Notes on the model theory of finite and pseudo-finite fields

Zoé Chatzidakis, CNRS - Université Paris 7.

These notes contain the material covered during a mini-course I taught at the Universidad Autónoma de Madrid, 15 - 25 November 2005. The course was funded by the Marie Curie RTN MODNET and by the Universidad Autónoma, and I thank them for their financial support and hospitality. I also thank the audience for their enthusiasm and comments, and Margarita Otero who made this course possible.

1. Preliminary results in model theory

In this section we recall some basic model theoretic notions and results which we will assume throughout the course. We let $\mathcal{L}$ be a first-order language, T an $\mathcal{L}$-theory, A, B, C, ... $\mathcal{L}$-structures.

(1.1) Morphisms. Let A and B be $\mathcal{L}$-structures, $F : A \to B$ a map.

(1) F is a *homomorphism* (of $\mathcal{L}$-structures) if for all n, n-ary relation R, n-ary function f and n-tuple a in A.

$$A \models R(a) \Rightarrow B \models R(F(a)) \text{ and } f(F(a)) = F(f(a)).$$

(2) F is an *embedding* iff F is an injective homomorphism and for every n, n-ary relation R and n-tuple a in A, we have

$$A \models R(a) \iff B \models R(F(a)).$$

(3) F is an *isomorphism* iff F is an embedding and is surjective.

(4) F is an *elementary embedding* iff $F(A) \prec B$. Equivalently, if for every formula $\varphi(x)$ and tuple a in A,

$$A \models \varphi(a) \iff B \models \varphi(F(a)).$$

(5) A *partial isomorphism* $g : A \to B$ is an embedding g of some substructure A_0 of A into B.

(1.2) A test for being an elementary substructure. Let $A \subseteq B$ be $\mathcal{L}$-structures. Then $A \prec B$ if and only if for every formula $\varphi(x, y)$ and tuple a in A,

$$B \models \exists y \; \varphi(a, y)$$

if and only if there is some tuple b in A such that

$$B \models \varphi(a, b).$$

(1.3) Some consequences. The following results are easy:

(1) If $A \prec B$ and $B \prec C$ then $A \prec C$.

(2) If $A \prec C$, $B \prec C$ and $A \subseteq B$, then $A \prec B$.

(3) Let $(A_i)_{i\in\mathbb{N}}$ be an increasing chain with $A_i \prec A_{i+1}$ for every i, and let $A_\omega = \bigcup_{i\in\mathbb{N}} A_i$. Then $A_i \prec A_\omega$ for every i.

(1.4) Some preservation results. We start with a compactness result.

Proposition. Let T_1 and T_2 be theories, with $T_1 \cup T_2$ consistent, and let Δ be a set of sentences closed under finite disjunctions. The following conditions are equivalent¿

(1) There is $\Gamma \subseteq \Delta$ such that $T_1 \cup \Gamma$ axiomatizes $T_1 \cup T_2$.

(2) For all models A and B of T_1, if $A \models T_2$ and B satisfies all sentences of Δ satisfied by A, then $B \models T_2$.

Proof. (1) implies (2) is clear. Assume that (2) holds, and let $\Gamma = \{\psi \in \Delta \mid T_1 \cup T_2 \vdash \psi\}$. We need to show that $T_1 \cup \Gamma \vdash T_2$, i.e., that every model of $T_1 \cup \Gamma$ is a model of T_2. Let $B \models T_1 \cup \Gamma$, and let

$$\Sigma = \{\neg\psi \mid \psi \in \Delta,\ B \models \neg\psi\}.$$

We will show that $T_1 \cup T_2 \cup \Sigma$ is consistent. Otherwise, using compactness, there are $\neg\psi_1, \dots, \neg\psi_m \in \Sigma$, $\varphi \in T_2$ such that $T_1 \cup \{\neg\psi_1, \dots, \neg\psi_m, \varphi\}$ is inconsistent, i.e.,

$$T_1 \cup \{\varphi\} \vdash \psi_1 \vee \cdots \vee \psi_m.$$

Since Δ is closed under disjunction, the sentence $\psi_1 \vee \cdots \vee \psi_m$ is in Γ, and therefore cannot be in Σ. This gives us the desired contradiction, and shows that $T_1 \cup T_2 \cup \Sigma$ is consistent.

Let A be a model of $T_1 \cup T_2 \cup \Sigma$. If $\psi \in \Delta$ holds in A, then $\neg\psi \notin \Sigma$, and therefore $B \models \psi$. By condition (2), we obtain that $B \models T_2$.

Corollary. Let T be a theory, $\varphi(x)$ a formula such that $T \cup \{\exists x \varphi(x)\}$ is consistent. Let Δ be a set of formulas in the variables x, closed under disjunctions. The following conditions are equivalent:

(1) There are formulas $\psi_1(x), \dots, \psi_m(x) \in \Delta$ such that

$$T \vdash \forall x\ [\varphi(x) \leftrightarrow (\psi(x) \wedge \cdots \psi_m(x))].$$

(2) Whenever A and B are models of T, and a, b are tuples in A, B respectively, if $A \models \varphi(a)$ and every formula $\psi(x) \in \Delta$ which is satisfied by a in A is also satisfied by b in B, then $B \models \varphi(b)$.

Proof. Assume x is of length n, and enlarge the language by adding an n-tuple c of (new) constants. We apply the theorem to the sentence $\varphi(c)$ and obtain formulas $\psi_1(x), \dots, \psi_m(x) \in \Delta$ such that $T \cup \{\psi_1(c), \dots \psi_m(c)\}$ axiomatizes $T \cup \{\varphi(c)\}$, in other words

$$T \vdash \varphi(c) \leftrightarrow (\psi_1(c) \wedge \cdots \wedge \psi_m(c)).$$

Since the symbols in c do not appear in T, we then obtain

$$T \vdash \forall x\ [\varphi(x) \leftrightarrow (\psi(x) \wedge \cdots \psi_m(x))].$$

(1.5) The above results can be used to show the following preservation theorem. For more details, see e.g. Chang and Keisler. We let T be a theory, $\varphi(x)$ a formula.

(1) The following conditions are equivalent:

(a) For all models A and B of T, homomorphism $F : A \to B$ and tuple a in A, if $A \models \varphi(a)$ then $B \models \varphi(F(a))$.

(b) There is a positive formula $\psi(x)$ (that is, ψ is built from atomic formulas using $\wedge$, $\vee$, $\exists$, $\forall$, but not using $\neg$) such that $T \vdash \forall x\ \varphi(x) \leftrightarrow \psi(x)$.

(2) The following conditions are equivalent:

(a) For all models $A \subseteq B$ of T and tuple a in A, if $A \models \varphi(a)$ then $B \models \varphi(b)$.

(b) There is an existential formula $\psi(x)$ such that $T \vdash \forall x\ \varphi(x) \leftrightarrow \psi(x)$.

(3) The following conditions are equivalent:

(a) For all models $A \subseteq B$ of T and tuple a in A, if $B \models \varphi(a)$ then $A \models \varphi(a)$.

(b) There is a universal formula $\psi(x)$ such that $T \vdash \forall x\ \varphi(x) \leftrightarrow \psi(x)$.

(4) The following conditions are equivalent:

(a) For every increasing chain $(A_i)_{i\in\mathbb{N}}$ of models of T and tuple a in A_0, if $A_i \models \varphi(a)$ for every i, then $\bigcup_{i\in\mathbb{N}} A_i \models \varphi(a)$.

(b) There is a $\forall\exists$-formula $\psi(x)$ such that $T \vdash \forall x\ \varphi(x) \leftrightarrow \psi(x)$.

(1.6) Model complete theories, elimination of quantifiers.

Recall that a theory T is *model complete* iff whenever A and B are models of T and $A \subseteq B$, then $A \prec B$.

The theory T *eliminates quantifiers* iff for every formula $\varphi(x)$ there is a quantifier-free formula $\psi(x)$ such that $T \vdash \forall x\ \varphi(x) \leftrightarrow \psi(x)$.

Remarks. Recall that if A is a subset of an $\mathcal{L}$-structure M, the language $\mathcal{L}(A)$ is obtained by adding to $\mathcal{L}$ new constant symbols, one for each element of A. We denote by $\Delta(A)$ (or $\Delta_M(A)$) the *quantifier-free diagram of* A, i.e., the set of quantifier-free sentences of $\mathcal{L}(A)$ which hold in M. Using compactness, one easily obtains

(1) T is model complete if and only if for every model A of T, the theory $T \cup \Delta(A)$ is complete (in $\mathcal{L}(A)$).

(2) T eliminates quantifiers if and only if for every subset A of a model M of T, the theory $T \cup \Delta_M(A)$ is complete (in $\mathcal{L}(A)$).

(3) If T is model complete, then it has an axiomatization by $\forall\exists$-sentences.

Proof. (1) B is a model of $T \cup \Delta(A)$ if and only if B contains an isomorphic copy of A. If $T \cup \Delta(A)$ is complete and $B \supseteq A$, then $(B, a)_{a\in A} \equiv (A, a)_{a\in A}$, i.e., $A \prec B$.

(2) The left-to-right implication is easy, and the right-to-left implication follows from Proposition (1.4).

(3) The union of an increasing chain of models of T is a model of T by (1.3). Apply (1.5).

(1.7) Proposition. Let T be a theory.

(1) To show that T is model complete, it is enough to show that whenever $A \subseteq B$ are models of T and $\varphi(x, y)$ is a quantifier-free formula, a is a tuple of elements of A, then

$$B \models \exists y\ \varphi(a, y) \iff A \models \exists y\ \varphi(a, y).$$

(One then says that A *is existentially closed in* B, also denoted by $A \prec_1 B$.)

(2) T is model complete if and only if every formula $\varphi(x)$ is equivalent modulo T to an existential formula $\psi(x)$ (i.e., $T \vdash \forall x(\varphi(x) \leftrightarrow \psi(x))$), if and only if every formula is equivalent modulo T to a universal formula.

(3) To show that T eliminates quantifiers, it is enough to show that if C is a substructure of the models A and B of T, then there is a model D of T which contains A and B as elementary substructures.

(4) (Assume that $\mathcal{L}$ is countable) To show that T eliminates quantifiers, it is enough to find a model M of T with the following properties:

(a) Every countable model of T embeds in M.

(b) If $F : A_0 \to B_0$ is a partial isomorphism, with A_0, B_0 finite subsets of M, and $c \in M$, then there is some partial isomorphism extending F and having c in its domain.

Part (a) follows from a more general result:

(1.8) Back and forth Theorem. Let A and B be $\mathcal{L}$-structures. Assume that there is a family $\mathcal{I}$ of partial isomorphism from A to B such that

(a) (Forth) Pour tout $a \in A$ and $F \in \mathcal{I}$, there is $F' \in \mathcal{I}$ which extends F and with $a \in dom(F')$.

(b) (Back) Pour tout $b \in B$ and $F \in \mathcal{I}$, there is $F' \in \mathcal{I}$ which extends F and with $b \in Im(F')$.

Then $A \equiv B$.

Definition. A family of partial isomorphisms between two $\mathcal{L}$-structures has the *property of the back-and-forth* if it satisfies the above conditions (a) and (b).

(1.9) *Proof of (1.7)* (1) Let $A \subseteq B$ be models of T. Since $A \prec_1 B$, $Th(A,a)_{a\in A} \cup \Delta(B)$ is consistent. Indeed, our hypothesis says that all universal formulas of $\mathcal{L}(A)$ which hold in A also hold in B. Let $A_1 \models Th(A,a)_{a\in A} \cup \Delta(B)$, we may assume that A_1 is an elementary extension of A which contains B. Then $B \prec_1 A_1$, and using the same reasoning, there is an elementary extension B_1 of B which contains A_1. Repeating the same argument, we build by induction a chain

$$A_0 = A \subseteq B_0 = B \subseteq A_1 \subseteq B_1 \subseteq \cdots \subseteq A_n \subseteq B_n \subseteq A_{n+1} \subseteq \cdots$$

where $A_n \prec A_{n+1}$ and $B_n \prec B_{n+1}$. Then $C = \bigcup_{n\in\mathbb{N}} A_n = \bigcup_{n\in\mathbb{N}} B_n$ is an elementary extension of A and of B, whence $A \prec B$ (see (1.3)).

(2) If every formula is equivalent modulo T to an existential formula, it is also equivalent modulo T to a universal formula (since the negation of an existential formula is a universal formula). This shows the equivalence of the last two conditions. The equivalence of the first two follows from the definition of model completeness and the preservation result (1.5)(2).

(3) If T does not eliminate quantifiers, there is a substructure C of models A and B of T such that $(A,c)_{c\in C} \not\equiv (B,c)_{c\in C}$. Then clearly one cannot find D in which both A and B elementarily embed.

(4) We will show the theorem of the back and forth. One shows by induction on the number of quantifiers of a formula $\varphi(x)$ in prenex form that if $F \in \mathcal{I}$ and a is a tuple in $dom(F)$, then

$$A \models \varphi(a) \iff B \models \varphi(F(a)).$$

For quantifier-free formulas, this is exactly the definition of partial isomorphism. Write the formula $\varphi(x)$ as $\exists y\ \psi(x,y)$, y a single variable, and assume that the result holds for $\psi(x,y)$. Let $F \in \mathcal{I}$, $a \in dom(F)$, and assume that $A \models \varphi(a)$. Then there is $c \in A$ such that $A \models \psi(a,c)$. By hypothesis (forth), there is $F' \in \mathcal{I}$ extending F and with $c \in dom(F')$. By induction hypothesis, $B \models \psi(F'(a), F'(c))$, i.e., $B \models \varphi(F(a))$. One shows the other direction in the same way, using now the "back" direction.

(1.10) Examples

(1) The theory of the field $\mathbb{C}$ (or more generally, the theory of algebraically closed fields) eliminates quantifiers in the language of rings $\{+, -, \cdot, 0, 1\}$.

(2) Let us now consider the field $\mathbb{R}$ in the language of rings. One can show that its theory is model complete. However, it does not eliminate quantifiers. Indeed let $C = \mathbb{Q}(\alpha)$, where $\alpha^2 = 2$. One can embed C into $\mathbb{R}$ in two different ways: by sending α to $\sqrt{2}$ (the positive square root of 2) or to $-\sqrt{2}$. But clearly $(\mathbb{R}, \sqrt{2}) \not\equiv (\mathbb{R}, -\sqrt{2})$ since $\sqrt{2}$ has a square root in $\mathbb{R}$, while $-\sqrt{2}$ does not.

(1.11) Ultraproducts

Definitions. Let I be a set, A_i, $i \in I$, a family of $\mathcal{L}$-structures, and $\mathcal{F}$ a subset of $\mathcal{P}(I)$ (the set of subsets of I).

(1) $\mathcal{F}$ is a *filter* (on I) iff: (i) $\emptyset \notin \mathcal{F}$; (ii) if $X, Y \in \mathcal{F}$ then $X \cap Y \in \mathcal{F}$; (iii) if $X \in \mathcal{F}$ and $Y \supseteq X$ then $Y \in \mathcal{F}$.

(2) $\mathcal{F}$ is an *ultrafilter* iff it is a maximal filter, i.e., is contained properly in no filter. One shows easily that a filter $\mathcal{F}$ is an ultrafilter if and only if, for every $X \subseteq I$, either X or $I \setminus X$ is in $\mathcal{F}$.

(3) A filter $\mathcal{F}$ is *principal* iff there is some $i \in I$ such that $\{i\} \in \mathcal{F}$. If there is no such i, it is called *non-principal*.

(4) We define an $\mathcal{L}$-structure on the Cartesian product $\prod_{i \in I} A_i$ as follows. We view an element a of $\prod_{i \in I} A_i$ as a function from I to the disjoint union of the A_i's, whose value at i is in A_i. If f is an n-ary function symbol, R is an n-ary function symbol, and $(a_1, \dots, a_n) \in \prod_{i \in I} A_i$, then $f(a_1, \dots, a_n)(i) = f(a_1(i), \dots, a_n(i))$, and $\prod_{i \in I} A_i \models R(a_1, \dots, a_n)$ iff $A_i \models R(a_1(i), \dots, a_n(i))$ for every $i \in I$. Finally, the interpretation of a constant c is the function which to i associates the interpretation of c in A_i.

(5) Let $\mathcal{F}$ be a filter on I. We define an equivalence relation $\equiv_{\mathcal{F}}$ on $\prod_{i \in I} A_i$ by setting

$$a \equiv_{\mathcal{F}} b \iff \{i \in I \mid a(i) = b(i)\} \in \mathcal{F}.$$

The equivalence class of $a \in \prod_{i \in I} A_i$ will be denoted by $[a]_{\mathcal{F}}$, and the set of equivalence classes by $\prod_{i \in I} A_i/\mathcal{F}$. $\prod_{i \in I} A_i/\mathcal{F}$ has a natural $\mathcal{L}$-structure: the constant c is interpreted by $[c]_{\mathcal{F}}$; $f([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}}) = [f(a_1, \dots, a_n)]_{\mathcal{F}}$, and $R([a_1]_{\mathcal{F}}, \dots, [a_n]_{\mathcal{F}})$ holds iff $\{i \in I \mid A_i \models R(a_1(i), \dots, a_n(i))\} \in \mathcal{F}$. The structure $\prod_{i \in I} A_i/\mathcal{F}$ is called the *reduced product of the structures A_i with respect to $\mathcal{F}$*. If $\mathcal{F}$ is an ultrafilter, then $\prod_{i \in I} A_i/\mathcal{F}$ is called the *ultraproduct* of the A_i with respect to $\mathcal{F}$. If all A_i are equal to the same structure A, then we talk of *reduced power* and of *ultrapower* of A.

(6) Observe that the natural map $\prod_{i \in I} A_i \to \prod_{i \in I} A_i/\mathcal{F}$ is a homomorphism of $\mathcal{L}$-structures.

(1.12) Examples. If I is finite, then all ultrafilters on I are principal. Note that if $\mathcal{F}$ is principal, say $\{j\} \in I$, then the ultraproduct $\prod_{i \in I} A_i/\mathcal{F}$ is naturally isomorphic to A_j.

The best known non-principal filter (on an infinite set I) is called the *Fréchet filter* and is the set of all subsets X of I such that $I \setminus X$ is finite. It is contained in all non-principal ultrafilters on I (Exercise).

Observe that if $A \subset I$ is infinite, then it intersects every cofinite subset of I; hence A and the Fréchet filter generate a (proper) filter, and A belongs to a non-principal ultrafilter.

