety of $\mathbb{A}^2 \times \mathbb{A}^2$ (for example, $x + w = yz$), $b = (b_1, b_2)$ a generic point of $\mathbf{Z}$, then $\operatorname{acl}(\Bbbk(b_1)) \cap \operatorname{acl}(\Bbbk(b_2)) = \operatorname{acl}(\Bbbk)$, but b_1 and b_2 are not generic. A similar argument shows that non of the (twisted) fixed fields in $ACFA$ is modular.

6.2.9. *Lemma.* Assume that the base in $\mathcal{T}_\sigma$ is algebraically closed. Let A and B be sets, $E = \operatorname{acl}_\sigma(A) \cap \operatorname{acl}_\sigma(B)$, $F \supseteq E$ independent from $A \cup B$ over E. Then $\operatorname{acl}_\sigma(A, F) \cap \operatorname{acl}_\sigma(B, F) = \operatorname{acl}_\sigma(F)$.

Proof. We first prove the analogous statement for $\mathcal{T}$. Let $a \in \operatorname{acl}(A, F) \cap \operatorname{acl}(B, F)$. We have $A \mathop{\smile\hskip-0.9em^| \ }_B F$ (since $E \subseteq \operatorname{acl}(B)$) and $a \in \operatorname{acl}(B, F)$, so $A \mathop{\smile\hskip-0.9em^| \ }_B F, a$. Likewise, $B \mathop{\smile\hskip-0.9em^| \ }_A F, a$. We claim that $F, a \mathop{\smile\hskip-0.9em^| \ }_E A, B$. Indeed, let p be the type of F and a over $A \cup B$. Then the above independence, it does not fork over both A and B. Hence the canonical base of p is in $\operatorname{acl}(A) \cap \operatorname{acl}(B) = E$. It now follows that $a \mathop{\smile\hskip-0.9em^| \ }_F A, B$, but then a belongs to both sides, so $a \in \operatorname{acl}(F)$.

We now go back to $\mathcal{T}_\sigma$. We may assume that all sets are definably closed (and in particular, σ structures). But then $\operatorname{acl}_\sigma = \operatorname{acl}$, so the statement follows from the same statement for $\mathcal{T}$. □

6.2.10. *Corollary.* Assume the base is algebraically closed in $\mathcal{T}_\sigma$.

(1) If R is modular, then for any $A \subseteq R$ and any B (not necessarily in R), $A \mathop{\smile\hskip-0.9em^| \ }_{\operatorname{acl}(A) \cap \operatorname{acl}(B)} B$.
(2) If S is a set of SU-rank 1 modular types, then $\bigcup_{p \in S} p(M)$ is modular (in particular, if $\mathbf{X}$ is a rank 1 definable set, such that each type in it is modular, then $\mathbf{X}$ itself is modular).
(3) If p and q are types of SU-rank 1 such that p is modular and q is non-orthogonal to p, then q is modular.

Proof. (1) As in 6.2.9, it is enough to prove the analogous statement for $\mathcal{T}$. We use the fact (mentioned in 3.2.35) that the canonical base of a type is algebraic over some set of realisations. Applying this to the type of A over B, we get the result.

(2) For simplicity we prove only that if a, b satisfy rank 1 modular types p and q, then for all tuples c we have $a, b \perp\!\!\!\perp_A c$, where $A = \operatorname{acl}(a, b) \cap \operatorname{acl}(c)$. The general case (where the left hand side has more elements) is similar. Furthermore, we may assume that the types p and q are non-orthogonal, since otherwise the result is easy.

Since p and q are non-orthogonal, we may choose an algebraically closed set B containing A and the parameters of p and q, independent from a, b and c, such that q has a realisation b_1, again independent from B, but dependent on a. Since p and q have rank 1, a and b_1 are inter-algebraic.

We now have by 6.2.9 that

$$\operatorname{acl}(b_1, b, B) \cap \operatorname{acl}(c, B) = \operatorname{acl}(a, b, B) \cap \operatorname{acl}(c, B) = B$$

On the other hand, by transitivity, a, b is not independent from c over B. Hence also b, b_1 is not independent from c over B. Since b and b_1 satisfy the same modular type, this is a contradiction to (1).

(3) Omitted.

□

6.2.11. *Definition.* A type p of *SU*-rank 1 is called *trivial* if, for any set A, any realisations $a_1, \dots, a_n$ of p which are independent from A, and such that $a_i \perp\!\!\!\perp_A a_j$ for $i < j$ are independent (i.e., any two disjoint subsets are independent).