(1.13) Łos' Theorem. Let I be infinite, $\mathcal{F}$ a filter on I, $(A_i)_{i\in I}$ a family of $\mathcal{L}$-structures, and $A = \prod_{i\in I} A_i/\mathcal{F}$.

(1) Let $\varphi(x)$ be a positive $\mathcal{L}$-formula, a a tuple in $\prod_{i\in I} A_i$. Then

$$A \models \varphi([a]_{\mathcal{F}}) \iff \{i \in I \mid A_i \models \varphi(a(i))\} \in \mathcal{F}.$$

(2) Assume now that $\mathcal{F}$ is an ultrafilter, and let $\varphi(x)$ be any formula, a a tuple in $\prod_{i\in I} A_i$. Then

$$A \models \varphi([a]_{\mathcal{F}}) \iff \{i \in I \mid A_i \models \varphi(a(i))\} \in \mathcal{F}.$$

This result is not difficult to prove, using induction on the complexity of the formulas. Note the restriction in (1) of $\varphi(x)$ being positive: the result definitely doesn't hold for formulas involving a negation, as can be shown by the following easy example. Let $a, b \in \prod_i A_i$. Then $[a]_{\mathcal{F}} = [b]_{\mathcal{F}} \iff \{i \in I \mid a(i) = b(i)\} \in \mathcal{F}$. But if $\mathcal{F}$ is not an ultrafilter, choose $A \subset I$ such that A and $I \setminus A$ are not in $\mathcal{F}$; then (assuming $|A_i| \geq 2$ for all i, choose b such that $\{i \in I \mid a(i) = b(i)\} = A$. Clearly $[a]_{\mathcal{F}} \neq [b]_{\mathcal{F}}$ but $\{i \in I \mid a(i) \neq b(i)\} \notin \mathcal{F}$.

(1.14) One immediate consequence of Łos' theorem is that if $\mathcal{F}$ is an ultrafilter on I, then the $\mathcal{L}$-structure A embeds elementarily into its ultrapower $A^I/\mathcal{F}$, via the map which to an element a associates $[\hat{a}]_{\mathcal{F}}$, where $\hat{a}$ is the function taking the value a on I.

(1.15) Theorem (Keisler-Shelah). Two $\mathcal{L}$-structures A and B are elementarily equivalent if and only if they have isomorphic ultrapowers.

(1.16) Let I be infinite, $\mathcal{F}$ a non-principal ultrafilter on I, and A an $\mathcal{L}$-structure. Assume that the language $\mathcal{L}$ is countable. Then $A^I/\mathcal{F}$ is *ω_1-saturated*, i.e., if $B \subset A^I/\mathcal{F}$ is countable, and $\Sigma(x)$ is a set of $\mathcal{L}(B)$-formulas which is finitely satisfiable in $A^I/\mathcal{F}$, then there is a tuple $b \in A^I/\mathcal{F}$ which satisfies all formulas of $\Sigma(x)$.

Proof. The set of $\mathcal{L}(B)$-formulas is countable, and therefore we may choose an enumeration $\psi_n(x)$ of $\Sigma(x)$. For each n, choose $b(n) \in A$ satisfying $\bigwedge_{i=1}^n \psi_i(x)$, and let $b = [(b(n))_n]_{\mathcal{F}}$.

(1.17) Exercise. Under the same hypotheses on $I, \mathcal{F}, \mathcal{L}$, show that the conclusion of (1.16) also holds if one assumes that the tuple x of variables has length ω.

(1.18) Exercise (harder). Let I be infinite, $\mathcal{F}$ an ultrafilter on I, and $(A_i)_{i\in I}$ a family of $\mathcal{L}$-structures of cardinality $\leq \aleph_0$. Show that $\prod_{i\in I} A_i/\mathcal{F}$ is either finite or of cardinality $\geq 2^{\aleph_0}$.

(1.19) Exercise. Let $\mathcal{F}$ be a filter on a set I, and let $J \in \mathcal{F}$, A_i, $i \in I$, a family of $\mathcal{L}$-structures. Let

$$\mathcal{G} = \{X \cap J \mid X \in \mathcal{F}\}.$$

(1) Show that $\mathcal{G}$ is a filter on J, and that if $\mathcal{F}$ is an ultrafilter, so is $\mathcal{G}$.

(2) Show that

$$\prod_{i\in I} A_i/\mathcal{F} \simeq \prod_{i\in J} A_i/\mathcal{G}.$$

(1.20) Exercise. Let I and J be sets, and $\mathcal{F}$, $\mathcal{G}$ be filters on I and J respectively.

(1) Let $\mathcal{D}$ be the set of subsets X of $I \times J$ satisfying:

$$\{i \in I \mid \{j \mid (i,j) \in X\} \in \mathcal{G}\} \in \mathcal{F}.$$

Verify that this set is a filter, and that $\mathcal{F}$ and $\mathcal{G}$ are ultrafilters, then so is $\mathcal{D}$.

(2) Verify that if $(A_{i,j})_{i\in I, j\in J}$ is a family of $\mathcal{L}$-structures, then the natural isomorphism

$$\prod_{(i,j)\in I\times J} A_{i,j} \;\to\; \prod_{i\in I}(\prod_{j\in J} A_{i,j})$$

induces an isomorphism

$$\prod_{(i,j)\in I\times J} A_{i,j}/\mathcal{D} \;\to\; \prod_{i\in I}(\prod_{j\in J} A_{i,j}/\mathcal{G})/\mathcal{F}.$$

(1.21) More on saturated models. Let κ be an infinite cardinal (greater that $|\mathcal{L}|$), M an $\mathcal{L}$-structure. Then M is *κ-saturated* iff for every n and subset A of M of cardinality $< \kappa$, if $\Sigma(x)$ (where x is an n-tuple of variables) is a set of $\mathcal{L}(A)$-formula which is finitely satisfiable in M, then there is an n-tuple a in M which satisfies all formulas in Σ. M is *saturated* if it is $|M|$-saturated. One can then show that an $\mathcal{L}$-structure M is κ-saturated if and only if the two following conditions hold:

(i) (κ-universality) Every model of $Th(M)$ of cardinality $< \kappa$ embeds elementarily into M.

(ii) (κ-homogeneity) If $\lambda < \kappa$ and $(a_i)_{i<\lambda}, (b_i)_{i<\lambda}$ are two sequences of elements of M such that $(M, a_i)_{i<\lambda} \equiv (M, b_i)_{i<\lambda}$, and if $a \in M$ then there is $b \in M$ such that $(M, a_i, a)_{i<\lambda} \equiv (M, b_i, b)_{i<\lambda}$.

(1.22) Theorem. Let M be an infinite $\mathcal{L}$-structure, and κ an infinite cardinal. Then M has an elementary extension which is κ-saturated.

Proof. Exercise. The proof is done in two steps.

(1) First show that M has an elementary extension M_1 such that for every n and subset A of M of cardinality $< \kappa$, if $\Sigma(x)$ (where x is an n-tuple of variables) is a set of $\mathcal{L}(A)$-formula which is finitely satisfiable in M, then there is an n-tuple a in M_1 which satisfies all formulas in Σ.

(2) Now, using step 1, build an increasing chain M_α, $\alpha < \kappa^+$ of elementary extensions of M such that

(i) $M_0 = M$;

(ii) if α is a limit ordinal then $M_\alpha = \bigcup_{\beta<\alpha} M_\beta$.

(iii) If $\alpha = \beta + 1$, then $M_\beta \prec M_\alpha$ and for every n and subset A of M_β of cardinality $< \kappa$, if $\Sigma(x)$ (where x is an n-tuple of variables) is a set of $\mathcal{L}(A)$-formula which is finitely satisfiable in M_β, then there is an n-tuple a in M_α which satisfies all formulas in Σ.

You then verify that $M_{\kappa^+} = \bigcup_{\alpha<\kappa^+} M_\alpha$ is κ-saturated and an elementary extension of M. Can you bound its cardinality? (Here we are using the fact that κ^+ is regular, and in particular that a subset of κ^+ of cardinality $< \kappa$ cannot be cofinal in κ^+.)

(1.23) Criterion for completeness of a theory. Let T be a theory in a language $\mathcal{L}$, and κ an infinite cardinal $\geq |\mathcal{L}|$. Then T is complete if and only if whenever M_1 and M_2 are two κ-saturated models of T, there exists a (non-empty) family $\mathcal{I}$ of partial isomorphisms $M_1 \to M_2$ which has the property of the back-and-forth.

(1.24) Criterion for elimination. Let T, $\mathcal{L}$, κ be as above, and Δ a set of $\mathcal{L}$-formulas closed under disjunction. The following conditions are equivalent:

(1) Every $\mathcal{L}$-formula is equivalent modulo T to a Boolean combination of formulas from Δ.

(2) For any two κ-saturated models M_1 and M_2 of T, consider the set $\mathcal{I}(M_1, M_2)$ of partial isomorphisms from M_1 to M_2, with domain of cardinality $< \kappa$ and which preserve formulas from Δ (i.e., if the tuple a is in the domain of $f \in \mathcal{I}$ and $\varphi(x) \in \Delta$, then $M_1 \models \varphi(a) \iff M_2 \models \varphi(f(a))$). Then either this set is empty, or it has the property of the back-and-forth.

Observe that this criterion works for theories which are not complete (e.g., the theory of algebraically closed fields). The condition that $\mathcal{I}$ be non-empty is equivalent to: M_1 and M_2 satisfy the same sentences in Δ.

(1.25) Definability of structures. Let M be an $\mathcal{L}$-structure, and N be an $\mathcal{L}'$-structure. We say that the $\mathcal{L}'$-structure N is *definable in* M iff there is some definable subset S of M^k for some k, and a bijection $F : N \to S$ such that:

– For each n-ary relation $R \in \mathcal{L}'$, $R^* = F(R)$ is definable.

– For each n-ary function $f \in \mathcal{L}'$, the image by F of the graph of f definable, and we will denote by f^* the function it defines on S^n.

Then the $\mathcal{L}'$-structures N and S are isomorphic.

Comments. By definable, I mean possibly with parameters from M. To be precise, one would say N is A-definable in M, whenever S, all $R^*, f^*, c^* = F(c)$ where c ranges over all constant symbols of $\mathcal{L}'$, are definable with parameters from A.

Exercise. Show that if N is definable in M (with parameters from $A \subset M$) and F is the isomorphism $N \to S$, then to every formula $\varphi(x) \in \mathcal{L}'$ we can associate an $\mathcal{L}(A)$-formula $\varphi^*(\bar{x})$ such that, for every tuple a in N^m, we have

$$N \models \varphi(a) \iff M \models \varphi^*(F(a)).$$

Exercise. Show that the field $\mathbb{C}$ is definable in the field $\mathbb{R}$.

(1.26) Interpretation of structures. Let M be an $\mathcal{L}$-structure, N an $\mathcal{L}'$-structure. Then N is *interpretable in* M iff there is a definable subset S of M^k for some k, a definable subset E of M^{2k} which defines an equivalence relation on S, and a bijection F between N and the set S/E of E-equivalences classes of S, such that:

– For each n-ary relation $R \in \mathcal{L}'$, the set of tuples $(a_1, \ldots, a_n) \in S^n$ such that $(F^{-1}(a_1/E), \ldots, F^{-1}(a_n/E)) \in R$, is definable in M.

– For each n-ary function $f \in \mathcal{L}'$, the set of tuples $(a_1, \ldots, a_n, b) \in S^{n+1}$ such that $f(F^{-1}(a_1/E), \ldots, F^{-1}(a_n/E)) = F^{-1}(b/E)$, is definable in M.

Exercise. Show that if N is interpretable in M (with parameters from $A \subset M$), and F is the isomorphism $N \to S/E$, then to every formula $\varphi(x) \in \mathcal{L}'$ we can associate an $\mathcal{L}(A)$-formula $\varphi^*(\bar{x})$ such that, for every tuple $a = (a_1, \dots, a_n)$ in N^n, and tuple $b = (b_1, \dots, b_n)$ where each $b_i \in F(a_i)$, we have

$$N \models \varphi(a) \iff M \models \varphi^*(b).$$

Exercise. Let G be a group (the language is the language of groups $\{\cdot, ^{-1}, 1\}$), and assume that H is a definable normal subgroup of G. Show that the quotient group G/H is interpretable in the group G.

These two exercises should convince you of the fact that "definability" and "interpretation" are just fancy names for very natural concepts that you have already used.

(1.27) Theory of a class, models of a theory.

We fix a first-order language $\mathcal{L}$. Given a class $\mathcal{K}$ of $\mathcal{L}$-structures, one can consider its theory, $\mathrm{Th}(\mathcal{K})$, consisting of the sentences which hold in all elements of $\mathcal{K}$. Dually, given an $\mathcal{L}$-theory T, we may consider the class $\mathcal{M}od(T)$ of all models of T.

Clearly one has $\mathcal{K} \subseteq \mathcal{M}od(\mathrm{Th}(\mathcal{K}))$, and the class $\mathcal{M}od(\mathrm{Th}(\mathcal{K}))$ is called the *elementary class generated by* $\mathcal{K}$. The inclusion is in general strict. If $\mathcal{K} = \mathcal{M}od(\mathrm{Th}(\mathcal{K}))$ then $\mathcal{K}$ is called an *elementary class.*

Proposition. Let $\mathcal{K}$ be a class of $\mathcal{L}$-structures. Then $\mathcal{K}$ is an elementary class if and only if it is closed under ultraproducts and elementary substructures.

Proof. The necessity of the condition is clear, since the class of models of a theory is closed under these operations. For the sufficiency, it suffices to show that every model of $\mathrm{Th}(\mathcal{K})$ embeds elementarily into an ultraproduct of members of $\mathcal{K}$. By (1.15) and the fact that an ultrapower of an ultraproduct of members of $\mathcal{K}$ is an ultraproduct of members of $\mathcal{K}$ (see exercise (1.20)), it suffices to show that every model of $\mathrm{Th}(\mathcal{K})$ is elementarily equivalent to an ultraproduct of members of $\mathcal{K}$.

So, let $M \models \mathrm{Th}(\mathcal{K})$, and let I be the set of all $\mathcal{L}$-sentences true in A. Observe that I is closed under finite conjunctions. If $\theta \in I$, then $\neg\theta \notin \mathrm{Th}(\mathcal{K})$, i.e., there is some $M_\theta \in \mathcal{K}$ such that $M_\theta \models \theta$, and we choose one such. For each $\psi \in I$, consider $X_\psi = \{\theta \in I \mid M_\theta \models \psi\}$. Then each X_ψ is non-empty, and furthermore, since I is closed under finite conjunctions, the set $\{X_\psi \mid \psi \in I\}$ has the finite intersection property. Let $\mathcal{U}$ be an ultrafilter on I which contains all X_ψ, $\psi \in I$. By Łos' theorem, $M \equiv \prod_{\theta \in I} M_\theta/\mathcal{U}$.

(1.28) Example. If $\mathcal{K} = \{M\}$, then $\mathcal{M}od(\mathrm{Th}(\mathcal{K}))$ is simply the set of $\mathcal{L}$-structures elementarily equivalent to M. By Keisler-Shelah (1.15), every member of $\mathcal{M}od(\mathrm{Th}(\mathcal{K}))$ embeds elementarily into an ultrapower of M.

(1.29) Example 2. Consider the class $\mathcal{K}$ of all finite fields, and T_f its theory. Then, every model of T_f which is finite is in $\mathcal{K}$, and every model of T_f which is infinite is elementarily equivalent to an ultraproduct of finite fields. Thus, to show that the pseudo-finite fields are exactly the infinite models of T_f, it suffices to show that if F is a pseudo-finite field, then it is elementarily equivalent to an ultraproduct of finite fields.

2. Finite fields - properties

(2.1) Basic properties

If p is a prime number, then $p\mathbb{Z}$ is a maximal ideal of $\mathbb{Z}$ and $\mathbb{Z}/p\mathbb{Z}$ is a field with p elements. This field is denoted by $\mathbb{F}_p$, and it is the *prime field of characteristic* p, i.e., it is contained in every field of characteristic p.

Let F be a finite field. Since $1 \in F$, it necessarily contains one of the fields $\mathbb{F}_p$, and is therefore a vector space over $\mathbb{F}_p$, whence of cardinality p^n for some $n \in \mathbb{N}$.

Let F be a field of characteristic p having $q = p^n$ elements, let K be an algebraically closed field containing F. Let us consider the multiplicative group $F^\times = F\backslash\{0\}$ of F. It has $q-1$ elements and hence every non-zero element of F satisfies the equation $X^{q-1} - 1 = 0$. Thus all elements of F satisfy $X^q - X = 0$. The derivative of this equation equals -1, all its roots are therefore simple roots, and we obtain

$$X^q - X = \prod_{a\in F}(X-a).$$

Conversely, let us consider the set $S \subset K$ of all solutions of $X^q - X = 0$. As above, its roots are all distinct. S is closed under multiplication, and $S \setminus \{0\}$ by multiplicative inverse. Because we are in characteristic p and q is a power of p, we have $(a+b)^p = a^p + b^p$ and $(a+b)^q = a^q + b^q$. This implies that S is closed under addition, and is therefore a subfield of K.

So, we have shown:

Theorem. Let F be a finite field. Then for some prime p and $q = p^n$, F has q elements. Its elements are exactly the roots of the equation $X^q - X = 0$.

(2.2) The multiplicative group of a finite field. Let $F = \mathbb{F}_q$ be a finite field. We will show that $F^\times$ is cyclic. It can be written as a direct sum of cyclic subgroups, and if it is not cyclic, then its exponent m is a proper divisor of $q-1$. But all roots of $X^{q-1} = 1$ are simple roots, whence all roots of $X^m = 1$ are simple as well (since $m \not\equiv 0$ modulo p). This implies that $q - 1 = m$.

(2.3) Perfect fields. Recall that a field F of characteristic $p > 0$ is *perfect* if every element of F has a p-th root. By convention, every field of characteristic 0 is perfect.

If $F = \mathbb{F}_{p^n}$ is finite, then the order of $F^\times$ is prime to p, which implies that every element is (multiplicatively) divisible by p, i.e., F is perfect.

An example of imperfect field is $\mathbb{F}_p(t)$, where t is transcendental over $\mathbb{F}_p$.

The *perfect hull* of a field F is the smallest perfect field containing F. If F is of characteristic $p > 0$ and is not perfect, it is obtained by adjoining to F all p^n-th roots of elements of F. It is then denoted by F^{1/p^∞}.