6.2.12. *Exercise.* Show that p (of *SU*-rank 1) is trivial if and only if for any $a_1, \dots, a_n$ realising P, $\operatorname{acl}_\sigma(\{a_i\}) = \bigcup \operatorname{acl}_\sigma(a_i)$. Conclude that any trivial type is modular. The trivial types behave rather differently from the modular non-trivial ones, and these cases usually viewed as separate (and for this reason the dichotomy below is also called trichotomy).

6.2.13. *Exercise.* Let G be a definable group, and let p be a type in G. Show that p is not trivial.

6.2.14. *The dichotomy theorem.* Let p be a type of *SU*-rank 1 over an algebraically closed subset. Then p is modular if and only if it is orthogonal to all the fixed fields.

In addition, in characteristic 0, any modular type has a unique non-forking extension to any set.

We note that by 6.2.10, it follows that a *definable set* of rank 1 is modular if and only if it is orthogonal to the fixed fields. Also, the easy direction of the theorem follows immediately from 6.2.10 and 6.2.8. End lecture 38

Lecture 39, Dec. 4, 2009

6.2.15. *Bounded types.* Given the dichotomy theorem, it is desirable to have a criterion that allows us to determine on which side of the dichotomy a particular type falls. The paper contains such a criterion for types of σ-degree 1 in characteristic 0.

By definition, a given element a is of σ-degree 1 over $B = \operatorname{acl}_\sigma(B)$ if and only if a is not in B, but $\sigma(a)$ is algebraic over $B(a)$. Since B is a σ-structure, this implies that $\sigma^2(a)$ is algebraic over $B(\sigma(a))$, which itself is algebraic over $B(a)$. Therefore $\sigma^2(a)$, and similarly $\sigma^k(a)$ are all algebraic over $B(a)$.

We say that a is *bounded* over B if the sequence of (separable) degrees $[B(\sigma^k(a)) : B(a)]_s$ of these extensions is bounded (for all $k \in \mathbb{Z}$). Clearly, this is a property of $\text{tp}_\sigma(a/B)$, so we may talk about bounded and unbounded types (of σ-degree 1).

6.2.16. *Theorem.* Let p be a type of σ-degree 1. If p is bounded, then it is non-orthogonal to the field. In characteristic 0, the converse also holds.

To prove one direction of the theorem, we will need the following lemma, which says that if p is bounded, then we may take the bound to be 1.

6.2.17. *Lemma.* If $\deg_\sigma(a/C) = 1$, where $C = \text{acl}_\sigma(C)$, and a is bounded over C, then a is inter-algebraic over C with some b such that $C(b)$ is a difference sub-field.

6.2.18. *Some geometry.* The *projective space* $\mathbb{P}^n_{\Bbbk}$ of dimension n over $\Bbbk$ consists, as a set, of the set of lines through the origin in $n+1$-dimensional space. If p is a homogeneous polynomial in $n+1$ variables, the set of lines on which $p = 0$ is well defined. As in the affine case, the collection of all such sets forms a basis of closed sets on $\mathbb{P}^n$, the Zariski topology. The space defined in this way has a natural structure of an algebraic variety, in the sense that it has an open cover, such that each open set is homeomorphic to an affine variety, and the algebras of functions on the intersections are isomorphic.

Any closed irreducible subset of a variety (in particular of $\mathbb{P}^n$) is again a variety in this sense. A variety is *projective* if it is isomorphic to a closed subset of some projective space. It is a fact that any projective variety $\mathbf{X}$ is *complete*: For any other variety $\mathbf{Y}$, the projection to $\mathbf{Y}$ is closed. In particular, any map from $\mathbf{X}$ is closed.

As in the affine case, we may associate to each variety $\mathbf{X}$ its field of functions $K(\mathbf{X})$ (this will be equal to the field of functions of an open affine dense subset), and a dominant map of varieties determines a map of function fields (in the other direction). It is a fact that this process determines an equivalence of categories between smooth projective curves over $\Bbbk = \Bbbk^a$ with dominant maps and finitely generated field extensions of $\Bbbk$ of transcendence degree 1. In other words, to any such field K we may associate a unique smooth projective curve $\mathbf{X}_K$ (smoothness is a condition that implies that when considering complex points, any point has a neighbourhood in the classical topology diffeomorphic to $\mathbb{C}^n$. We will not use this condition explicitly. We note also that for complete curves, a map is dominant if and only if it is non-constant).

Given a smooth projective curve $\mathbf{X}$, points $x_1, \dots, x_m \in \mathbf{X}$ and integers $d_1, \dots, d_m$, let $V \subseteq K(\mathbf{X})$ be the subspace of rational functions that have poles of order at most d_i at x_i (and nowhere else). It turns out that this space is finite dimensional, and is non-zero precisely if the sum d of the d_i is non-zero (Hartshorne [16, p. IV.1.2]). Let $s_0, \dots, s_k$ be a basis for V. The smoothness of $\mathbf{X}$ implies that locally, one may always choose a rational function f that has zeroes of precisely order d_i at x_i (Hartshorne [16, p. II.6.2]), and so $fs_0, \dots, fs_k$ are well defined functions on $\mathbf{X}$ (locally), giving a map $F(x) = [fs_0(x) : \cdots : fs_k(x)]$ from $\mathbf{X}$ to $\mathbb{P}^k$, provided that the fs_i are not all simultaneously 0. It is a fact that given the points x_i, for large enough d this process indeed defines a map into projective space, which is an embedding. Furthermore, fixing d, changing to another basis amounts to composing the embedding with an automorphism of the projective space, given by a linear transformation of the corresponding vector space (in fact, all automorphisms of $\mathbb{P}^k$ come from linear automorphisms).

6.2.19. *proof of 6.2.16 (one direction).* By the lemma, and since we are working up to non-orthogonality, it is enough to prove that if E is a model of $ACFA$, and b is such that $E(b)$ is a difference field, then $\operatorname{tp}(b/E)$ is non-orthogonal to the fixed field. By the above, we may assume that b is a generic point of a smooth projective curve $\mathbf{X}$ over E. Since $\sigma(b) \in E(b)$, there is a function $g : \mathbf{X} \to \mathbf{X}^\sigma$, such that $\sigma(b) = g(b)$. Hence $b = \sigma^{-1}(g(b))$, so g is one-to-one, and since $\mathbf{X}$ is complete, it is an isomorphism.

Since E is a model of $ACFA$, there is a point $a \in E$ such that $\sigma(a) = g(a)$. Since $\mathbf{X}$ is projective, there is, by the overview above, a natural number m, and functions $f_0, \dots, f_n$ on $\mathbf{X}$ with at most poles of degree m on a (and no poles elsewhere), that define a closed embedding of $\mathbf{X}$ into $\mathbb{P}^n$. Applying σ to F, we get an embedding $F^\sigma : \mathbf{X}^\sigma \to \mathbb{P}^n$. Then $F^\sigma \circ g$ is another embedding with poles in a, and so by the above, $F^\sigma \circ g = A \circ F$ for some matrix A, viewed as an automorphism of $\mathbb{P}^n$. If A is the identity, we get that $\sigma(F(b)) = F(b)$, and since F is an isomorphism, $E(F(b))$ is isomorphic to $E(b)$.

If A is not the identity, since E is a model of $ACFA$, we may find a matrix B such that $B = \sigma(B)A$. Then setting $F' = B \circ F$, we reduce to the previous case.

End lecture 39

Lecture 40, Dec. 7, 2009

7. The finite rank theory

7.1. **Analysis of finite rank types.** Given the classification of types of rank 1, the next step is to analyse types of finite rank in terms of them. Ideally, we would like to say that we may find coordinates on $\operatorname{tp}_\sigma(a/E)$ with values in rank 1 sets, i.e., a tuple $(b_1, \dots, b_n)$, inter-definable with a over E, with such that each b_i has rank 1 over E. In reality, the situation is similar, but more complicated: inter-definable is replaced with inter-algebraic, and the b_i give a sequence of fibrations, rather than a product. This sequence is called the analysis of a.

We start with a result that gives us the beginning of the analysis.

7.1.1. *Proposition.* Let p be a non-algebraic type of finite rank over an algebraically closed base in $\mathcal{T}_\sigma$. Then p is non-orthogonal to some type q of rank 1. If q is modular, then it can be taken to the type of an element algebraic over a realisation of p.

Proof. We prove only the first part. Assume that $SU(p) = n > 0$. There is then a tuple a and an extension p_1 of p to a with $SU(p_1) = n - 1$. By first taking a non-forking extension of p, we may assume that a has finite rank. We may further assume that a is of minimal rank with this property.

Let b_1 and b_2 be two realisations of p_1, independent over a. We claim that b_1 and b_2 are not independent over 0. Otherwise, $\operatorname{tp}_\sigma(b_1/b_2)$ is a non-forking extension of p_1, we have that $SU(b_1/a, b_2) = SU(b_1/a) = n - 1$ and $SU(a/b_2) < SU(a)$ (since b_2 and a are not independent). This contradicts that minimality of $SU(a)$.