(2.4) The algebraic closure of $\mathbb{F}_p$.

Let m, n be positive integers, p a prime. Then

$$\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n} \iff m \text{ divides } n,$$

and in that case we have $[\mathbb{F}_{p^n} : \mathbb{F}_{p^m}] = n/m$.

Indeed, if $\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n}$ then $\mathbb{F}_{p^n}$ is a $\mathbb{F}_{p^m}$-vector space, hence m divides n and $[\mathbb{F}_{p^n} : \mathbb{F}_{p^m}] = n/m$. Conversely, if m divides n, then $p^m - 1$ divides $p^n - 1$, whence all roots of $X^{p^m-1} = 1$ are contained in $\mathbb{F}_{p^n}$, i.e., $\mathbb{F}_{p^m} \subseteq \mathbb{F}_{p^n}$.

It follows easily that for any $m, n \geq 1$, $\mathbb{F}_{p^m} \cap \mathbb{F}_{p^n} = \mathbb{F}_{p^d}$ and $\mathbb{F}_{p^m}\mathbb{F}_{p^n}$ (the composite field of $\mathbb{F}_{p^m}$ and $\mathbb{F}_{p^n}$) $= \mathbb{F}_{p^e}$ where d is the greatest common divisor of m and n, and e is the least common multiple of m and n.

Let α be algebraic over $\mathbb{F}_p$. Then $\mathbb{F}_p(\alpha)$ is a finite-dimensional $\mathbb{F}_p$-vector space, and is therefore also finite. This implies that the algebraic closure $\mathbb{F}_p^{alg}$ of $\mathbb{F}_p$ is $\bigcup_{n\in\mathbb{N}} \mathbb{F}_{p^n}$.

(2.5) The Frobenius map. Fix a prime p, and let F be a field of characteristic p. Using the binomial rule, one has, for every $a, b \in F$, $(a + b)^p = a^p + b^p$. Since $x^p = 0$ implies $x = 0$, this means that the map $x \mapsto x^p$ is an injective endomorphism of the field F. This map is called the *Frobenius map*. It is the identity on $\mathbb{F}_p$ (since every element of $\mathbb{F}_p$ satisfies $X^p - X = 0$), and defines an automorphism of each $\mathbb{F}_{p^n}$. Hence it defines an element φ of $\mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. Observe that if $d \in \mathbb{N}$, the elements of $\mathbb{F}_p^{alg}$ which are fixed by φ^d are precisely the elements of $\mathbb{F}_{p^d}$. Furthermore, one checks that the restriction $\varphi|_{\mathbb{F}_{p^d}}$ of φ to $\mathbb{F}_{p^d}$ has order exactly d: φ^ℓ being the identity on $\mathbb{F}_{p^d}$ means exactly that all elements of $\mathbb{F}_{p^d}$ satisfy $X^{p^\ell} = X$, and therefore that d divides ℓ.

This will allow us to describe $\mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p)$ in a nice fashion, see below. But first we need to do a little infinite Galois theory and introduce pro-finite groups.

3. Topological groups, profinite groups and infinite Galois theory

(3.1) Topological groups. A topological group is a group G endowed with a topology, such that group multiplication : $G \times G \to G$ and the inverse map $G \to G$ are continuous ($G \times G$ is endowed with the product topology). It then follows that multiplication (on the left or on the right) by an element g of the group defines a homeomorphism of G. Observe also that in order to define the topology, it suffices to give a basis of open sets containing 1: if $g \in G$, then translating the basis by g will give us a basis of open sets containing g.

Some properties. Let G be a topological group.

(1) Let H be a subgroup of G. Then the closure $\bar{H}$ of H in G is also a subgroup of G. If H is normal in G so is $\bar{H}$.

(2) Let H be an open subgroup of G. Then H is *clopen*, i.e., open and closed. If H is a closed subgroup of G of finite index, then H is clopen.

(3) Assume that G is compact, and let H be an open subgroup of G. Then $[G : H] < \infty$. [Note: in English, compact does not imply Hausdorff].

(4) Let $U \subset G$ and assume that U and V are dense open subsets of G. Then $U \cdot V = G$.

(5) Let $U, V \subset G$. Then $\bar{U} \cdot \bar{V} \subset \overline{U \cdot V}$.

(6) Let H be a subgroup of G and assume that H is dense in G, and that G is T_1 (i.e., every singleton is closed). Then $Z(G) \cap H = Z(H)$. [Recall that the center of a group G, denoted by $Z(G)$, is the set of elements commuting with all elements of G].

(7) Assume that G is T_1, and let $g \in G$. Then the centraliser of g in G, $C_G(g)$, is a closed subgroup of G. If H is abelian, then so is $\bar{H}$. Similarly for nilpotent and solvable.

Proof. (1) If X is a closed set containing H, then X^{-1} $(= \{g^{-1} \mid g \in X\})$ contains $H^{-1} = H$. Hence $\bar{H}$ is closed under the inverse map. Similarly, if $g \in H$, then $g\bar{H} = \overline{gH} \subseteq \bar{H}$, so that $H \cdot \bar{H} \subseteq \bar{H}$. Similarly, if $h \in \bar{H}$, then $Hh \subset \overline{Hh} = \bar{H}h$, and $Hh \subset \bar{H}$, so that $\bar{H} \cdot \bar{H} = \bar{H}$, i.e., $\bar{H}$ is a subgroup of G.

If H is normal in G, then for every $g \in G$ we have $gH = Hg$; hence also $g\bar{H} = \bar{H}g$.

(2) G is the disjoint union of the cosets gH, where g runs through a set of representatives of G/H. If H is open, so is every coset, and the union of the cosets gH for $g \notin H$ is open, and is the complement of H in G. Hence H is also closed. Similar proof for the second assertion, using the fact that there are only finitely many cosets.

(3) G is the union of the open sets gH, and therefore, by compactness, there are only finitely many of them, i.e., $[G : H] < \infty$.

(4) Let $g \in G$. Then V^{-1} and gV^{-1} are also dense open subsets of G. Hence $U \cap gV \neq \emptyset$, i.e., there are $u \in U$ and $v \in V$ such that $u = gv^{-1}$, and $g = uv$.

(5) As in the proof of (1), we first get that $U \cdot \bar{V} \subseteq \overline{U \cdot V}$, then that $\bar{U} \cdot \bar{V} \subseteq \overline{U \cdot V}$.

(6) Clearly $Z(G) \cap H \subseteq Z(H)$. Assume $h \in H$, $h \notin Z(H)$, and let $g \in G$ be such that $[g, h] = g^{-1}h^{-1}gh \neq 1$. Choose some open set U containing $[g, h]$ and not containing 1 (here we use the T_1-property of G). By continuity of the commutator map $(x, y) \mapsto [x, y]$, there are some open subset U_1 and U_2 of G, with $h \in U_1$ and $g \in U_2$, and such that the image of $U_1 \times U_2$ under the commutator map is contained in U. In particular, $[h, U_2] \subset U$; by density of H, $U_2 \cap H \neq \emptyset$, and this implies that $h \notin Z(H)$.

(7) $C_G(g) = \{h \in G \mid [g, h] = 1\}$. The set X of pairs (h_1, h_2) such that $[h_1, h_2] = 1$ is closed (because $\{1\}$ is closed), hence so is $X \cap \{g\} \times G$, and the latter is homeomorphic to $C_G(g)$.

Hence, if H is abelian, then for every $h \in H$, $C_G(h)$ is a closed subgroup of G which contains H, and therefore $\bar{H}$.

(3.2) Profinite groups and infinite Galois groups. There are several equivalent definitions of profinite groups. A *profinite* group is a topological group, which is compact, Hausdorff, and totally disconnected (i.e., has a basis of open sets which consists of clopen sets). Equivalently, it is an inverse limit of a projective system of finite groups with the corresponding topology. This might not be very informative, we will see how it works for Galois groups.

Let F be a field. We know that F^{alg} is the union of all finite normal algebraic extensions of F. We are interested in $Aut(F^{alg}/F)$, the group of automorphisms of F^{alg} which fix the elements of F, and I will (abusively) denote it by $\mathcal{G}al(F^{alg}/F)$. Note that if F is of characteristic $p > 0$, $\sigma \in Aut(F^{alg})$, and $a^p = b \in F^{alg}$ then necessarily $\sigma(a)^p = \sigma(b)$, so that knowing σ on b forces the value of σ on a. Thus every automorphism of the *separable closure* F^{sep} of F (the set of elements of F^{alg} which satisfy a separable equation over F) which is the identity on F will extend **uniquely** to an element of $Aut(F^{alg}/F)$.

Let $\mathcal{N}$ be the family of all finite Galois extensions of F, and if $L \subseteq M \in \mathcal{N}$, let $\pi_{LM} : \mathcal{G}al(M/F) \to \mathcal{G}al(L/F)$ be the restriction map. Then these maps are epimorphisms, and if $L \subseteq M \subseteq N \in \mathcal{N}$, we have $\pi_{MN}\pi_{LM} = \pi_{LN}$. We endow each (finite) group $\mathcal{G}al(L/F)$ with the discrete topology, and consider the group $\prod_{L \in \mathcal{N}} \mathcal{G}al(L/F)$ (pointwise multiplication), endowed with the product topology (recall that a basic set of the product topology on $\prod_{i \in I} X_i$ is a set $\prod_{i \in I} U_i$ where each U_i is open in X_i and all but finitely many of the U_i's equal X_i). This space therefore has a basis consisting of clopen sets.

Furthermore, since a finite discrete topological space is compact and Hausdorff, the group $\prod_{L\in\mathcal{N}} \mathcal{G}al(L/F)$ is also compact Hausdorff.

We then consider the subgroup

$$H = \{(\sigma_L)_L \in \prod_{L\in\mathcal{N}} \mathcal{G}al(L/F) \mid \text{ if } L \subseteq M \in \mathcal{N} \text{ then } \sigma_M|_L = \sigma_L\}.$$

One verifies that H is a closed subgroup of $\prod_{L\in\mathcal{N}} \mathcal{G}al(L/F)$, and that the map $\mathcal{G}al(F^{alg}/F) \to \prod_L \mathcal{G}al(L/F)$ defined by $\sigma \mapsto (\sigma|_L)_L$ embeds $\mathcal{G}al(F^{alg}/F)$ into H. Furthermore, one verifies that each element of H defines an automorphism of F^{alg} which fixes F, and therefore this map is in fact an isomorphism.

The embedding of $\mathcal{G}al(F^{alg}/F)$ into $\prod_L \mathcal{G}al(L/F)$ gives it the structure of a topological group, and it is compact, Hausdorff and totally disconnected. A basis for the topology is given by the translates of the subgroups $\mathcal{G}al(F^{alg}/L)$, $L \in \mathcal{N}$ (exercise). It follows that every open set containing the identity contains an open set of the form $\mathcal{G}al(F^{alg}/L)$ for some $L \in \mathcal{N}$.

The way of describing $\mathcal{G}al(F^{alg}/F)$ as an inverse limit of a projective system of finite groups is as follows: the inverse system is the set $\{\mathcal{G}al(L/F), \pi_{LM} \mid L \subseteq M \in \mathcal{N}\}$, and we write $\mathcal{G}al(F^{alg}/F) = \lim_{\leftarrow} \mathcal{G}al(L/F)$.

(3.3) Description of $\mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p)$. The family of all finite Galois extensions of $\mathbb{F}_p$ can be indexed by the integers $n \geq 2$, and we have seen that $\mathcal{G}al(\mathbb{F}_{p^n}/\mathbb{F}_p) \simeq \mathbb{Z}/n\mathbb{Z}$. The maps π_{mn} for m dividing n are then simply the epimorphisms $\mathbb{Z}/n\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$ given by reducing modulo m. Thus

$$\mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p) = \{(a_n)_{n\geq 2} \in \prod_{n\geq 2} \mathbb{Z}/n\mathbb{Z} \mid \text{ if } m|n \text{ then } a_m \equiv a_n \text{mod}(m)\}.$$

This group is also denoted by $\hat{\mathbb{Z}}$, or $\lim_{\leftarrow} \mathbb{Z}/n\mathbb{Z}$. One can show that it is isomorphic to $\prod_{\ell \text{ aprime}} \mathbb{Z}_\ell$, where $\mathbb{Z}_\ell$ denotes the set of ℓ-adic integers, i.e., $\mathbb{Z}_\ell = \lim_{\leftarrow} \mathbb{Z}/\ell^n\mathbb{Z}$.

Exercise. Verify that the subgroup $\langle\varphi\rangle$ of $\hat{\mathbb{Z}}$ generated by $(1)_n$ is dense, i.e., that it intersects every open subsets of $\hat{\mathbb{Z}}$. (Observe that $(1)_n$ corresponds to the Frobenius in the isomorphism $\mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p) \simeq \hat{\mathbb{Z}}$).

(3.4) The Galois correspondence. The Galois correspondence works almost as well as in the finite case, but we need to consider **closed** subgroups instead of arbitrary ones. I.e., separable algebraic extensions of F correspond to closed subgroups of $\mathcal{G}al(F^{alg}/F)$.

Thus if H_0 is a subgroup of $\mathcal{G}al(F^{alg}/F)$, E the subfield of F^{sep} of elements fixed by all elements of H_0, then $\mathcal{G}al(F^{alg}/E)$ will be the closure of H_0 in $\mathcal{G}al(F^{alg}/F)$ for the topology.

Note also that if E is a (maybe infinite) Galois extension of F, then one can define $\mathcal{G}al(E/F)$ in a similar way, by restricting one's attention to the family $\mathcal{N}_E$ of finite Galois extensions of F contained in E. Then the subgroup $\mathcal{G}al(F^{alg}/E)$ will be a normal (closed) subgroup of $\mathcal{G}al(F^{alg}/F)$, and the quotient $\mathcal{G}al(F^{alg}/F)/\mathcal{G}al(F^{alg}/E)$ is naturally isomorphic to $\mathcal{G}al(E/F)$, the restriction map $\mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(E/F)$ being continuous for the topology.

Thus we have a correspondence between quotients of $\mathcal{G}al(F^{alg}/F)$ by continuous homomorphisms and Galois extensions of F.

An aside: to check that a group epimorphism $G \to H$ between profinite groups is continuous, it suffices to check that the kernel of this map is a closed subgroup of G.

(3.5) Example. Consider again the Frobenius map $\varphi \in \mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p)$, and let H_0 be the subgroup it generates. Then the elements of $\mathbb{F}_p^{alg}$ which are fixed by φ are all in $\mathbb{F}_p$. However, H_0 being countable, it certainly doesn't equal $\hat{\mathbb{Z}}$ (which has cardinality $2^{\aleph_0}$, exercise).

(3.6) More on $\hat{\mathbb{Z}}$.

We will discuss a few properties of $\hat{\mathbb{Z}}$, and in particular show that if F is a perfect field, then $\mathcal{G}al(F^{alg}/F) \simeq \hat{\mathbb{Z}}$ if and only if F has exactly one (separable) extension of degree n for every $n > 1$. (Since our field F is perfect, we do not need to worry about separability).

First observe that this condition is certainly necessary: for every n, $\hat{\mathbb{Z}}$ has a unique open subgroup H_n of index n (namely, in the notation of (3.3), the set of sequences $(a_i)_{i\geq 2}$ such that a_i is divisible by n in $\mathbb{Z}/i\mathbb{Z}$. Equivalently, if $\bar{1}$ denotes the sequence which equals 1 everywhere, then this subgroup is the closure of the subgroup generated by $n\bar{1}$ in $\hat{\mathbb{Z}}$, and we will denote it by $n\hat{\mathbb{Z}}$. Hence, by Galois theory, this implies that F has exactly one separable extension of degree n: the subfield of F^{sep} consisting of elements fixed by H_n.

Let us now assume that F has a unique separable extension of degree n for each $n > 1$, and let us see what this implies. Let L be a finite Galois extension of F and G its Galois group over F. First note that our hypothesis implies that all subgroups of G are normal (since the conjugate of a subgroup of G has the same order as that subgroup). We will show that G is cyclic. This is done by induction on its order, call it n. If n is prime, there is nothing to prove, since $G \simeq \mathbb{Z}/n\mathbb{Z}$. Otherwise, let p be a prime dividing n, and $a \in G$ an element of order p. Then the subgroup $\langle a\rangle$ is a normal subgroup of G, and by induction hypothesis, the quotient $G/\langle a\rangle$ is cyclic. Hence there is $b \in G$ such that $G = \langle a, b\rangle$. If $a \in \langle b\rangle$, then $G = \langle b\rangle$ and we are done. Otherwise, we have $\langle a\rangle \cap \langle b\rangle = (1)$ (because a has order p), and the commutator $[a, b] = a^{-1}b^{-1}ab$ belongs to the intersection of the normal subgroups $\langle a\rangle$ and $\langle b\rangle$, i.e., a and b commute. Thus $G = \langle a, b\rangle = \langle a\rangle \times \langle b\rangle$. Because G has a unique subgroup of order p, we get that p does not divide the order of b, and therefore G is cyclic.

We have therefore shown the following: assume that the perfect field F has exactly one extension of degree n for each $n > 1$. Then all (continuous) quotients of the profinite $\mathcal{G}al(F^{alg}/F)$ are cyclic, and for every $n > 1$ it has a quotient $\simeq \mathbb{Z}/n\mathbb{Z}$.

(3.7) Proposition. Let G be a profinite group.

(1) If all continuous finite quotients of G are cyclic, then there is a (continuous) epimorphism $\hat{\mathbb{Z}} \to G$.

(2) If moreover for every $n > 1$, G has a continuous quotient isomorphic to $\mathbb{Z}/n\mathbb{Z}$, then any continuous epimorphism from $\hat{\mathbb{Z}}$ onto G is an isomorphism.

Proof. (1) Note first of all that G is abelian. Using the compactness of G, we can therefore find an element $a \in G$ such that for every open subgroup H of G, aH generates G/H. Indeed, for each open subgroup H of G, let S_H be the set of elements whose image in G/H is a generator of G/H. This set is closed (since it is a union of cosets of H), non-empty, and if $H \subseteq H'$ then $S_{H'} \subset S_H$. Hence, if H_1 and H_2 are open normal subgroups of G, then $S_{H_1H_2} \subseteq S_{H_1} \cap S_{H_2}$, and eveery finite intersection of S_H is non-empty. By compactness of G, there is some element a in the intersection of all S_H.