Hence $SU(b_2/b_1) = n - 1$. In particular, b_2 is independent from a over b_1 (since $SU(b_2/a, b_1) = n - 1$, because b_1 and b_2 are independent over a), and by assumption b_2 is independent from b_1 over a. As in the proof of 6.2.9, it follows that b_2 is independent from a and b_1 over $\operatorname{acl}_\sigma(a) \cap \operatorname{acl}_\sigma(b_1)$. Since b_2 and b_1 are not independent over 0, we must have some $c \in \operatorname{acl}_\sigma(a) \cap \operatorname{acl}_\sigma(b_1)$, non-algebraic, with $SU(b_1/c) = n - 1$. It follows that $SU(c) = 1$. Then $q = \operatorname{tp}_\sigma(c)$ is a type of rank 1 non-orthogonal to p (it is of rank n, realised by b_1, and $SU(b_1/c) = n - 1$). □

7.1.2. *Corollary.* For any element a of finite rank in $ACFA_0$ (over an algebraically closed base), there is a tuple $(b_1, \dots, b_n)$, inter-algebraic with a, such that either $\mathrm{tp}_\sigma(b_{i+1}/b_1, \dots, b_i)$ is modular of rank 1, or there is a tuple c independent from a over $b_1, \dots, b_i$ such that b_{i+1} is definable over c and the fixed field.

Proof. If a is algebraic over the base, we are done. Otherwise, by 7.1.1, $\mathrm{tp}_\sigma(a)$ is non-orthogonal to some type q of rank 1. By 6.2.14, q is either modular or non-orthogonal to the fixed field.

Assume that q is modular. Again by 7.1.1, we may assume that q is realised by $b \in \mathrm{acl}_\sigma(a)$, hence $SU(a/b) < SU(a)$.

If q is non-orthogonal to the fixed field, so is $\mathrm{tp}_\sigma(a)$. Hence, there is an element d_0 in the fixed field, and a tuple c_0 independent from a and d_0, such that a and d_0 are not independent over c_0, i.e., $d_0 \in \mathrm{acl}_\sigma(c_0, a)$ (but $d_0 \notin \mathrm{acl}_\sigma(c_0)$). Furthermore, we may replace d_0 by the code for its finite set of conjugates and assume that $d_0 \in \mathrm{dcl}_\sigma(c_0, a)$. Let C be the σ-structure generated by c_0, d_0, and let $b = \mathcal{C}b(C/\mathrm{acl}_\sigma(a))$ (in ACF). Then by definition $b \in \mathrm{acl}_\sigma(a)$, but b is not algebraic over the base: otherwise we would have that C is independent from a, but we know that c_0 is independent from a, so this would give that d_0 is independent from c_0 and a, contradicting that $d_0 \in \mathrm{acl}_\sigma(c_0, a)$. Hence again $SU(a/b) < SU(a)$. By 3.2.35, b is definable over some independent (over a) conjugates $C_1, \dots, C_k$ of C. We take c to be a finite subset over which this is true.

Since a has finite rank, and in each case the rank goes down, we are done by induction. □

7.1.3. *Remark.* It is possible to study more explicitly non-orthogonality to the fixed field and conclude that in the second case the conclusion can be replaced by "b_{i+1} is in the difference field generated by c and the fixed field".

7.1.4. *Definition.* Let $\mathbf{X}$ and $\mathbf{Y}$ be two (infinitely) definable sets (in any theory). $\mathbf{X}$ is *internal* to $\mathbf{Y}$ if there is a small set A of parameters such that each element of $\mathbf{X}$ is definable over A and $\mathbf{Y}(M)$ (for a saturated model M).

7.1.5. *Proposition.* Let a be a tuple of finite rank in $ACFA$. Assume that there is a tuple c independent from a, and a tuple b in the fixed field, such that a is in the difference field generated by b and c. Then $\mathrm{tp}_\sigma(a)$ is internal to the fixed field. End lecture 40

7.1.6. *Definition.* A type p (in a simple theory) is (fully) *stable* if any of its (possibly forking) extensions has a unique non-forking extension to each set. We note that in this case, p is stably embedded (exercise).

7.1.7. *Proposition.* Any stable type of rank 1 in $ACFA_0$ is modular.

7.1.8. *Corollary.* Any element of finite rank in $ACFA_0$ has an analysis with each step either stable and modular or internal to the fixed field (with a quantifier free formula).

7.1.9. *Theorem.* A type of finite rank in $ACFA_0$ is stable if and only if any of its extensions is orthogonal to the fixed field. A formula that is orthogonal to the fixed field is stable, and the induced theory on it is modular.

7.2. **Internality.** We now study what happens in the setting of internality (7.1.4). More complete accounts are available in Hrushovski [18], Hrushovski [17] or **defaut**.

In this sub-section, we assume that $\mathbf{X}$ is internal to $\mathbf{C}$ (in an arbitrary theory), and for simplicity, we assume that $\mathbf{X}$ (and hence $\mathbf{C}$) are definable, rather than type definable. It follows that $\mathbf{X}$ is rather close to $\mathbf{C}$. For example, if the set of parameters A in the definition is empty, we get that that $\mathbf{X}$ is a quotient of a definable subset of $\mathbf{C}^k$. In the general case, we have non-trivial automorphism group, but we will see that it is "small".

We assume in addition that $\mathbf{C}$ is stably embedded. We note that given any $\mathbf{C}$, there is a unique smallest (ind-) definable set $\tilde{\mathbf{C}}$, such that $\tilde{\mathbf{C}}$ is stably embedded. Furthermore, any automorphism fixing $\mathbf{C}$ pointwise, also fixes $\tilde{\mathbf{C}}$ in the same way. Thus, there is no harm in passing to $\tilde{\mathbf{C}}$.

7.2.1. *Exercise.* Assume EI and that $\mathbf{C}$ is stably embedded. Show that for any type p over $\mathbf{C}(M)$ (where M) is any model) there is a small subset $C_0 \subseteq \mathbf{C}(M)$ such that p is the only extension of its restriction to C_0. Show that if M is strongly homogeneous, then the automorphism group of M over $\mathbf{C}(M)$ acts transitively on the realisations of p. Thus, a stably embedded $\mathbf{C}$ behaves like a small subset.

7.2.2. The assumption says that there is a parameter a, and an a-definable surjective map $g_a : \mathbf{C}_a \to \mathbf{X}$, where $\mathbf{C}_a$ is a (type) definable subset of $\mathbf{C}^k$. Since we assume elimination of imaginaries, we may assume that g_a is a bijection.

Given a model M containing a, let $p = \mathrm{tp}(a/\mathbf{C}(M))$. Since $\mathbf{C}$ is stably embedded, the code for $\mathbf{C}_a$ is in $\mathbf{C}$, and so is determined in the type p. Since the property of being a bijective function is also first order, we get that g_b is a bijection from $\mathbf{C}_a$ to $\mathbf{X}$ for all b realising p. From now on we write $\mathbf{C}$ instead of $\mathbf{C}_a$. Thus, we have a definable function $g : \mathbf{C} \times p \to \mathbf{X}$, which gives a bijection over each realisation of p.

7.2.3. Let τ be an automorphism of M fixing $\mathbf{C}(M)$. Then τ maps $p(M)$ to itself. If τ fixes a realisation a of p, then $\tau(g(c,a)) = g(\tau(c), \tau(a)) = g(c,a)$ for all $c \in \mathbf{C}(M)$, and since g_a is surjective, this means that $\mathbf{X}$ is fixed pointwise by τ.

It follows that if we define $\mathbf{G}(M) = \mathrm{Aut}({}^{\mathbf{X}(M), p(M), \mathbf{C}(M)}\!/_{\mathbf{C}(M)})$, where we consider the induced structure on these sets, then $\mathbf{G}(M)$ acts freely on p. If M is saturated, then since p is a type over $\mathbf{C}$ (and $\mathbf{C}$ is stably embedded), it also acts transitively, so p is a $\mathbf{G}$-*torsor*, i.e., then combined map $\mathbf{G} \times p \to p \times p$, obtained from the projection and the action, is a definable isomorphism. Note that this implies a bijection on the level of points (of a sufficiently saturated model), but in the interesting case, it will not be a bijection on points over 0. We will see that the assumption on M is unnecessary, and the action is transitive for any M.

7.2.4. Let a and b be realisations of p. The composition $g_b \circ {g_a}^{-1}$ is a bijection from $\mathbf{X}$ to itself. We claim that this map coincides with the (unique) automorphism $\tau \in \mathbf{G}(M)$ taking a to b. Indeed, if $x \in \mathbf{X}(M)$ is any element, then $x = g_a(c)$ for some $c \in \mathbf{C}(M)$

$$\tau(x) = \tau(g_a(c)) = \tau(g(a,c)) = g(\tau(a), \tau(c)) = g_b(c) = g_b({g_a}^{-1}(x)) \qquad (13)$$

Likewise, it is easy to see that the action of τ on p is given by $c \mapsto d$, where $g_d = g_b \circ {g_a}^{-1} \circ g_c$.