Consider the set I of numbers n such that G has an open subgroup H_n of index n. Then G embeds (continuously) into $\prod_{n\in I} G/H_n$ [by definition of a profinite group, the intersection of all open subgroups of a profinite group is 1; moreover, by compactness, every open subgroup is of finite index and is therefore also closed]. By compactness, we get that G is in fact isomorphic to

$$\{(b_n) \in \prod_{n\in I} G/H_n \mid \text{if } n|m \text{ then } b_n \equiv b_m \bmod n\}.$$

Thus the continuous epimorphism $\prod_{n>1} \mathbb{Z}/n\mathbb{Z} \to \prod_{n\in I} G/H_n$ restricts to an epimorphism $f : \hat{\mathbb{Z}} \to G$.

(2) Our hypothesis implies that the set I defined above coincides with the set of integers > 1, whence the map f is injective, and $G \simeq \hat{\mathbb{Z}}$.

So, we need to show that any (continuous) epimorphism $g : \hat{\mathbb{Z}} \to \hat{\mathbb{Z}}$ is injective. Let $N = Ker(g)$. Then N is a closed subgroup of $\hat{\mathbb{Z}}$. Since g is an epimorphism, for every $n > 1$ we have that $g^{-1}(n\hat{\mathbb{Z}})$ is the unique subgroup of $\hat{\mathbb{Z}}$ of index n, i.e., equals $n\hat{\mathbb{Z}}$. As $N \subset g^{-1}(n\hat{\mathbb{Z}})$ for every $n > 1$, we get that $N \subset \bigcap_{n>1} n\hat{\mathbb{Z}} = (1)$, i.e., g is injective.

Remarks. (1) The result of (2) extends easily to a profinite group G such that for every $n > 1$, G has only finitely many open subgroups of index $\leq n$.

(2) $\hat{\mathbb{Z}}$ is also called the free profinite group on one generator. It can be defined using universal properties. Similarly, one can define free profinite groups on $2, 3, \dots, \aleph_0, \dots$ generators.

(3.8) Definition. If G is a profinite group, then $X \subset G$ *generates topologically* G, or *is a set of topological generators*, iff the subgroup of G generated by X is dense in G. This means: if N is an open subgroup of G, then the image of X in G/N generates G/N. We saw already that $\bar{1}$ is a topological generator of $\hat{\mathbb{Z}}$.

Exercise.

(1) Find a necessary and sufficient condition on the sequence $(a_n)_{n\geq 2} \in \hat{\mathbb{Z}} \subset \prod_{n\geq 2} \mathbb{Z}/n\mathbb{Z}$ to be a topological generator of $\hat{\mathbb{Z}}$. Show that if $(a_n)_{n\geq 2}$ is a topological generator of $\hat{\mathbb{Z}}$, then there is an automorphism of $\hat{\mathbb{Z}}$ which sends $\bar{1}$ to $(a_n)_{n\geq 2}$.

(2) Show that if $f : \hat{\mathbb{Z}} \to G$ is a continuous epimorphism, where G is a profinite group, and σ is a topological generator of G, then $f^{-1}(\sigma)$ contains a topological generator of $\hat{\mathbb{Z}}$. [Hint: show that if $g : A \to B$ is an epimorphism between two finite cyclic groups, then every generator of B lifts to a generator of A. This result is true in more generality and is called Gaschütz lemma: let $g : A \to B$ be an epimorphism of finite groups, and assume that A is generated by d elements. If $X \subset B$ is of size

d and generates B, then there is some set Y generating A, of size d, and such that $g(Y) = X$.]

4. Zariski topology, etc.

In this section I recall some basic terminology coming from old-fashioned algebraic geometry, a good reference is Chapter III in Lang's book *Introduction to algebraic geometry.* We fix two algebraically closed field $K \subseteq \Omega$. Let $n \in \mathbb{N}$, $X = (X_1, \dots, X_n)$.

(4.1) The Zariski topology. The Zariski topology on K^n is given by the following basic closed sets

$$V(f_1(X), \dots, f_m(X)) = \{a \in K^n \mid f_1(a) = \cdots = f_m(a) = 0\},$$

where $f_1(X), \dots, f_m(X) \in K[X]$. Similarly, if $B \subset K[X]$, we define $V(B) = \{a \in K^n \mid f(a) = 0 \text{ for all } f(X) \in B\}$. Zariski closed sets are also called *algebraic sets.*

Dually, given a subset S of K^n, we define $I(S) = \{f(X) \in K[X] \mid f(a) = 0 \text{ for all } a \in S\}$. Note that $I(S)$ is a *radical ideal* (i.e., if it contains f^n then it contains f). Hilbert's Nullstellensatz tells us that if I is a proper radical ideal of $K[X]$, then $V(I) \neq \emptyset$, and that if $f_1(X), \dots, f_m(X) \in K[X]$, then $I(V(f_1(X), \dots, f_m(X)))$ is the radical of the ideal generated by $f_1(X), \dots, f_m(X)$.

Dually, if $S \subset K^n$, then $V(I(S))$ is the smallest Zariski-closed subset of K^n containing S, it is called the *Zariski closure of* S, and denoted by $\bar{S}$.

Thus there is a correspondence between Zariski closed subsets of K^n and radical ideals of $K[X]$. Since $K[X]$ is Noetherian, we get that every strictly descending chain of closed subsets of K^n is finite.

One can show that the Zariski topology on K^n is the topology induced by the Zariski topology on Ω^n.

(4.2) Irreducible components, (affine) varieties. Recall that a closed set is irreducible if it is not the union of two proper closed subsets. One verifies (fairly easily) that the Zariski closed set S is irreducible if and only if the ideal $I(S)$ is prime.

The Noetherianity of the topology then implies that every closed subset S of K^n can be expressed as a finite union of irreducible closed sets, and these closed sets correspond to the minimal prime ideals containing $I(S)$. They will be called the *irreducible components* of the closed set S.

By convention, an *(affine) variety* is an irreducible closed subset of K^n.

(4.3) Coordinate ring, field of definition of an algebraic set. Let $S \subset K^n$ be an algebraic set, $I = I(S)$. We then form the ring $K[S] = K[X]/I$, this is the *coordinate ring of* S (over K). If S is a variety, then $K[S]$ will be a domain, and the *dimension* of S is the transcendence degree of the field of fractions $K(S)$ of $K[S]$ over K. If S is an arbitrary algebraic set, then $dim(S)$ will be the sup of the dimensions of the irreducible components of S.

If F is a subfield of K such that I is generated by $I \cap F[X]$, then we say that S is *defined over* F, and the ring $F[S] = F[X]/I \cap F[X]$ is the coordinate ring of S over

F. There is a smallest subfield of K over which S is defined, and it is called the *field of definition of* S.

Warning. The algebraic set S may be *definable* over F without being defined over F, but this phenomenon only occurs in positive characteristic. Here is an example: if K is of characteristic p, and $t \in K$ is transcendental over $\mathbb{F}_p$, then the closed set $\{t\}$ is *definable* over $F = \mathbb{F}_p(t^p)$ (by the equation $X^p = t^p$), but it is *not defined over* F, since the ideal of $K[X]$ generated by $X^p - t^p$ does not contain the element $(X - t)$.

(4.4) Rational points. Let $S \subseteq K^n$ be an algebraic set, and F a subfield (or subring) of K. Then $S(F) = S \cap F^n$ is the set of *S-rational points of* S. We define $S(\Omega)$ to be the set of points in Ω^n which satisfy $f(X) = 0$ for all $f \in I(S)$. (We usually write S instead of $S(\Omega)$). If S is a variety, a point a in $S(\Omega)$ is *generic* over K iff $tr.deg(K(a)/K) = dim(S)$.

(4.5) Tensor products of F-algebras. Let F be a field. Recall that an F-algebra is simply a ring containing F, and it is therefore an F-vector space. Let A and B be two F-algebras. The F-algebra $A \otimes_F B$ is defined as follows:

Let $\mathcal{E}_A$ and $\mathcal{E}_B$ be bases of the F-vector spaces A and B respectively; we assume that they both contain 1. Then as a vector space, $A \otimes_F B$ is the F-vector space with basis $\mathcal{E} = \{a \otimes b \mid a \in \mathcal{E}_A, b \in \mathcal{E}_B\}$.

We now need to define multiplication on $A \otimes_F B$. By definition, every element of $A \otimes_F B$ will be a finite F-linear combination of elements of the basis. Let $c, c' \in \mathcal{E}_A$ and $d, d' \in \mathcal{E}_B$. Write $cc' = \sum_{a \in \mathcal{E}_A} \alpha_a a$, $dd' = \sum_{b \in \mathcal{E}_B} \beta_b b$, where the α_a and β_b are in F. Then define

$$(c \otimes d) \cdot (c' \otimes d') = \sum_{a \in \mathcal{E}_A, b \in \mathcal{E}_B} \alpha_a \beta_b a \otimes b.$$

Note that this sum makes sense since almost all α_a and β_b are 0.

One extends multiplication to $A \otimes_F B$ in the unique fashion so that multiplication is associative, distributive with the addition. Observe that A embeds into $A \otimes_F B$ by identifying $a \in \mathcal{E}_A$ with $a \otimes 1$, and similarly for B (identify $b \in \mathcal{E}_B$ with $1 \otimes b$). One denotes these copies of A and B by $A \otimes 1$ and $1 \otimes B$, and their elements by $a \otimes 1$, $1 \otimes b$ respectively. One then defines $a \otimes b$ to be the element $(a \otimes 1) \cdot (1 \otimes b)$. One important property is that if $a \in A$, $b \in B$ and $c \in F$, then $c(a \otimes b) = ca \otimes b = a \otimes cb$.

Exercise. Convince yourself of the fact that the isomorphism type of $A \otimes_F B$ does not depend on the choice of the bases $\mathcal{E}_A$ and $\mathcal{E}_B$.

(4.6) Regular extensions. Let $F \subseteq E$ be fields. Then E is a *regular extension of* F iff E and F^{alg} are linearly disjoint over F, i.e.: if elements of the F-vector space E are linearly independent, then they are also linearly independent in the F^{alg}-vector space EF^{alg}. One can show that this condition is symmetric. Equivalently, this happens iff the F-algebra $E \otimes_F F^{alg}$ is a domain. We will use this equivalent formulation.

If F is of characteristic 0 or is perfect, then E is a regular extension of F if and only if $E \cap F^{alg} = F$. [This criterion does not work for imperfect fields.]

Let V be an algebraic set defined over F, and assume it is F-irreducible, i.e., cannot be written as the union of two proper closed subsets defined over F, or equivalently, the ideal $I(V) \cap F[X]$ is prime. Consider the ring $F[V] = F[X]/I(V) \cap F[X]$, and its field of fractions $F(V)$. Then one has:

$$V \text{ is a variety} \iff F(V) \text{ is a regular extension of } F$$

(4.7) Exercise. Show that if F is a subfield of E and L, and L is a Galois extension of F, then E and L are linearly disjoint over F if and only $E \cap L = F$.

5. Bounds on ideals in polynomial rings, and applications

We will want to express in a field F the following property: if V is a variety defined over F, then V has a point with its coordinates in F. I.e., for all m, n, $X = (X_1, \dots, X_n)$, we want to find axioms which are satisfied by F if and only if

For all $f_1(X), \dots, f_m(X) \in F[X]$ which generate a prime ideal in $F^{alg}[X]$, there is an n-tuple a such that $f_1(a) = \dots = f_m(a) = 0$.

The difficulty is of course to express that $f_1(X), \dots, f_m(X)$ generate a prime ideal in $F^{alg}[X]$.

(5.1) Total degree. Let $i = (i_1, \dots, i_n) \in \mathbb{N}^n$. We define $X^i = X_1^{i_1} X_2^{i_2} \cdots X_n^{i_n}$, and its (total) degree to be $i_1 + \dots + i_n$. A polynomial in $F[X]$ is an F-linear combination of monomials of this form, and we define its total degree to be the sup of the degrees of the monomial occurring in it. For each integer d, the set of polynomials of total degree $\leq d$ forms a finite-dimensional vector space over F, and is denoted by $F[X]_{\leq d}$. Say its dimension is $N(d)$. Then the vector space $F[X]_{\leq d}$ is definable in F (using $N(d)$-tuples), as is the graph of the multiplication $F[X]_{\leq d} \times F[X]_{\leq d} \to F[X]_{\leq 2d}$. [For that one fixes an enumeration m_i of all monomials in X such that if the degree of m_i is strictly less than the degree of m_j then $i < j$; thus the polynomial $f(X) = \sum_{i=1}^{N(d)} a_i m_i \in F[X]_{\leq} d$ is encoded by the $N(d)$-tuple $(a_1, \dots, a_{N(d)})$.]

(5.2) Bounds. The following result is classical, it was proved first by Hermann [He], later by Seidenberg [S] (and there is a proof using ultraproducts in [DS]). I state it in its entirety, although I will only need items (1) and (4).

Theorem. Let n, d be positive integers, $X = (X_1, \dots, X_n)$.

(1) There is a constant $A = A(n, d)$ such that for every field F, polynomials $f_1, \dots, f_m, g \in F[X]_{\leq d}$, if g belongs to the ideal of $F[X]$ generated by $f_1, \dots, f_m$, then there are $h_1, \dots, h_m \in F[X]_{\leq A}$ such that $g = \sum_{i=1}^m f_i h_i$.

(2) There is a constant $B = B(n, d)$ such that for every field F, for every ideal I of $F[X]$ generated by elements of $F[X]_{\leq d}$ and for every $g \in F[X]_{\leq d}$, if $g^k \in I$ for some integer k, then $g^B \in I$.

(3) There is a constant $C = C(n, d)$ such that for every field F, ideals I and J generated by elements of $F[X]_{\leq d}$, the ideals $I \cap J$ and $J : I = \{f \in F[X] \mid fI \subseteq J\}$ are generated by elements of $F[X]_{\leq C}$.

(4) There is a constant $D = D(n, d)$ such that for every field F and ideal I of $F[X]$ generated by elements of $F[X]_{\leq d}$, if I is not prime, then there are $g, h \in F[X]_{\leq D}$ such that $gh \in I$ but $g, h \notin I$.

(5) There is a constant $E = E(n, d)$ such that for every field F and ideal I of $F[X]$ generated by elements of $F[X]_{\leq d}$, there are at most E minimal prime ideals containing I, and they are generated by elements of $F[X]_{\leq E}$.

Note that the number m will be bounded by the number $N(d)$ of monomials of degree $\leq d$.

(5.3) Corollary. Let $n \geq 1$. There is a formula $\varphi(y)$, y an $mN(d)$-tuple of variables, such that in every field F, for every $mN(d)$-tuple a in F, if $f_1, \dots, f_m$ is the m-tuple of elements of $F[X]_{\leq d}$ encoded by a, then

$$F \models \varphi(a) \iff \text{ the ideal of } F[X] \text{ generated by } f_1, \dots, f_m \text{ is prime.}$$

Proof. Let $D = D(n,d)$, $A = A(n,D)$. Then $f_1, \dots, f_m$ generate a prime ideal I in $F[X]$

if and only if for all $g, h \in F[X]_{\leq D}$, either $gh \notin I$ or one of g, h is in I,

if and only if for all $g, h \in F[X]_{\leq D}$, either for all $h_1, \dots, h_m \in F[X]_{\leq A}$, $gh \neq \sum_{i=1}^m h_i f_i$, or there are $h_1, \dots, h_m \in F[X]_{\leq A}$ such that $[g = \sum_{i=1}^m h_i f_i$ or $h = \sum_{i=1}^m h_i f_i]$.

This last statement is clearly an elementary property of the $mN(d)$-tuple a of coefficients of $f_1, \dots, f_m$.

(5.4) Corollary. Let $n \geq 1$. There is a **quantifier-free** formula $\psi(y)$, y an $mN(d)$-tuple of variables such that in every field F, for every $mN(d)$-tuple a in F, if $f_1, \dots, f_m$ is the m-tuple of elements of $F[X]_{\leq d}$ encoded by a, then

$$F \models \psi(a) \iff \text{ the ideal of } F^{alg}[X] \text{ generated by } f_1, \dots, f_m \text{ is prime.}$$

Proof. Take the formula $\varphi(y)$ given by (5.3). By quantifier-elimination of the theory of algebraically closed fields, there is a quantifier-free formula $\psi(y)$ such that in every algebraically closed field K, for every $mN(d)$-tuple a in K we have

$$K \models \varphi(a) \iff K \models \psi(a).$$

But if the tuple a is in the subfield F of K, we have

$$K \models \psi(a) \iff F \models \psi(a).$$

Thus $F \models \psi(a)$ if and only if the m-tuple $(f_1, \dots, f_m)$ of $F[X]_{\leq d}$ encoded by a generates a prime ideal in $F^{alg}[X]$.

(5.5) Pseudo-algebraically closed fields. A field F is pseudo-algebraically closed (abbreviated by PAC) if every variety V defined over F has an F-rational point. Using the above, we therefore get:

Corollary. There is a theory (in the language of rings) whose models are exactly the PAC fields.

(5.6) Side comments on PAC fields. Examples are: real closed fields, separably closed fields, and we will see also that ultraproducts of finite fields are PAC.

An algebraic extension of a PAC field is also PAC.

In order to show that a field F is PAC, it suffices to show: every (irreducible) curve C defined over F has infinitely many points with their coordinates in F.

6. Pseudo-finite fields and their elementary theory

In this section, we will introduce pseudo-finite fields, study their theory, and show that they are the infinite models of the theory of finite fields.

(6.1) Definition. A field F is *pseudo-finite* iff it satisfies the following three conditions.

P1 F is perfect (i.e., either $char(F) = 0$ or if $char(F) = p > 0$ then every element is a p-th power).

P2 For every $n > 1$, $\mathcal{G}al(F^{alg}/F) \simeq \hat{\mathbb{Z}}$.

P3 Every variety V defined over F has an F-rational point.

(6.2) Lemma. There is a set of sentences (in the language of rings) whose models are exactly the pseudo-finite fields, and we denote this set (or rather, the deductive closure of this set) by Psf.

Proof. In other words, we need to show that these properties can be expressed by sentences.

Property P1 is easy: you add a collection of sentences saying that if $p = 0$ then $\forall x \exists y\ y^p = x$ (p is the term $1 + 1 + \cdots + 1$, p-times).

Property P3 is an infinite collection of sentences $\theta_{n,d}$, one for each pair (n, d). The sentence $\theta_{n,d}$ will say: for every variety V defined by polynomials in $F[X_1, \dots, X_n]_{\leq d}$, there exists $x \in V$.