We now know that if a and b are conjugate under the automorphism group, then $g_b \circ {g_a}^{-1}$ defines an automorphism. We claim that this holds in general. To show that, we may embed

7.2.5. *Corollary.* $\mathbf{G}$ is represented by an ω-definable group, i.e., there is an ω-definable group $\mathbf{G}_0$, and a definable group action of $\mathbf{G}_0$ on $\mathbf{X}$ and p, identifying $\mathbf{G}_0(M)$ with $\mathbf{G}(M)$ for any model M. Furthermore, $\mathbf{G}_0$ is the intersection of definable groups $\mathbf{G}_i$, and p is the intersection of a corresponding system p_i of definable torsors of $\mathbf{G}_i$.

Proof. $\mathbf{G}_0$ is given as a definable set by the quotient of $p \times p$ by the relation of giving the same function on $\mathbf{X}$. Thus, any two distinct elements of $\mathbf{G}_0$ determine distinct bijections of $\mathbf{X}$ with itself. The group operation is given by composition of functions, and the action is described above, as well as the proof that $\mathbf{G}_0$ is identified with $\mathbf{G}$. The last statement follows by approximating p by finite sets of formulas. The elements of $\mathbf{G}_i$ are "automorphisms" preserving only the structure in the approximation p_i. □

7.2.6. *Example.* Let $\mathbf{C}$ be an algebraically closed field, and let $\mathbf{X}$ be a definable vector space over $\mathbf{C}$, of dimension n. Let $b : \mathbf{X} \times \mathbf{X} \to \mathbf{C}$ be a non-degenerate symmetric bilinear form. Then $\mathbf{X}$ is internal to $\mathbf{C}$, as can be shown by picking any basis a of $\mathbf{X}$ over $\mathbf{C}$. Given such an $a = (a_1, \dots, a_n)$, let $c_{i,j} = b(a_i, a_j)$. These statements belong to the type of a, and in fact determine that type. The bijection determined by a is the one taking the standard basis of $\mathbf{C}^n$ to $a_1, \dots, a_n$. The element corresponding to the pair of realisations (a, b) is therefore the unique linear map of $\mathbf{X}$ that takes a to b. The fact that they both satisfy the equations above (for the same $c_{i,j}$) is equivalent to this linear map preserving the bilinear form. Hence the definable group is the group of all b-preserving linear automorphisms of $\mathbf{X}$, as one expects. In this example, the type is determined by a finite number of formulas, so the group is definable, rather than just ω-definable.

7.2.7. *Example.* Let $\mathbf{X}$ be the set in $ACFA$ defined by $\sigma(x) = Ax$ for some (invertible) $n \times n$ matrix A over the base. By 2.3.6, the set of solutions is an n-dimensional vector space over the fixed field $\mathbf{C}$. Hence the group of automorphisms of $\mathbf{X}$ over $\mathbf{C}$ is an ω-definable subgroup of $GL(\mathbf{X})$ (and might not be definable).

References

[1] James Ax. "The elementary theory of finite fields". In: *Ann. of Math. (2)* 88 (1968), pp. 239–271. ISSN: 0003-486X (cit. on pp. 4, 12).

[2] John T. Baldwin. *Fundamentals of stability theory.* Perspectives in Mathematical Logic. Berlin: Springer-Verlag, 1988, pp. xiv+447. ISBN: 3-540-15298-9 (cit. on p. 19).

[3] John T. Baldwin and Saharon Shelah. "Model companions of T_{Aut} for stable T". In: *Notre Dame J. Formal Logic* 42.3 (2001), 129–142 (2003). ISSN: 0029-4527 (cit. on p. 26).

[4] zlem Beyarslan. "Random hypergraphs in pseudofinite fields". In: *Journal of the Institute of Mathematics of Jussieu* (2006), pp. 1–19. DOI: `10.1017/S1474748009000073`. arXiv: `math/0701029` (cit. on p. 21).

[5] zlem Beyarslan and Ehud Hrushovski. *On algebraic closure in pseudofinite fields.* Sept. 2009. arXiv: `0909.2189` (cit. on p. 14).

[6] Elisabeth Bouscaren, ed. *Model theory and algebraic geometry.* Lecture Notes in Mathematics 1696. An introduction to E. Hrushovski's proof of the geometric Mordell-Lang conjecture. Berlin: Springer-Verlag, 1998. ISBN: 3-540-64863-1 (cit. on p. 39).

[7] Steven Buechler. *Essential stability theory.* Perspectives in Mathematical Logic. Berlin: Springer-Verlag, 1996. ISBN: 3-540-61011-1 (cit. on pp. 16, 19, 41).

[8] C. C. Chang and H. J. Keisler. *Model theory.* 3rd ed. Studies in Logic and the Foundations of Mathematics 73. Amsterdam: North-Holland Publishing Co., 1990, pp. xvi+650. ISBN: 0-444-88054-2 (cit. on p. 36).

[9] Zo Chatzidakis. "Model theory of difference fields". In: *The Notre Dame lectures.* Vol. 18. Lect. Notes Log. Urbana, IL: Assoc. Symbol. Logic, 2005, pp. 45–96 (cit. on p. 5).

[10] Zo Chatzidakis and Ehud Hrushovski. "Model theory of difference fields". In: *Trans. Amer. Math. Soc.* 351.8 (1999), pp. 2997–3071. ISSN: 0002-9947 (cit. on pp. 1, 3, 4, 38, 41, 42).

[11] Zo Chatzidakis and Anand Pillay. "Generic structures and simple theories". In: *Ann. Pure Appl. Logic* 95.1-3 (1998), pp. 71–92. ISSN: 0168-0072 (cit. on pp. 5, 25, 26).

[12] Zo Chatzidakis et al. "Model theory of difference fields. II. Periodic ideals and the trichotomy in all characteristics". In: *Proc. London Math. Soc. (3)* 85.2 (2002), pp. 257–311. ISSN: 0024-6115 (cit. on pp. 4, 38).

[13] Lou van den Dries and K. Schmidt. "Bounds in the theory of polynomial rings over fields. A nonstandard approach". In: *Invent. Math.* 76.1 (1984), pp. 77–91. ISSN: 0020-9910 (cit. on p. 9).

[14] Jean-Louis Duret. "Les corps faiblement algbriquement clos non sparablement clos ont la proprit d'indpendence". In: *Model theory of algebra and arithmetic.* Proceedings of the Conference on Applications of Logic to Algebra and Arithmetic. Ed. by L. Pacholski and J. Wierzejewski. Lecture Notes in Mathematics 834. Berlin: Springer, 1980, pp. 136–162. ISBN: 3-540-10269-8 (cit. on p. 21).

[15] David Eisenbud. *Commutative algebra. With a view toward algebraic geometry.* Graduate Texts in Mathematics 150. New York: Springer-Verlag, 1995. ISBN: 0-387-94268-8; 0-387-94269-6 (cit. on p. 12).

[16] Robin Hartshorne. *Algebraic geometry.* Graduate Texts in Mathematics, No. 52. New York: Springer-Verlag, 1977. ISBN: 0-387-90244-9 (cit. on pp. 7, 45).

[17] Ehud Hrushovski. "Computing the Galois group of a linear differential equation". In: *Differential Galois theory.* Ed. by Teresa Crespo and Zbigniew Hajto. Banach Center Publications 58. May 28–June 1, 2001. Warsaw: Polish Academy of Sciences Institute of Mathematics, 2002, pp. 97–138 (cit. on pp. 4, 48).

[18] Ehud Hrushovski. *Groupoids, imaginaries and internal covers.* Mar. 2011. arXiv: `math/0603413` (cit. on pp. 5, 31, 48).

[19] Ehud Hrushovski. *The Elementary Theory of the Frobenius Automorphisms.* arXiv: `math/0406514` (cit. on p. 4).

[20] Ehud Hrushovski. "The Manin-Mumford conjecture and the model theory of difference fields". In: *Ann. Pure Appl. Logic* 112.1 (2001), pp. 43–115. ISSN: 0168-0072 (cit. on pp. 1, 2).

[21] Ehud Hrushovski and Boris Zilber. “Zariski geometries”. In: *J. Amer. Math. Soc.* 9.1 (1996), pp. 1–56. ISSN: 0894-0347. DOI: `10.1090/S0894-0347-96-00180-4` (cit. on p. 39).

[22] Byunghan Kim and Anand Pillay. “Simple theories”. In: *Ann. Pure Appl. Logic* 88.2-3 (1997). Joint AILA-KGS Model Theory Meeting (Florence, 1995), pp. 149–164. ISSN: 0168-0072 (cit. on pp. 19, 34).