Property P2 is a little more complicated. We will show below that it is enough to find a collection of sentences, one for each $n > 1$, expressing that F has exactly one algebraic extension of degree n. Let us assume this result, and fix some $n > 1$.

Consider the formula $Irr(y)$ where $y = (y_1, \dots, y_n)$ which says: $\forall z_1, \dots, z_n$, for all $1 \leq d < n$ the polynomials $X^n + y_1 X^{n-1} + \cdots + y_n$ and $(X^d + z_1 X^{d-1} + \cdots + z_d)(X^{n-d} + z_{d+1} X^{d-1} + \cdots + z_n)$ are not equal.

It is clear that it is a first-order formula, and it exactly says that the polynomial $X^n + y_1 X^{n-1} + \cdots + y_n$ is irreducible (over F).

Thus the sentence $\exists y\ Irr(y)$ will guarantee that F has at least one algebraic extension of degree n.

In order to show that having exactly one extension of degree n is an elementary property of the field F, we will express the following fact in a first-order way: if $Irr(y)$ and $Irr(y')$ hold, then the algebraic extension obtained from F by adjoining a root of the polynomial $X^n + \sum_{i=1}^n y_i X^{n-i}$ contains a root of the polynomial $X^n + \sum_{i=1}^n y'_i X^{n-i}$.

To do that, we will show that if $Irr(a)$ holds in F, and α is a root of the polynomial $X^n + \sum_{i=1}^n a_i X^{n-i}$, then the structure $(F(\alpha), +, \times, 0, 1, P_F)$ is definable in F, where $+, \times, 0, 1$ are the usual addition, multiplication and constants on the field $F(\alpha)$, and P_F is a predicate for the subfield F. The idea is to view $F(\alpha)$ as an F-vector space with basis $\{1, \alpha, \dots, \alpha^{n-1}\}$, and to define in this vector space operations that make it isomorphic to the field $F(\alpha)$.

We let $S = F^n$, $+^*$ the usual addition on the vector space S, and $0^* = (0, 0, \dots, 0)$, $1^* = (1, 0, \dots, 0)$, P_F^* the set of elements $\{(b, 0, \dots, 0) \mid b \in F\}$. Clearly these sets, elements and relations are definable in F, with no parameters.

Multiplication by α induces a linear transformation of the vector space $F(\alpha)$, and its matrix is

$$M_\alpha = \begin{pmatrix} 0 & 0 & \cdots & 0 & -a_n \\ 1 & 0 & \cdots & 0 & -a_{n-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -a_1 \end{pmatrix}$$

since $\alpha^n = -\sum_{i=1}^n a_i \alpha^{n-i}$. Note that multiplication by α^i is also a linear transformation, and its matrix is simply M_α^i. So, we define $\times^*$ as follows

$$(x_1, \dots, x_n) \times^* (y_1, \dots, y_n) = (x_1 I_n + x_2 M_\alpha + \cdots x_n M_\alpha^{n-1}) \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}.$$

Here I_n denotes the identity $(n \times n)$-matrix. Observe that the definition of $\times^*$ uses the tuple $(a_1, \dots, a_n)$, but is totally uniform.

Hence, there is a formula $\theta^*(y)$ of the language of fields, such that if $Irr(b)$ holds for some n-tuple b in F and β is a root of $X^n + \sum_{i=1}^{n-1} b_i X^{n-i}$, then

$F \models \theta^*(b)$ if and only if $F(\beta)$ satisfies that if c is an n-tuple of elements of F which satisfies $Irr(x)$ in F, then $\exists z\ z^n + \sum_{i=1}^n c_i z^{n-i}$.

(6.3) Lemma. Let F be a perfect field, and assume that $\mathcal{G}al(F^{alg}/F) \simeq \hat{\mathbb{Z}}$. Let E be an elementary extension of F. Show that $\mathcal{G}al(E^{alg}/E) = \hat{\mathbb{Z}}$, that $F^{alg} \cap E = F$ and that $E^{alg} = EF^{alg}$.

Proof. Exercise.

(6.4) Theorem. Let $\mathcal{Q}$ be the set of all prime powers, and let $\mathcal{U}$ be a non-principal ultrafilter on $\mathcal{Q}$. Then the field $F^* = \prod_{q \in \mathcal{Q}} \mathbb{F}_q$ is a pseudo-finite field.

Proof. Since each $\mathbb{F}_q$ is prefect, F^* is perfect. Also, each $\mathbb{F}_q$ satisfies all axioms of P2, and therefore so does F^*. However, no finite field $\mathbb{F}_q$ will satisfy all axioms in P3. Indeed, fix q, and consider the algebraic set V defined by

$$Y \prod_{0 \leq i < j \leq q} (X_i - X_j) = 1.$$

One verifies easily that this set is a variety (it coordinate ring is $\mathbb{F}_q[X_0, \dots, X_q, \prod_{0 \leq i < j \leq q}(X_i - X_j)^{-1}]$, which is a domain, and whose field of fractions is purely transcendental over $\mathbb{F}_q$), but it cannot have an $\mathbb{F}_q$-rational point, since $\mathbb{F}_q$ only has q distinct elements. However, this is essentially the only problem. A result of Lang-Weil (see below (6.5)) will tell us that any axiom $\theta_{n,d}$ occurring in P3 will hold in all sufficiently large $\mathbb{F}_q$'s. Hence given n, d, the set of $q \in \mathcal{Q}$ such that $\mathbb{F}_q$ satisfies $\theta_{n,d}$ is cofinite, and therefore must be in $\mathcal{U}$. By Łos' theorem, F^* satisfies all $\theta_{n,d}$.

(6.5) The theorem of Lang-Weil. For every positive integers n, d, there is positive constant C $(= C(n,d))$ such that for every finite field $\mathbb{F}_q$ and variety V defined by polynomials in $\mathbb{F}_q[X_1, \dots, X_n]_{\leq d}$,

$$\big||V(\mathbb{F}_q)| - q^{\dim(V)}\big| \leq Cq^{\dim(V)-1/2}.$$

[Recall that $V(\mathbb{F}_q)$ is the set of points of $V \cap \mathbb{F}_q^n$, and $\dim(V)$ is the dimension of V, i.e., $tr.deg(\mathbb{F}_q(V)/\mathbb{F}_q)$.]

In particular, if $q > C^2$, then any variety V as above will have a rational point in $\mathbb{F}_q$. Indeed, we get

$$0 < -Cq^{\dim(V)-1/2} + q^{\dim(V)} \leq |V(\mathbb{F}_q)|.$$

The constant C can in fact be effectively computed.

(6.6) Recall that we want to show that the pseudo-finite fields are exactly the models of the theory T_f consisting of all sentences true in all finite fields. For that, we will study the completions of Psf, and show that any pseudo-finite field is elementarily equivalent to an ultraproduct of finite fields. In fact a stronger statement is true: a pseudo-finite field of characteristic 0 is elementarily equivalent to an ultraproduct of fields $\mathbb{F}_p$, p ranging over all primes. But for that, we need first to find the elementary invariants of pseudo-finite fields. The results given below will also allow us to show that the completions of Psf are model complete if one enlarges the language in a certain way.

(6.7) Lemma. Let F be a perfect PAC field.

(1) Let a be a tuple in some field containing F, and assume that the field $F(a)$ is a regular extension of F. Then there is an F-morphism $F[a] \to F$.

(2) Assume that F is $\aleph_1$-saturated, that A is a countable subset of some field containing F and that $F(A)$ is a regular extension of F. Then there is an F-morphism $F[A] \to F$.

Proof. (1) Consider the ideal $I(a/F) = \{f(X) \in F[X] \mid f(a) = 0\}$, where X is a tuple of the same length as a. Then $F[a] \simeq F[X]/I(a/F)$, and our regularity assumption tells us that the algebraic set V defined by $I(a/F)$ is a variety. Hence there is $b \in V(F)$. The map which sends a to b extends to an F-morphism $F[a] \to F$: by definition of V, b satisfies all polynomial equations over F which are satisfied by a.

(2) Finding this map is the same thing as finding some subset B of F which satisfies all equations satisfied by the (infinite) tuple A. Use (1) and $\aleph_1$-saturation.

(6.8) The embedding lemma. Let K, E, F be perfect fields, with $K \subset E, F$, F and E are regular extensions of K, E is countable and F is an $\aleph_1$-saturated pseudo-finite field. Assume that we have a continuous isomorphism $\Phi : \mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(E^{alg}/E)$ such that for every $a \in K^{alg}$ and $\sigma \in \mathcal{G}al(F^{alg}/F)$, $\Phi(\sigma)(a) = \sigma(a)$.

Then there is a K^{alg}-embedding $\varphi : E^{alg} \to F^{alg}$ such that F is a regular extension of $\varphi(E)$, and for every $a \in E^{alg}$ and $\sigma \in \mathcal{G}al(F^{alg}/F)$,

$$\varphi(\Phi(\sigma)(a)) = \sigma(\varphi(a)).$$

Comments. (1) Since all fields are perfect, the regularity of F over K is equivalent to $K^{alg} \cap F = F$.

(2) What is important in this result, is the fact that there is a K-embedding φ of E into F such that F is a regular extension of $\varphi(E)$.

(3) Because E and F are regular extensions of K, the restriction maps $\mathcal{G}al(E^{alg}/E) \to \mathcal{G}al(K^{alg}/K)$ and $\mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(K^{alg}/K)$ are onto. Thus the first condition on Φ simply tells us that the following diagram commutes:

$$\begin{tikzcd}
\mathcal{G}al(F^{alg}/F) \arrow[rr, "\Phi"] \arrow[dr] & & \mathcal{G}al(E^{alg}/E) \arrow[dl] \\
& \mathcal{G}al(K^{alg}/K) &
\end{tikzcd}$$

The second condition on Φ simply tells us that it coincides with the map induced by φ from $\mathcal{G}al(F^{alg}/F)$ to $\mathcal{G}al(E^{alg}/E)$, $\sigma \mapsto \varphi^{-1}\sigma\varphi$.

Proof of the embedding lemma. We work inside a large algebraically closed field Ω, and we may actually "move" E and assume that it is linearly disjoint from F over K, so that E^{alg} and F^{alg} are linearly disjoint over K^{alg}. Indeed, replacing E^{alg} by a K^{alg}-isomorphic copy $\psi(E)$, it suffices to solve the problem for $\psi(E)$, replacing Φ by the map $\Psi^{-1}\Phi$, where $\Psi : \mathcal{G}al(\psi(E^{alg}/\psi(E)) \to \mathcal{G}al(E^{alg}/E)$ is the map $\sigma \mapsto \psi^{-1}\sigma\psi$. Getting the embedding φ of $\psi(E^{alg})$ into F^{alg} will then give us the desired embedding of E^{alg} into F: simply take $\varphi\psi$.

Because E^{alg} and F^{alg} are linearly disjoint over K^{alg}, and because E and F are regular extensions of K, we get that E and F are linearly disjoint over K. We also get that

$$\mathcal{G}al(E^{alg}F^{alg}/EF) \simeq \mathcal{G}al(E^{alg}/E) \times_{\mathcal{G}al(K^{alg}/K)} \mathcal{G}al(F^{alg}/F),$$

where the right-hand-side group denotes the set of elements (σ_1, σ_2) of $\mathcal{G}al(E^{alg}/E) \times \mathcal{G}al(F^{alg}/F)$ such that $\sigma_1|_{K^{alg}} = \sigma_2|_{K^{alg}}$. Indeed, one shows (easily) that the Galois extensions $E^{alg}F$ and EF^{alg} are linearly disjoint over EFK^{alg}, and the result follows by Galois theory.

An alternate way of looking at it is to notice that the field $E^{alg}F^{alg}$ is simply the field of fractions of $E^{alg} \otimes_{K^{alg}} F^{alg}$. If $\sigma_1 \in \mathcal{G}al(E^{alg}/E)$ and $\sigma_2 \in \mathcal{G}al(F^{alg}/F)$ have the same restriction to K^{alg}, one then defines $(\sigma_1, \sigma_2)(a \otimes b) = \sigma_1(a) \otimes \sigma_2(b)$ for $a \in E^{alg}$ and $b \in F^{alg}$, and then checks that this extends to an automorphism of $E^{alg} \otimes_{K^{alg}} F^{alg}$.

Under this identification, because of our condition on Φ, the graph of Φ^{-1} can be naturally viewed as a closed subgroup H of $\mathcal{G}al(E^{alg}F^{alg}/EF)$. Moreover, if $\sigma_0 \in \mathcal{G}al(F^{alg}/F)$ is a topological generator of $\mathcal{G}al(F^{alg}/F)$, then the element $\tau_0 = (\Phi(\sigma_0), \sigma_0)$ is a topological generator of H (and of course $\Phi(\sigma_0)$ is a topological generator of $\mathcal{G}al(E^{alg}/E)$). Observe that the map $\tau_0 \mapsto \sigma_0$ defines an isomorphism between H and $\mathcal{G}al(F^{alg}/F)$.

Extend τ_0 to an automorphism τ of $(EF)^{alg}$, and let M be the subfield of $(EF)^{alg}$ of elements fixed by τ. Then, because $\tau|_{E^{alg}} = \Phi(\sigma_0)$ and $\tau|_{F^{alg}} = \sigma_0$, we have

$$M \cap E^{alg} = E, \quad M \cap F^{alg} = F,$$

i.e., M is a regular extension of E and of F. Observe also that the restriction maps $H = \mathcal{G}al((EF)^{alg}/M) \to \mathcal{G}al(E^{alg}/E)$ and $H \to \mathcal{G}al(F^{alg}/F)$ being isomorphisms (because they are epimorphisms, and all groups are isomorphic to $\hat{\mathbb{Z}}$), we obtain that

$$(EF)^{alg} = ME^{alg} = MF^{alg}.$$

Because M contains F, the field MF^{alg} is in fact the ring generated by M and F^{alg}. Because E is countable, there is some countable subfield M_0 of M containing E and such that $F^{alg}[M_0]$ contains E^{alg}. Then, $F(M_0)$ is contained in M and is therefore a regular extension of F. By Lemma (6.7), there is an F-morphism $F[M_0] \to F$, and because F^{alg} and FM_0 are linearly disjoint over F, this morphism extends to an F^{alg}-morphism $\varphi : F^{alg}[M_0] \to F^{alg}$. This is our desired φ.

Observe first that since M_0 contains E, we have $\varphi(E) \subset F$. We need to show that Φ coincides with the map induced by φ on the Galois groups. So, let $a \in E^{alg}$, and write it $a = \sum_i m_i b_i$ for some $m_i \in M_0$ and $b_i \in F^{alg}$. Then

$$\Phi(\sigma_0)(a) = \tau(a) = \sum_i \tau(m_i)\tau(b_i) = \sum_i m_i \sigma_0(b_i)$$

because τ is the identity on M and $\tau|_{F^{alg}} = \sigma_0$. Thus

$$\varphi(\Phi(\sigma_0)(a)) = \sum_i \varphi(m_i)\sigma_0(b_i)$$

because φ is the identity on F^{alg}. On the other hand we have

$$\sigma_0(\varphi(a)) = \sigma_0(\sum_i \varphi(m_i)b_i) = \sum_i \varphi(m_i)\sigma_0(b_i)$$

because the $\varphi(m_i)$ are in F. This gives us the desired equality. Because σ_0 is a topological generator of $\mathcal{G}al(F^{alg}/F)$, if $\sigma \in \mathcal{G}al(F^{alg}/F)$ and $a \in E^{alg}$, then $\sigma(\varphi(a)) = \sigma_0^j(\varphi(a))$ for some integer j, and thus we get the desired result.

To show that F is a regular extension of $\varphi(E)$, it suffices to show that $\varphi(E)^{alg} \cap F = \varphi(E)$. Let $a \in E^{alg}$. Then

$$\begin{aligned} \varphi(a) \in F &\iff \sigma_0(\varphi(a)) = \varphi(a) \\ &\iff \varphi(\Phi(\sigma_0)(a)) = \varphi(a) \\ &\iff \Phi(\sigma_0)(a) = a \\ &\iff a \in E. \end{aligned}$$

(6.9) Comments. This is a rather weak version of the embedding lemma. The full version of this lemma can be obtained by changing some of the hypotheses and the conclusion. Here are the changes to be made:

– F does not need to be pseudo-finite, only PAC (but it is still $\aleph_1$-saturated and perfect).

– Φ is a (continuous) epimorphism from $\mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(L/E)$, where L is a Galois extension of E.

The conclusion on the regularity of the extension F of $\varphi(E)$ is then replaced by: $\varphi(L) \cap F = \varphi(E)$.

The strategy of the proof is essentially the same. Again, one has the graph H of Φ sitting as a closed subgroup of $\mathcal{G}al(F^{alg}/F) \times_{\mathcal{G}al(K^{alg}/K)} \mathcal{G}al(E^{alg}/E)$, and one uses the fact that $H \simeq \mathcal{G}al(F^{alg}/F)$ is projective to "lift" H to a closed subgroup of $\mathcal{G}al((EF)^{alg}/EF)$. See the book *Field arithmetic* by Fried and Jarden [FJ].

There is also a version of this lemma for imperfect PAC fields.

(6.10) Corollary. Let K_1, K_2, E, F be perfect fields, with E a regular extension of K_1, F a regular extension of K_2, E countable and F pseudo-finite $\aleph_1$-saturated. Assume that we have an isomorphism $\Phi : \mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(E^{alg}/E)$ and an isomorphism $\varphi_0 : K_1^{alg} \to K_2^{alg}$ satisfying for all $a \in K_1^{alg}$ and $\sigma \in \mathcal{G}al(F^{alg}/F)$,

$$\varphi_0(\Phi(\sigma)(a)) = \sigma(\varphi_0(a)).$$

Then φ_0 extends to an embedding $\varphi : E \to F$ such that F is a regular extension of $\varphi(E)$, and for all $a \in E^{alg}$ and $\sigma \in \mathcal{G}al(F^{alg}/F)$,

$$\varphi(\Phi(\sigma)(a)) = \sigma(\varphi(a)).$$

Proof. If $K_1^{alg} = K_2^{alg}$ and φ_0 is the identity, then this is simply (6.8). We will show how to reduce to the situation of that lemma. Extend φ_0 to an isomorphism ψ defined on E^{alg}. Now apply the previous lemma to the situation $(\psi(E^{alg}), F^{alg}, \Psi^{-1}\Phi)$ where $\Psi : \mathcal{G}al(\psi(E)^{alg}/\psi(E)) \to \mathcal{G}al(E^{alg}/E)$ is the map $\sigma \mapsto \psi^{-1}\sigma\psi$. One needs of course to verify that they satisfy the hypotheses of Lemma (6.8), but this is easily done, since ψ extends φ_0. One then finds a K_2^{alg}-embedding ρ of $\psi(E)^{alg}$ into F^{alg} satisfying the conclusions of (6.8). Take $\varphi = \rho\psi$, and verify that it works.