[23] Daniel Lascar. “Autour de la proprit du petit indice”. In: *Proc. London Math. Soc. (3)* 62.1 (1991), pp. 25–53. ISSN: 0024-6115 (cit. on p. 29).

[24] Daniel Lascar. *Stability in model theory.* Pitman Monographs and Surveys in Pure and Applied Mathematics 36. Translated from the French by J. E. Wallington. Harlow: Longman Scientific & Technical, 1987, pp. viii+193. ISBN: 0-582-99463-2 (cit. on p. 19).

[25] David Marker. *Model theory: An introduction.* Graduate Texts in Mathematics 217. New York: Springer-Verlag, 2002. ISBN: 0-387-98760-6 (cit. on pp. 5, 15).

[26] David Marker et al. *Model theory of fields.* 2nd ed. Lecture Notes in Logic 5. La Jolla, CA: Association for Symbolic Logic, 2006. ISBN: 978-1-56881-282-3; 1-56881-282-5 (cit. on p. 15).

[27] Alice Medvedev. “Minimal sets in ACFA”. PhD thesis. UC Berkeley, 2007. URL: `http://www.math.uic.edu/~alice/MedvedevThesis.pdf` (cit. on p. 3).

[28] Alice Medvedev and Thomas Scanlon. *Polynomial dynamics.* 2009. arXiv: `0901.2352` (cit. on p. 3).

[29] James S. Milne. *Algebraic Geometry.* Course lecture notes. Mar. 2008. URL: `http://jmilne.org/math/` (cit. on p. 7).

[30] Elisabetta Pastori and Pablo Spiga. *Failure of n-uniqueness: A family of examples.* Sept. 2009. arXiv: `0909.0589` (cit. on p. 32).

[31] Anand Pillay. *Applied stability theory.* Course lecture notes, differential fields. Dec. 2003. URL: `http://www.math.uiuc.edu/People/pillay/lecturenotes.applied.pdf` (cit. on p. 39).

[32] Anand Pillay. *Geometric stability theory.* Oxford Logic Guides 32. Oxford Science Publications. New York: The Clarendon Press Oxford University Press, 1996. ISBN: 0-19-853437-X (cit. on pp. 19, 25).

[33] Anand Pillay. *Model theory.* Course lecture notes. 2002. URL: `http://www.math.uiuc.edu/People/pillay/lecturenotes_modeltheory.pdf` (cit. on p. 5).

[34] Anand Pillay. *Stability theory.* Course lecture notes. Sept. 2003. URL: `http://www.math.uiuc.edu/People/pillay/lecturenotes.stability.pdf` (cit. on p. 19).

[35] Tristram de Piro et al. “Constructing the hyperdefinable group from the group configuration”. In: *J. Math. Log.* 6.2 (2006), pp. 121–139. ISSN: 0219-0613. DOI: `10.1142/S0219061306000530`. arXiv: `math/0508583` (cit. on p. 32).

[36] Jean-Pierre Serre and John Tate. “Good reduction of abelian varieties”. In: *Ann. of Math. (2)* 88 (1968), pp. 492–517. ISSN: 0003-486X (cit. on p. 2).

[37] Saharon Shelah. *Classification theory and the number of nonisomorphic models.* 2nd ed. Studies in Logic and the Foundations of Mathematics 92. Amsterdam: North-Holland Publishing Co., 1990. ISBN: 0-444-70260-1 (cit. on p. 19).

[38] J. R. Shoenfield. “The problem of predicativity”. In: *Essays on the foundations of mathematics*. Ed. by Y. Bar-Hillel et al. Dedicated to A. A. Fraenkel on his seventieth anniversary. The Hebrew University, Jerusalem: The Magnes Press, 1961, pp. 132–139 (cit. on p. 16).

[39] Frank O. Wagner. *Simple theories*. Mathematics and its Applications 503. Dordrecht: Kluwer Academic Publishers, 2000, pp. xii+260. ISBN: 0-7923-6221-7 (cit. on p. 40).

[40] Andr Weil. *Varits abliennes et courbes algbriques*. Actualits Sci. Ind., Publ. Inst. Math. Univ. Strasbourg 8 (1946) 1064. Hermann & Cie., Paris, 1948, p. 165 (cit. on p. 3).

Index

Department of Pure Math, Notre-Dame University, IN, USA
E-mail address: `mailto:kamensky.1@nd.edu`
URL: `http://mkamensky.notlong.com`