(6.11) Proposition. Let E and F be pseudo-finite fields, K a common subfield which is perfect and such that E and F are regular extensions of K. Then

$$E \equiv_K F.$$

Proof. By $\equiv_K$ we mean that E and F are elementarily equivalent in the language $\mathcal{L}(K)$ of rings to which one has added constants for the elements of K. First of all we may assume that K is countable (since K can be expressed as a union of perfect subfields which are countable and relatively algebraically closed in E and F). Then, we may assume that E and F are $\aleph_1$-saturated: indeed, let E^* and F^* be elementary extensions of E and F respectively, and which are $\aleph_1$-saturated (see (1.22) for the existence). Then

$$E \equiv_K F \iff E^* \equiv_K F^*.$$

Step 1. Let $E_0 \prec E$ be countable and containing K. Then E is a regular extension of E_0, and E_0 is a regular extension of K. We want to apply the embedding Lemma (6.8), but for that we need to find the isomorphism Φ. Let σ_1 be a topological generator of $\mathcal{G}al(F^{alg}/F)$, and consider its restriction σ_0 to K^{alg}. Then σ_0 extends to a topological generator σ_2 of $\mathcal{G}al(E^{alg}/E)$ (see (3.8)) and sending σ_1 to σ_2 will define a group isomorphism Φ between $\mathcal{G}al(F^{alg}/F)$ and $\mathcal{G}al(E^{alg}/E)$ (which are both isomorphic to $\hat{\mathbb{Z}}$). By definition, σ_1 and σ_2 have the same action on K^{alg}.

By Lemma (6.3), the restriction map $\mathcal{G}al(E^{alg}/E) \to \mathcal{G}al(E_0^{alg}/E_0)$ is an isomorphism, and we let Φ_0 be the composition of Φ with this isomorphism.

By (6.10), there is a K^{alg}-embedding $\varphi_0 : E_0^{alg} \to F^{alg}$, such that F is a regular extension of $\varphi_0(E_0)$, and the map $\mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(E_0^{alg}/E_0)$ dual to φ_0 coincides with Φ_0.

Step 2. Let $F_0 \prec F$ be countable and contain $\varphi_0(E_0)$. We want to find an embedding of F_0 into E which extends φ_0^{-1} and has certain properties. Obviously we will use (6.10). Because $F_0 \prec F$, the restriction map $\mathcal{G}al(F^{alg}/F) \to \mathcal{G}al(F_0^{alg}/F_0)$ is an isomorphism, and this allows us (by composing it with Φ^{-1}) to get an isomorphism $\Psi_0 : \mathcal{G}al(E^{alg}/E) \to \mathcal{G}al(F_0^{alg}/F_0)$ such that the tuple $(E, F_0, \varphi_0^{-1}, \Psi_0)$ satisfies the hypotheses of (6.10).

Hence there is some $\psi_0 : F_0^{alg} \to E^{alg}$, which extends φ_0^{-1}, and is such that E is a regular extension of $\psi_0(F_0)$, and the map $\mathcal{G}al(E^{alg}/E) \to \mathcal{G}al(F_0^{alg}/F)$ dual to ψ_0 coincides with Ψ_0.

Step 3. We now use the same technique as in step 2 to build inductively sequences of partial isomorphisms φ_i and ψ_i such that:

(i) The domain of φ_i is an elementary substructure E_i of E and F is a regular extension of the image of φ_i; the domain of ψ_i is an elementary substructure of F, and E is a regular extension of the image of ψ_i.

(ii) Each ψ_i extends φ_i^{-1}, and each φ_{i+1} extends ψ_i^{-1}.

Consider now $E_\omega = \bigcup_{i\in\mathbb{N}} E_i$ and $F_\omega = \bigcup_{i\in\mathbb{N}} F_i$. Then $E_\omega \prec E$ and $F_\omega \prec E$. Moreover, $E_\omega \simeq_K F_\omega$: take $\bigcup_i \varphi_i$ $(= \bigcup_i \psi_i^{-1})$. Since E_ω and F_ω are K-isomorphic, we have $E_\omega \equiv_K F_\omega$. From $E_\omega \prec E$ and $F_\omega \prec F$, we deduce $E \equiv_K F$.

(6.12) Corollary. Let $E \subset F$ be pseudo-finite fields. Then $E \prec F \iff E^{alg} \cap F = E$.

Proof. Immediate by (6.11).

(6.13) Theorem. Let E and F be pseudo-finite fields, and K a common subfield. Then

$$E \equiv_K F \iff E \cap K^{alg} \simeq_K F \cap K^{alg}.$$

Proof. First note that since E and F are perfect, $E \cap K^{sep} \simeq_K F \cap K^{sep} \iff E \cap K^{alg} \simeq_K F \cap K^{alg}$.

Assume that $E \equiv_K F$, and let $\mathcal{N}$ be the set of all finite Galois extensions of K. For each $L \in \mathcal{N}$, we consider the set $S_L = \{\sigma \in \mathcal{G}al(K^{alg}/K) \mid \sigma(E \cap L) = F \cap L\}$. This set is a union of cosets of the open subgroup $\mathcal{G}al(K^{alg}/L)$, and is therefore open and closed. If $L \subseteq M \in \mathcal{N}$, then $S_M \subseteq S_L$; hence if $L, M \in \mathcal{N}$, then $S_L \cap S_M \supseteq S_{LM}$. It therefore suffices to show that each S_L is non-empty: if we do that, we will have shown that any finite intersection of S_L's is non-empty. By compactness of $\mathcal{G}al(K^{alg}/K)$, it will follow that $\bigcap_{L\in\mathcal{N}} S_L$ is non-empty, i.e., that $E \cap K^{alg} \simeq_K F \cap K^{alg}$. Fix $L \in \mathcal{N}$, and $\alpha \in L$ such that $E \cap L = K(\alpha)$ (this is where we use the separability of L over K). Let $f(X) \in K[X]$ be the minimal polynomial of α over K. Then $E \models \exists x\ f(x) = 0$, and therefore $F \models \exists x\ f(x) = 0$. Thus there is a K-embedding of $E \cap L$ into $F \cap L$, by some $\sigma \in \mathcal{G}al(L/K)$, and hence $[F \cap L : K] \geq [E \cap L : K]$. Reasoning similarly with $F \cap L$, we obtain that $[E \cap L : K] \geq [F \cap L : K]$ and therefore $[E \cap L : K] = [F \cap L : K]$, i.e., $\sigma(E \cap L) = F \cap L$, and $S_L \neq \emptyset$.

This shows one direction. For the other, fix some K-isomorphism $\varphi_0 : E \cap K^{alg} \to F \cap K^{alg}$, and extend it to a K-isomorphism φ defined on E. Then $E \equiv_K \varphi(E)$ because φ is a K-isomorphism, and $\varphi(E) \equiv_{F\cap K^{alg}} F$ by Proposition (6.11).

Remark. The left-to-right direction holds for arbitrary fields. Let E and F be fields, K a common subfield. Then $E \equiv_K F$ implies $E \cap K^{alg} \simeq_K F \cap K^{alg}$. If E and F are perfect, then the above proof works with no change, and also if K is perfect. If K is not perfect, then in particular K is infinite and one sets $\mathcal{N}$ to be the set of all finite normal extensions of K, and defines S_L as before. If $L_1, \dots, L_n$ are the conjugates of $L \cap F$ over K (under the action of $\mathrm{Aut}(L/K)$), one then shows (as in the separable case) that for every $a \in E \cap L$, there is some i such that $a \in L_i$. Then the result follows from

Lemma. Let K be an infinite field, and $V, V_1, \dots, V_n$ subspaces of the K-vector space W. If $V \subseteq V_1 \cup \cdots \cup V_n$, then $V \subseteq V_i$ for some i.

The proof of this lemma is by induction on n. If $n = 1$ there is nothing to prove. If $V_1 \cap V \subseteq V_2 \cup \cdots \cup V_n$, then actually $V \subseteq V_2 \cup \cdots \cup V_n$ and we are done by induction hypothesis. So, assume that $V_1 \cap V \not\subset V_2 \cup \cdots \cup V_n$, take $a \in V \cap V_1$, $a \notin V_2 \cup \cdots \cup V_n$ and $b \in V$. For every $k \in K$, we have $a + bk \in V$, and therefore there is some i such that $a + bk \in V_i$. Because K is infinite, there are $k \neq \ell \in K$ and i such that $a + bk$ and $a + b\ell$ belong to the same V_i. Then $(a + bk) - (a + b\ell) = b(k - \ell) \in V_i$, which implies that $b \in V_i$, $a \in V_i$ and therefore $i = 1$, $V \subseteq V_1$.

(6.14) Theorem. The completions of Psf are obtained by specifying the characteristic, and adding a collection of sentences of the form $\exists t\ f(t) = 0$ and $\forall t\ f(t) \neq 0$, where $f(T)$ ranges over all polynomials in the single variable T and with coefficients in $\mathbb{Z}$.

Proof. This is a direct consequence of (6.13), or rather of its proof: the isomorphism type of the relative algebraic closure of the prime field inside a field can be described by a conjunction of sentences of that form.

(6.15) Theorem. Let $\varphi(x)$ be a formula. Then there is a formula $\psi(x)$ which is a Boolean combination of formulas of the form $\exists t\ f(x,t) = 0$, where $f(X,T)$ is a polynomial over $\mathbb{Z}$, T a single variable (and X a tuple of variables of the same length as x).

Proof. This follows from Corollary (1.4) and (the proof of) Theorem (6.13). Indeed, let Δ be the set of Boolean combinations of formulas of the form $\exists t\ f(x,t) = 0$, where $f(X,T)$ is a polynomial over $\mathbb{Z}$, T a single variable. This set is clearly closed under disjunctions. By Corollary (1.4), we need to show that if E is a pseudo-finite field, a a tuple in E satisfying $\varphi(x)$, and if F is another pseudo-finite field, and b a tuple in F which satisfies (in F) all formulas in Δ which are satisfied by a in E, then $F \models \varphi(b)$.

So, let E, a and F, b be as above. Note that the formula expressing that $p = 0$ is in Δ (nothing prevents the polynomial $f(X,T)$ to be constant). Hence our hypothesis implies that the prime subfields k of E and F are isomorphic. Similarly, the tuples a and b satisfy exactly the same equations over $\mathbb{Z}$ (hence over $\mathbb{Z}/p\mathbb{Z}$ if the characteristic is p), and this implies that there is an isomorphism ψ between the subrings $k[a]$ of E and $k[b]$ of F which sends a to b. This isomorphism extends to an isomorphism ψ between the subfields $k(a)$ and $k(b)$. Our assumption on a and b satisfying the same formulas from Δ says that if $f(a,T) \in k[a][T]$, then

$$E \models \exists t\ f(a,t) = 0 \iff F \models \exists t\ f(b,t) = 0.$$

It follows that the isomorphism ψ extends to an isomorphism between $k(a)^{alg} \cap E$ and $k(b)^{alg} \cap F$. The fact that $F \models \varphi(b)$ now follows from the more general version of Theorem (6.13):

Theorem (6.13)'. Let E and F be pseudo-finite fields, K_1 a subfield of E and K_2 a subfield of F. Assume that we have an isomorphism ψ between K_1 and K_2. Then

$$(E, a)_{a \in K_1} \equiv (F, \psi(a))_{a \in K_1} \iff \text{there is } \psi' \supset \psi \text{ such that } \psi(E \cap K_1^{alg}) = F \cap K_2^{alg}.$$

[Here $(E, a)_{a \in K_1} \equiv (F, \psi(a))_{a \in K_1}$ means that in the language $\mathcal{L}(K_1)$, the constant corresponding to $a \in K_1$ is interpreted in F by the element $\psi(a)$.]

Proof. Extend ψ to some isomorphism θ defined on E. Then $(E, a)_{a \in K_1} \equiv (\theta(E), \psi(a))_{a \in K_1}$. By (6.13), we have

$$(\theta(E), \psi(a))_{a \in K_1} \equiv (F, \psi(a))_{a \in K_1} \iff \theta(E) \cap K_2^{alg} \simeq_{K_2} F \cap K_2^{alg}$$

and the result follows: K_2-isomorphisms between $\theta(E) \cap K_2^{alg}$ and $F \cap K_2^{alg}$ are in one-to-one correspondence with isomorphisms between $E \cap K_1^{alg}$ and $F \cap K_2^{alg}$ which extend ψ.

(6.16) Model completeness. We form the language $\mathcal{L}_c$ by adjoining to the language $\mathcal{L}$ of rings new constant symbols $c_{i,n}$, where $2 \leq n \in \mathbb{N}$ and $1 \leq i \leq n$. The theory Psf_c is obtained by adding to the theory Psf for each n an axiom stating that the polynomial $X^n + \sum_{i=1}^{n} c_{i,n} X^{n-i}$ is irreducible.

Note that every pseudo-finite field expands to a model of Psf_c: if F is pseudo-finite, for each n choose the $c_{i,n}$ to be the coefficients of some (monic) irreducible polynomial of degree n.

Theorem. The theory Psf_c is model complete.

Proof. Let $E \subseteq F$ be models of Psf_c. If L is an algebraic extension of E of degree n, then L is generated over E by a solution of the equation $X^n + \sum_{i=1}^{n} c_{i,n} X^{n-i}$. Since $F \models \mathrm{Psf}_c$, this polynomial stays irreducible over F, i.e., $F \cap L = E$. Then Corollary (6.12) gives us $E \equiv_E F$, i.e., $E \prec F$.

(6.17) Theorem. Let F be a pseudo-finite field, and $S \subset F^n$ be definable. Then there is an algebraic set $W \subseteq F^{n+m}$ such that, if $\pi : F^{n+m} \to F^n$ is the natural projection, then $\pi(W) = S$, and for each $y \in S$, the fiber $\pi^{-1}(y) \cap W$ is finite.

Proof. We may assume that S is $\emptyset$-definable. Expand F to a model of Psf_c. By model completeness of Psf_c, we know that every formula $\varphi(x)$ is equivalent modulo Psf_c to an existential formula (of $\mathcal{L}_c$). Since an inequation $x \neq 0$ is equivalent (modulo the theory of fields) to the formula $\exists y \; xy = 1$, we may assume that $\psi(x)$ is positive, and therefore there is an algebraic set $W \subset F^{n+m}$ such that $\pi(W) = S$. To check that W can be chosen so that the fibers are finite, we need to look a little closer at the way one obtains W.

By (6.15), we know that S is definable by a Boolean combination of formulas $\exists t \; f(x,t) = 0$, where $f(X,T) \in \mathbb{Z}[X,T]$, T a single variable. I.e., it is equivalent to a conjunction of disjunctions of formulas $\varphi_i(x)$, where $\varphi_i(x)$ is either $\exists t \; f_i(x,t) = 0$, or $\forall t \; f_i(x,t) \neq 0$.

Step 1. Replace each formula $\varphi_i(x)$ of the form $\exists t\ f_i(x,t)=0$ by an $\mathcal{L}$-formula $\psi_i(x)$ expressing that either the polynomial $f_i(x,T)$ is identically 0, (i.e., viewed as a polynomial in T its coefficients are all 0), or that some coefficient of $f_i(x,T)$ has a multiplicative inverse and $\exists t\ f_i(x,t)=0$. Thus $\psi_i(x)$ is the conjunction of a (positive) quantifier-free formula in x, and of an existential formula $\exists y \rho_i(y,x)$ which is such that, in any field F, for any tuple a in F, the formula $\rho_i(y,a)$ has only finitely many solutions.

Step 2. Assume now that $\varphi_i(x)$ is of the form $\forall t\ f(x,t)\neq 0$.

We will show that it is equivalent modulo Psf_c to the disjunction of a formula which says that the polynomial $f(x,T)$ is constant and takes a non-zero value (i.e., the conjunction of some equations with an existential formula saying that the constant coefficient is invertible), and of an existential formula $\exists y\ \rho_i(x,y)=0$ where y is a tuple of variables, $\rho_i(x,y)$ is a conjunction of equations, and for every field F and tuple a in F, the formula $\rho_i(x,a)$ has only finitely many solutions in F.

For that, it suffices to say the following: let m be the degree of $f(x,T)$ in T, and assume that $f(x,T)$ is not identically constant. Then F does not contain a root of this polynomial if and only adding a root of this polynomial defines a proper extension of F. I.e., if and only if, in the Galois extension L of F of degree $m!$ (which contains therefore all roots of $f(x,T)$), the polynomial $f(x,T)$ can be written $\prod_{i=1}^{m}(T-a_i)$ where the a_i's are not in F. Using the interpretation of L in F (using the constants $c_{i,m!}$), we see that this is expressible by an existential formula (see (6.2)). The variables which appear in the formula ρ_i with an existential quantifier will be:

– variables for the coefficients $b_{i,j}$ of the a_i's with respect to the basis $\{1,\alpha,\dots,\alpha^{m!-1}$, where α is a root of $X^{m!}+\sum c_{i,m!}X^{m!-i}$. Up to a permutation of the a_i's, these are uniquely determined.

– variables c_j with $1\leq j\leq m$, which appear in the equation $\prod_{i=2}^{m!} b_{i,j}c_j-1$ which implies that the element a_j is not in F.

This finishes the proof of the second step.

Step 3. So we have shown that $\varphi(x)$ is equivalent modulo Psf_c to a conjunction of disjunctions of positive existential formulas $\exists y\psi_i(x,y)$ where for each field F and tuple a, the set of elements satisfying $\psi_i(a,y)$ is finite. The conclusion follows observing that the conjunction of two formulas $\exists y \bigwedge f_1(x,y)=\cdots=f_m(x,y)=0$ and $\exists y \bigwedge g_1(x,y)=\cdots=g_n(x,y)=0$ is (logically) equivalent to $\exists y,z\ f_1(x,y)=\cdots=f_m(x,y)=g_1(x,z)=\cdots=g_n(x,z)=0$, and that their disjunction is equivalent (modulo the theory of fields) to $\exists y,z\ \bigwedge_{i,j} f_i(x,y)g_j(x,z)=0$.

An alternate and more model-theoretic proof, is simply to use Corollary (1.4) and Theorem (6.13). But one still needs to use the trick given in Step 2. Try it as an exercise, considering the set Δ of existential formulas of the form $\exists y\ \psi(x,y)$, where $\psi(x,y)$ is a conjunction of equations, and satisfying that for every field F and tuple a the set $\psi(a,y)$ is finite.

(6.18) Theorem. The pseudo-finite fields are exactly the infinite models of T_f. The pseudo-finite fields of characteristic 0 are exactly the infinite models of the set of sentences true in all prime fields.

Proof. In order to prove the first assertion, by Proposition (1.27) we need to show that if

F is a pseudo-finite field, then there is some non-principal ultrafilter $\mathcal{U}$ on the set $\mathcal{Q}$ of all powers of prime numbers, and such that

$$F \equiv \prod_{q\in\mathcal{Q}} \mathbb{F}_q/\mathcal{U}, \text{ i.e., by (6.14)}, \quad F \cap k^{alg} \simeq_k \left(\prod_{q\in\mathcal{Q}} \mathbb{F}_q/\mathcal{U}\right) \cap k^{alg},$$

where k denotes the prime subfield of F.

The proof is actually easier when $char(F) > 0$, but we will first do the case where $char(F) = 0$. In that case, note that the second assertion implies the first.

Let F be a pseudo-finite field of characteristic 0 and write $\mathbb{Q}^{alg}$ as the union of an increasing chain L_n, $n \in \mathbb{N}$, of finite Galois extensions of $\mathbb{Q}$. For each n, let $E_n = L_n \cap F$, and let $I(n)$ be the (finite) set of subfields of L_n which properly contain E_n. We will find a sentence θ_n which describes $L_n \cap F$. Choose a generator α of E_n over $\mathbb{Q}$, and let $f_n(T)$ be its minimal polynomial over $\mathbb{Q}$. Similarly, for each $M \in I(n)$, choose a generator β_M of M over $\mathbb{Q}$, let $g_M(T)$ be the minimal polynomial of β_M over $\mathbb{Q}$, and define $g_n(T) = \prod_{M\in I(n)} g_M(T)$. Consider now the sentence $\theta_n : \exists t\ f_n(t) = 0 \wedge \forall t\ g_n(t) \neq 0$. This is a sentence satisfied by F, and if E is any field of characteristic 0, then $E \models \theta_n \iff E \cap L_n \simeq E_n$.

The formula $\theta_0 \wedge \cdots \wedge \theta_n$ is still of the right form (since $\bigwedge_i \forall t\ g_i(t) \neq 0$ is equivalent to $\forall t\ \prod_i g_i(t) \neq 0$). It suffices to show that for every n, the set $A_n = \{p \text{ prime} \mid \mathbb{F}_p \models \theta_0 \wedge \cdots \wedge \theta_n\}$ is non-empty and infinite. Indeed, if $\mathcal{U}$ is an ultrafilter on the set $\mathcal{P}$ of all prime numbers and which contains all A_n's, then for each n, we will have, by Łos' theorem

$$F^* = \prod_{p\in\mathcal{P}} \mathbb{F}_p/\mathcal{U} \models \theta_n,$$

so that we will obtain

$$F^* \cap \mathbb{Q}^{alg} \simeq F \cap \mathbb{Q}^{alg}, \text{ i.e.,} F^* \equiv F.$$

The infiniteness of the sets A_n's follows from Tchebotarev's theorem. The consequence of Tchebotarev's theorem which we will use is :

Let $f_1(T), \dots, f_m(T), g(T) \in \mathbb{Z}[T]$, T a single variable. Let L be the Galois extension of $\mathbb{Q}$ obtained by adjoining all roots of the polynomials $f_i(T)$, $i = 1, \dots, m$. Assume that there is a subfield E of L such that $\mathcal{G}al(L/E)$ is cyclic and

$$E \models \bigwedge_{i=1}^{m} \exists t\ f_i(t) = 0 \wedge \forall t\ g(t) \neq 0.$$

Then the set of prime numbers p such that $\mathbb{F}_p \models \bigwedge_{i=1}^m \exists t\ f_i(t) = 0 \wedge \forall t\ g(t) \neq 0$ is infinite.

Let us now assume that F is of characteristic $p > 0$. If $F \cap \mathbb{F}_p^{alg}$ is infinite, then let I be an infinite sequence of integers such that if $n < m \in I$ then $n|m$ and $F \cap \mathbb{F}_p^{alg} = \bigcup_{n\in I} \mathbb{F}_{p^n}$. Take for $\mathcal{U}$ any non-principal ultrafilter containing $\{p^n \mid n \in I\}$. One then verifies easily that a polynomial $f(T) \in \mathbb{F}_p[T]$ has a root in F if and only if it has a root in all but finitely many of the $\mathbb{F}_{p^n}$, $n \in I$, if and only if it has a root in $\prod_{n\in\mathcal{Q}} \mathbb{F}_{p^n}/\mathcal{U}$.

If $F \cap \mathbb{F}_p^{alg}$ is finite, say of size p^n, then consider the set $I = \{p^{\ell n} \mid \ell \text{ a prime }\}$, and take for $\mathcal{U}$ any non-principal ultrafilter containing I. Again, one verifies that if $F^* = \prod_{m\in\mathbb{N}} \mathbb{F}_{p^m}/\mathcal{U}$, then $F^* \cap \mathbb{F}_p^{alg} \simeq \mathbb{F}_{p^n}$.

(6.19) Corollary/Remark. Observe that we have shown the following: let k be a prime field, E and algebraic extension of k which has at most one algebraic extension of each degree. Then there is some pseudo-finite field F such that $F \cap k^{alg} \simeq E$.

(6.20) Decidability issues. Observe first that by Theorem (6.18) we have

$$\text{Psf} \subset \text{Psf}_0 \subset T_{\text{prime}} \text{ and } \text{Psf} \subset T_f \subset T_{\text{prime}}.$$

We will first show that the theory Psf is decidable, that is, that there is an algorithm which decides, given a sentence θ, whether it is true in all pseudo-finite fields or not. From this we will be able to derive the decidability of the other theories.

We have an enumeration of a set Γ consisting of axioms for the theory Psf (this assumes that the bounds given in (5.2) on degrees of polynomials can be computed effectively, but they can). Hence, we can produce an enumeration of the set of all proofs made using axioms of Γ, and therefore of the theory Psf (by the completeness theorem, if a sentence is true in all pseudo-finite fields, then it is provable from Γ). Similarly, we have an enumeration of a set Γ_0 of axioms for the theory Psf_0 of all pseudo-finite fields of characteristic 0, and of the theory Psf_0. Note that $\Gamma_0 = \Gamma \cup \{p \neq 0 \mid p \text{ a prime}\}$.

This tells us that if θ is in Psf, then going through the enumeration of Psf we will find it. However, we need another procedure to decide if $\theta \notin \text{Psf}$. This is what we will do below. Let us fix a sentence θ.

Let ψ_n, $n \in \mathbb{N}$, be an enumeration of all sentences which are Boolean combinations of sentences of the form $\exists t\ f(t) = 0$, where $f(T) \in \mathbb{Z}[T]$. By (6.15), we know that $\text{Psf} \vdash \theta \leftrightarrow \psi_n$ for some n, i.e., $\theta \leftrightarrow \psi_n \in \Delta$, and therefore we can effectively find this ψ_n. Note that the proof of $\theta \leftrightarrow \psi_n$ uses only a finite number of axioms expressing the PAC property, and we can therefore find a constant C_1 (given by Lang-Weil (6.5)) such that in all finite field $\mathbb{F}_q$ with $q > C_1$ we have

$$\mathbb{F}_q \models \theta \leftrightarrow \psi_n.$$

It now remains to decide whether ψ_n is true in all pseudo-finite fields. I.e., we need to show that if k is a prime field, and $E \subset k^{alg}$ has at most one algebraic extension of each degree, then $E \models \psi_n$.

Step 1 Decide whether $\psi_n \in \text{Psf}_0$ or not.

Observe that ψ_n is (equivalent to) a disjunction of sentences of the form $\bigwedge_i \exists t\ f_i(t) = 0 \wedge \forall t\ g(t) \neq 0$. If $g(T)$ is not identically constant, then $\mathbb{Q}^{alg} \not\models \forall t\ g(t) \neq 0$, and therefore, we can assume that ψ_n is a disjunction of sentences of the form $\bigwedge_i \exists t\ f_i(t) = 0$. Let L be the splitting field of all polynomials appearing in ψ_n. Then one can compute effectively $\mathcal{G}al(L/\mathbb{Q})$, as well as those subfields E of L such that $\mathcal{G}al(L/E)$ is cyclic. Hence we an decide whether or not ψ_n is true in all subfields E of L such that $\mathcal{G}al(L/E)$ is cyclic. If it is not, then $\psi_n \notin \text{Psf}_0$ and therefore $\psi_n \notin \text{Psf}$, i.e., $\theta \notin \text{Psf}_0$, $\theta \notin \text{Psf}$.

Step 2 Decide whether $\psi_n \in \text{Psf}$.

Assume that $\psi_n \in \mathrm{Psf}_0$. The proof of ψ_n from Γ_0 only uses finitely many axioms expressing that the characteristic is $\neq p$, and therefore there is a constant C_2 such that ψ_n holds in all pseudo-finite fields of characteristic $p > C_2$. It therefore remains to check whether ψ_n holds in all those of characteristic $p \leq C_2$. Fix one such p. Then, as in step 1 we can assume that ψ_n is positive, and (because $\mathcal{G}al(\mathbb{F}_p^{alg}/\mathbb{F}_p)$ is pro-cyclic), it suffices to check that $\mathbb{F}_p \models \psi_n$. This is certainly decidable.

Step 3. Decidability of T_f and of T_{prime}.

We know that the equivalence $\theta \leftrightarrow \psi_n$ is true in all fields of size $> C_1$. If $\psi_n \notin \mathrm{Psf}$, then $\theta \notin T_f$ and we are done. Assume that $\psi_n \in \mathrm{Psf}$. Then θ is true in all fields of size $> C_1$, and we can decide whether or not it is true in all fields of size $\leq C_1$. This gives the decidability of T_f.

Similarly for T_{prime}, we reduce to the case where $\psi_n \in \mathrm{Psf}_0$. Then θ is true in all fields $\mathbb{F}_p$ with $p > C_1, C_2$, and it suffices to check what happens in the others. This shows the decidability of T_{prime}.

7. Measure, definability, and other applications

(7.1) Counting points. We saw in Theorem (6.18) that every pseudo-finite field is elementarily equivalent to an ultraproduct of finite fields. This implies in fact that every pseudo-finite field elementarily embeds into an ultraproduct of finite fields (an ultrapower of ultraproducts is an ultraproduct). Now, every finite field can be equipped with a measure (the counting measure), and one would think that the ultraproduct of these measures might define something interesting on F. It turns out that this is the case, and we will see below how it works. The main tool is the following

Theorem. Let $\varphi(x, y)$ be a formula, x an n-tuple of variables (y an m-tuple of variables). Then there is a finite set $D \subset \{0, 1, \dots, n\} \times \mathbb{Q}^{>0} \cup \{(0, 0)\}$ of pairs (d, μ), and a constant $C > 0$, formulas $\varphi_{d,\mu}(y)$ for $(d, \mu) \in D$ such that:

(1) If $\mathbb{F}_q$ is a finite field and a an m-tuple in $\mathbb{F}_q$, then there is some $(d, \mu) \in D$ such that

$$\big||\varphi(\mathbb{F}_q, a)| - \mu q^d\big| < Cq^{d-1/2}. \qquad (*)$$

[Here $\varphi(\mathbb{F}_q, a)$ denotes the set $\{b \in \mathbb{F}_q^n \mid \mathbb{F}_q \models \varphi(b, a)\}$.]

(2) The formula $\varphi_{d,\mu}(y)$ defines in each $\mathbb{F}_q$ the set of tuples a such that $(*)$ holds.

I am not going to give a proof of this result, although I will later sketch a strategy for the proof. With some work one can show that the constant C can be found effectively, see [FHJ2], and also [[FS], [FHJ1]. First a few remarks.

(7.2) Remarks.

(1) Observe that the pair $(0, 0)$ has been put in D to take care of the case when $\varphi(\mathbb{F}_q, a)$ is empty.

(2) If $\varphi(x, a)$ defines a variety V, then this is simply the Theorem of Lang-Weil, with $d = \dim(V)$ and $\mu = 1$.

(3) Thus, if $\varphi(x, a)$ defines an algebraic set W, all of whose irreducible components are defined over $\mathbb{F}_q$, then d will be the maximal dimension of the irreducible components of W, and μ the number of these components of maximal dimension. Note that

therefore, if $\varphi(x, y)$ is quantifier-free, then the associated set of pairs will be contained in $\{0, \ldots, n\} \times \mathbb{N}^{>0} \cup \{(0, 0)\}$.

(4) If q is sufficiently large, the formulas $\varphi_{d,\mu}(y)$ will define a partition of the parameter set $\mathbb{F}_q^m$.

(5) If $n = 1$, then there are positive numbers $A \in \mathbb{N}$ and $r \in \mathbb{Q}$ such that for every $\mathbb{F}_q$ and tuple a in $\mathbb{F}_q$,

$$\text{either } |\varphi(\mathbb{F}_q, a)| < A \text{ or } |\varphi(\mathbb{F}_q, a)| \geq rq.$$

Indeed, let D be the set of pairs (d, μ) associated to $\varphi(x, y)$; define $A_0 = \sup\{\mu \mid (0, \mu) \in D\}$, $r_0 = \inf\{\mu \mid (1, \mu) \in D\}$. Let $r = r_0/2$ and $A = \sup\{A_0 + C, 4C^2/r_0^2\}$. Using $(*)$, this gives the assertion.

(6) Observe that if q is sufficiently large, $(0, \mu) \in D$ and $\mathbb{F}_q \models \varphi_{0,\mu}(a)$, then, because $q^{-1/2}$ becomes very small, and in particular $< 1/2$, the number μ must give the exact size of the set $\varphi(\mathbb{F}_q, a)$ defined by $\varphi(x, a)$.

(7.3) Some simple applications of this result.

(1) There is no formula of the language of rings which defines in each field $\mathbb{F}_q^2$ the subfield $\mathbb{F}_q$.

(2) We know that the multiplicative group of $\mathbb{F}_q$ is cyclic, of order $q - 1$. There is no formula which defines in all fields $\mathbb{F}_q$ the set of generators of the multiplicative group $\mathbb{F}_q^\times$.

(3) Let G, H be groups definable in the pseudo-finite field F, and assume that $f : G \to H$ is definable, $Ker(f)$ is finite, and $\dim(G) = \dim(H) = d$. Then

$$\mu(G)[H : f(G)] = \mu(H)|Ker(f)|.$$

Proof. (1) If $\varphi(x)$ is a formula, there are $A > 0$ and $r \in \mathbb{Q}^{>0}$ such that for every finite field $\mathbb{F}_q$, the size of the set defined by φ is either $\leq A$ or greater than rq. hence, we cannot have a formula which defines in all $\mathbb{F}_{q^2}$ a set of size $\sqrt{q^2}$.

(2) The function ϕ (called the Euler function) giving the number of generating elements of a cyclic group can be computed. Note that if m, n are relatively prime integers then $\phi(nm) = \phi(n)\phi(m)$ (since $\mathbb{Z}/mn\mathbb{Z} \simeq \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}$). Also one has that $\phi(p^n) = (p-1)p^{n-1}$, since any lifting of a generator of $\mathbb{Z}/p\mathbb{Z}$ to $\mathbb{Z}/p^n/\mathbb{Z}$ is a generator of $\mathbb{Z}/p^n\mathbb{Z}$.

First observe that if $p^n > 2$, then $\phi(p^n) \geq \sqrt{n}$. Hence, for every $A \in \mathbb{N}$, the set of integers n such that $\phi(n) < A$ is finite.

We will now show that for every $\varepsilon > 0$, there is some prime power q such that $\phi(q-1) < \varepsilon(q-1)$. Observe that

$$\phi(n)/n = \prod_{\ell \text{ a prime divisor of } n} (1 - \frac{1}{\ell}).$$

Fix some prime p, and let $\ell_1, \ldots, \ell_m$ be distinct prime numbers, $M = \prod_{i=1}^m (\ell_i - 1)$. Then for every i, we have $p^M \equiv 1 \bmod(\ell_i)$ and therefore $\phi(p^M - 1) \leq (p^M - 1) \prod_{i=1}^m (1 - 1/\ell_i)$. Hence we can find arbitrarily small values of $\frac{\phi(p^M-1)}{p^M-1}$, which shows our assertion.

The existence of a formula defining the set of generators in all $\mathbb{F}_q$ would then, as in (1), contradict (7.2)(5).

(3) Let $F^* = \prod_{i\in I} \mathbb{F}_{q_i}/\mathcal{U}$ be an elementary extension of F, let a be a tuple of elements of F needed to define f, G and H (and their group law, and $(a(i))_i$ a sequence such that $[a(i)_i]_{\mathcal{U}} = a$.

Let $\varphi_1(x, a)$ be the formula defining G, $\psi_1(x, y, z, a)$ the one defining its group law, $\varphi_2(x, a)$ the formula defining H, $\psi_2(x, y, z, a)$ the one defining its group law, and $\theta(x, y, a)$ the formula defining the graph of f. The following property is then a first order property of the parameter a:

$\psi_i(x, y, z, a)$ is the graph of a group operation on the set defined by $\varphi_i(x, a)$ $(i = 1, 2)$, and $\theta(x, y, a)$ is the graph of a group morphism between the set defined by $\varphi_1(x, a)$ and the set defined by $\varphi_2(x, a)$, whose kernel is of size m.

Hence, by Łos' theorem, for a set $J \in \mathcal{U}$, we have, for all $j \in J$, that the following statement holds in $\mathbb{F}_{q_j}$:
$\psi_1(x, y, z, a(j))$ is the graph of a group operation on the set G_j defined by $\varphi_1(x, a(j))$, $\psi_2(x, y, z, a(j))$ is the graph of a group operation on the set H_j defined by $\varphi_2(x, a(j))$, and $\theta(x, y, a(j))$ is the graph of a group morphism $f_j : G_j \to H_j$, whose kernel is of size m.

But G_j and H_j are finite!! Hence we have $|G_j|[H_j : f_j(G_j)] = |H_j||Ker(f_j)|$. For q_j sufficiently large, dividing by q_j^d, we get $\mu(G_j)[H_j : f_j(G_j)] = \mu(H_j)|Ker(f)|$.

There is a first-order formula which expresses that fact, is satisfied in all $\mathbb{F}_{q_j}$ for $j \in J$, and therefore is satisfied by a in F^*, whence also in F. This gives the result.

(7.4) Very rough sketch of the proof of Theorem (7.1). The result is proved by induction on the complexity of formulas.

Let us first assume that $\varphi(x, y)$ is positive quantifier-free, that is, it is a disjunction of conjunction of equations (over $\mathbb{Z}$).

Let $\mathbb{F}_q$ be a finite field, and a a tuple in $\mathbb{F}_q$. Consider the set S defined by $\varphi(x, a)$. Then $S = W(\mathbb{F}_q)$, where W is the algebraic set given by the equations of $\varphi(x, a)$. However, we do not know that the Theorem of Lang-Weil can tell us the estimate of how many points there are: we will be able to apply this theorem only if *all irreducible components of* W *are defined over* $\mathbb{F}_q$. In order to be able to use Lang-Weil, we must therefore find an algebraic set W' such that $W'(\mathbb{F}_q) = W(\mathbb{F}_q)$ and all irreducible components of W' are defined over $\mathbb{F}_q$. This is done in the following fashion:

Write $W = W_1 \cup \cdots \cup W_m$ where each W_i is irreducible over $\mathbb{F}_q$. If W_i is a variety, then we know by Lang-Weil (6.5) that $|W(\mathbb{F}_q)| \sim q^{\dim(W_i)}$ and we do nothing. If W_i is not a variety, then W_i has several irreducible components, and any point in $W_i(\mathbb{F}_q)$ will belong to the intersection W_i' of these components, and we replace W_i by W_i'. We repeat the procedure and find eventually an algebraic set W', all of whose irreducible components are defined over $\mathbb{F}_q$ and such that $W'(\mathbb{F}_q) = W(\mathbb{F}_q)$. This procedure is effective, and using the results on bounds in polynomial rings, and we can write $W' = W_1' \cup \cdots \cup W_\ell'$, where the W_i' are varieties defined over $\mathbb{F}_q$. If d is the maximum of the dimensions of the W_i', and μ is the number of components of W' of dimension d, the result of Lang-Weil will then give us that $|W(\mathbb{F}_q)| \sim \mu q^d$. One also knows that having dimension d is an elementary property of the coefficients of a set of polynomials defining a variety. Thus, there is a formula $\varphi_a(y)$ satisfied by a in $\mathbb{F}_q$ and which expresses how we obtained W' from W, and that W' has exactly μ components of maximal dimension d. For each pair $(\mathbb{F}_q, a)$ we can find such a formula. By compactness, there are a finite number of those, say $\varphi_1(y), \dots, \varphi_k(y)$ such that in any finite field $\mathbb{F}$, we have $\mathbb{F} \models \forall y\ (\exists x \varphi(x,y)) \leftrightarrow (\bigvee_j \varphi_j(y))$. To each formula $\varphi_j(y)$ is associated a pair (d,μ), and we group them to obtain the desired $\varphi_{d,\mu}$.

The case of a quantifier-free formula $\varphi(x,y)$ follows, observing that modulo the theory of fields, an inequation $z \neq 0$ is equivalent to $\exists y\ yz = 1$. Thus, every quantifier-free definable set is in bijection, via a projection, with an algebraic set. We then use the first case.

Let us now assume that $\varphi(x,y)$ is arbitrary. Then, Theorem (6.17) tells us, using compactness, that there are positive quantifier-free $\mathcal{L}_c$-formulas $\psi_1(x,y,z), \dots, \psi_m(x,y,z)$ such that

$$\mathrm{Psf} \vdash \forall x, y\ (\varphi(x,y) \leftrightarrow \exists z \bigvee_j \psi_j(x,y,z)),$$

and furthermore such that for some integer N, in any field F one has

$$F \models \forall x, y\ (\exists z\ \psi_j(x,y,z) \to \exists^{\leq N} z\ \psi_j(x,y,z)).$$

The same equivalence holds in sufficiently large finite fields, say of size $\geq C'$ for some C' (only depending on $\varphi(x,y)$). Given some sufficiently large finite field $\mathbb{F}$ and tuple a in $\mathbb{F}$, we know by the previous steps how to estimate the size of the sets defined by the formulas $\psi_i(x,a,z)$. The problem is that the set defined by $\bigvee_j \psi_j(x,a,z)$ is not in bijection with the set defined by $\varphi(x,a)$: given some x in that set, there may be several z such that $\psi_j(x,a,z)$ holds. One uses a trick to transform the algebraic sets defined by the ψ_j, in such a way that we are able to count how many z are sitting above an x. Then we use some counting arguments and induction to conclude. The constant C of the Theorem will be sufficiently large so that, in field of size smaller than C' (and in which we do not necessarily have the equivalence), the inequality still holds. E.g., one can choose $C \geq C'^n$, where $n = |x|$.

(7.5) Definition of the measure on pseudo-finite fields. Let $\varphi(x,y)$ be a formula (x and n-tuple of variables), and D, $\varphi_{d,\mu}(y)$ the set and formulas given by Theorem (7.1). It follows from Remark (7.2)(6) that if F is a pseudo-finite field and a a tuple in F, then there will be a unique pair $(d,\mu) \in D$ such that $F \models \varphi_{d,\mu}(a)$. We then define $\dim(\varphi(x,a)) = d$ and $\mu(\varphi(x,a)) = \mu$. If S is the set defined by $\varphi(x,a)$, then we also write $\dim(S)$ and $\mu(S)$ respectively.

Proposition. Let F be a pseudo-finite field, S, T two definable sets.

(1) If V is a variety defined over F, then $\dim(V(F)) = \dim(V)$ and $\mu(V(F)) = 1$.

(2) Assume that $T \cap S = \emptyset$. Then

$$\mu(S \cup T) = \begin{cases} \mu(S) + \mu(T) & \text{if } \dim(S) = \dim(T), \\ \mu(S) & \text{if } \dim(S) > \dim(T), \\ \mu(T) & \text{if } \dim(S) < \dim(T). \end{cases}$$

(3) Assume that $f : S \to T$ is a definable function, which is onto. If for all $y \in T$, $\dim(f^{-1}(y)) = d$, then $\dim(S) = \dim(T) + d$. If moreover for every $y \in T$, $\mu(f^{-1}(y)) = m$, then $\mu(S) = m\mu(T)$.

(4) Let us define a function m_S on definable subsets of S as follows. Assume that $T \subset S$ is definable, and let $(d, \mu) = (\dim(S), \mu(S))$, $(e, \nu) = (\dim(T), \mu(T))$. Then

$$m_S(T) = \begin{cases} 0 & \text{if } e < d, \\ \nu/\mu & \text{if } d = e. \end{cases}$$

Then m_S is a finitely additive measure on the set of definable subsets of S.

(5) Let $\bar{S}$ be the Zariski closure of S (in F^{alg}. I.e., the smallest Zariski closed set containing S. It is defined over F). Then $\dim(S) = \dim(\bar{S})$. [That is, we are saying that the algebraic dimension of the algebraic set $\bar{S}$ coincides with the model-theoretic dimension of the set S]

Proof. (1) is clear.

Recall that F embeds elementarily in some ultraproduct $\prod_{q \in \mathcal{Q}} /\mathcal{U}$ of finite fields. Assume that S is defined by $\varphi(x, a)$, write $a = [a_q]_{\mathcal{U}}$, and S_q for the subset of $\mathbb{F}_q^n$ defined by $\varphi(x, a_q)$. Note that for some set $A \in \mathcal{U}$, we will then have $\mathbb{F}_q \models \varphi_{d,\mu}(a_q)$ for all $q \in A$, and therefore $|S_q| \sim \mu q^d$. A moment's thought shows that this gives items (2) - (4).

(5) By (6.17), there is an algebraic set $W(F) \subset F^{n+\ell}$ such that $S = \pi(W(F))$ and the restriction of the projection π to W is finite-to-one. Without loss of generality, $W(F)$ is Zariski dense in W, and by (3) we obtain that $\dim(W) = \dim(S)$. Working now in F^{alg}, we have that π is also finite-to-one on a Zariski-dense open subset of W, and therefore $\dim(W) = \dim(V)$ (algebraic dimensions). Since $V \supseteq \bar{S}$, we get that $\dim(V) = \dim(\bar{S})$.

(7.6) Existence of certain bounds. Let $\varphi(x, y)$ be a formula.

(1) There is a number M such that in any finite or pseudo-finite field F, the length of a chain of definable subsets of F^n defined by formulas $\varphi(x, a)$ for some tuples a in F, is bounded by M.

(2) There is a number M such that in any finite field or pseudo-finite field F, if S is a definable set and $(a_i)_{i \in I}$ is a set of tuples such that each $\varphi(x, a_i)$ defines a subset of S of the same dimension d as S, and for $i \neq j$, $\dim(\varphi(x, a_i) \wedge \varphi(x, a_j)) < d$, then $|I| \leq M$.

Proof. These two facts follow from general properties of measures. It suffices to show them for all pseudo-finite fields, since then they will be true in all sufficiently large finite fields, whence, taking into account the finitely many small finite fields, we will get the bound M.

(1) Assume that this is not the case, i.e., that there are such chains of arbitrarily large length. Then, going to a sufficiently saturated pseudo-finite field F, we can find a sequence $(a_i)_{i \in \mathbb{N}}$ of tuples in F such that if $i < j$ then the set S_j defined by $\varphi(x, a_j)$ is strictly contained in the set S_i defined by $\varphi(x, a_i)$. Let D be the finite set of pairs associated to φ. Because D is finite, we may, going to a subsequence, assume that for every $i \in \mathbb{N}$, $\dim(S_i) = d$ and $\mu(S_i) = \mu$. The proof is by induction on d.

If $d = 0$, then we know that μ is the size of the set S_i, and therefore $|I| = 1$. Assume $d > 0$ and that the result holds for all definable sets of smaller dimension. For $i > 0$ let $T_i = S_0 \setminus S_i$. Then the sets T_i, $i \in \mathbb{N}$, form a strictly increasing chain of subsets of S_0, and we have $\dim(T_i) < d$ (since $(\dim(S_i), \mu(S_i)) = (\dim(S_0), \mu(S_0))$). This contradicts the induction hypothesis and proves the result.

(2) Let D be the set of pairs associated to the formula $\varphi(x, y)$, and let ν be the inf of all μ such that $(d, \mu) \in D$. If $\varphi(x, a_i)$, $i \in I$, define subsets S_i of S such that $\dim(S_i) = d$ and $\dim(S_i \cap S_j) < d$, then we get $m_S(S_i) \geq \nu/\mu(S)$ and $m_S(S_i \cap S_j) = 0$. This gives $|I| \leq \mu(S)/\nu$.

8. Applications to finite fields

Applications are usually obtained in the following fashion. We have a family of sets defined over finite fields (for instance, $GL_n(\mathbb{F}_q)$, where n is fixed, q varies), and we know it has a certain property P. From this we deduce that the corresponding set defined in pseudo-finite fields (say, $GL_n(F)$) has also property P. From this one deduces a certain property Q of the set, and then conclude that Q is also true in all sufficiently large finite fields. We saw already an application of this strategy in the previous section, where we found bounds on the size of certain families of definable sets. Here are more applications. Most of this material comes from the paper [HP2] by Hrushovski-Pillay. The proofs are fairly elementary, and pleasant to read. They introduce some techniques which are classical in simplicity theory.

(8.1) Projective varieties. We defined affine varieties in section 4. Projective varieties are a more general concept. Let K be an algebraically closed field. We define the n-dimensional projective space over K, $\mathbb{P}^n(K)$ as follows: consider $K^{n+1} \setminus \{0\}$, and quotient by the equivalence relation

$$(x_0, \dots, x_n)E(y_0, \dots, y_n) \iff \exists\lambda\ x_0 = \lambda y_0, \dots, x_n = \lambda y_n.$$

The resulting object is $\mathbb{P}^n(K)$, and the E-equivalence class of $(x_0, \dots, x_n)$ is (often) denoted by $(x_0 : \dots : x_n)$. (Observe that $\mathbb{P}^n(K)$ is interpretable in K). Zariski-closed subsets of $\mathbb{P}^n(K)$ are the images (in K^{n+1}/E) of zero-sets of finite sets of homogeneous polynomials. There is a covering of $\mathbb{P}^n(K)$ by affine subsets: namely, for each $0 \leq i \leq n$, consider the set U_i of elements $(x_0 : \dots : x_n)$ with $x_i \neq 0$. There is a bijection $f_i : U_i \to K^n$ obtained by sending $(x_0 : \dots : x_n)$ to $(x_0/x_i, \dots, x_{i-1}/x_i, x_{i+1}/x_i, \dots, x_n/x_i)$ (i.e., one divides all coordinates of a representing element by the i-th one, then deletes the i-th coordinate). This bijection is continuous for the Zariski topology, and a subset W of $\mathbb{P}^n(K)$ is Zariski-closed if and only if for every i, $f_i(W \cap U_i)$ is Zariski-closed (in K^n).

A morphism $f : V \to W$ between two varieties is a function from $V(K)$ to $W(K)$, whose coordinate functions are given locally by polynomials (i.e., for every $a \in V(K)$, $f(a) \in W(K)$, and there is some open subset U of $V(K)$ containing a, and on which f is given by a tuple of elements of $K[V \cap U]$.

(8.2) Algebraic groups. Recall that an algebraic group is an algebraic set G (affine or projective), with a group law such that multiplication $G \times G \to G$ and the inverse map $G \to G$ are morphisms. This implies that these two maps are continuous (for the Zariski topology), and that G is a topological group. Hence, all results of section 3 apply.

The *connected component of* G is the (unique) irreducible component of the algebraic set G which contains the identity element 1, and it is denoted by G^0. If G is defined over K, then so is G^0 (because $1 \in G(K)$). The irreducible components of G are disjoint, and are the cosets of G^0 in G.

If $G = G^0$, then one says that G is *connected.*

Examples. The best-known examples are probably $GL_n(K)$ (the group of $n \times n$ invertible matrices), the additive group of K (usually denoted by $\mathbb{G}_a(K)$) and the multiplicative group $K^\times$ (usually denoted by $\mathbb{G}_m(K)$). A priori, $GL_n(K)$ is an *open subset* of K^n, since it is defined by $\det(x_{i,j}) \neq 0$, but it is in definable bijection with the subvariety of K^{n+1} defined by $\det(x_{i,j}y - 1 = 0$. Similarly for $K^\times$. Non-affine examples are elliptic curves, for instance the projective variety defined by $y^2z = x(x-z)(x-2z)$ (which can also be viewed as the affine variety defined by $y^2 = x(x-1)(x-2)$ to which one adds the point at infinity $(0:1:0)$).

(8.3) Some properties of algebraic groups. Let G be an algebraic group. Then all properties listed in (3.1) apply, since the Zariski topology on G is T_1. Moreover we have:

(1) The center of G, $Z(G)$, is an algebraic subgroup of G.

(2) A finite normal subgroup of a connected algebraic group is central.

Proof. (1) If $g \in G$, the centraliser $C_G(g)$ of g in G is a Zariski-closed subgroup of G (see (3.1)(7)). Now, $Z(G)$ is the intersection of all $C_G(g)$ and is therefore Zariski-closed.

(2) If H is a finite normal subgroup of G, and $h \in H$, then h has only finitely many conjugates under the action of G (since they are all in the finite group H), and this implies that $C_G(h)$ has finite index in G. Because G is connected, this implies that $C_G(h) = G$, i.e., that $h \in Z(G)$.

(8.4) Definable subgroups of algebraic groups. Let F be a pseudo-finite field, G an algebraic group defined over F, and H a definable subgroup of $G(F)$.

(1) Let $\bar{H}$ be the Zariski closure of H. Then $[\bar{H}(F) : H]$ is finite.

(2) Assume that $\bar{H} = G$. Then there is an algebraic group G' and a surjective morphism $f : G' \to G$, everything defined over F, such that $f(G'(F))$ is a subgroup of finite index of H (and therefore of $G(F)$), and $Ker(f)$ is finite (central).

Proof. (1) By Proposition (7.5)(5), we know that $\dim(\bar{H}) = \dim(H) = \dim(\bar{H}(F))$, and therefore $m_{\bar{H}(F)}(H) > 0$. Since the cosets of H in $\bar{H}(F)$ are disjoint, there are at most finitely many, i.e., $[H(F) : H]$ is finite.

(2) The proof of this result can be found in [HP2] and is fairly technical, I will not give it. In view of Theorem (6.17) it is not very suprising.

Remarks. (1) A morphism $f : G \to H$ of algebraic groups which is onto and has finite kernel is called an *isogeny*, and the groups G and H are said to be *isogenous.* Observe that since G is connected, the kernel of f is necessarily central: $Ker(f)$ is a finite normal subgroup of $G(K)$.

(2) Note the following immediate application of the fact that pseudo-finite fields behave like finite fields: if $f : G \to H$ is an isogeny of (connected) algebraic groups defined over the pseudo-finite field F, then $|Ker(f)(F)| = [H(F) : f(G(F))]$, see also (7.3).

(8.5) The following result is very useful. It is a generalisation of a theorem of Chevalley for algebraic groups, and also a generalisation of a model-theoretic result called *Zilber's irreducibility theorem* and which holds in superstable groups of finite U-rank.

Theorem. Let F be a pseudo-finite field, G an algebraic group defined over F. Let $X(i), i \in I$ be a family of definable subsets of $G(F)$ (no uniformity is assumed).

(1) Then there is a definable subgroup H of $G(F)$ and indices $i_1, \dots, i_M \in I$ such that

$$H \subset X(i_1)^{\pm 1} \cdot X(i_2)^{\pm 1} \cdots X i_M{}^{\pm 1},$$

and for every $i \in I$, $X(i)H/H$ is finite.

(2) If for every $i \in I$, the Zariski closure $\overline{X(i)}$ of $X(i)$ is a variety and contains 1 (the identity element of G), then in fact H equals the subgroup of $G(F)$ generated by all $X(i)$, $i \in I$, and its Zariski closure $\bar{H}$ equals the algebraic subgroup of G generated by all $\overline{X(i)}$, $i \in I$, and is contained in $\overline{X(i_1)}^{\pm 1} \cdot \overline{X(i_2)^{\pm 1}} \cdots \overline{X i_M{}^{\pm 1}}$.

Very rough sketch of proof. I will not give the proof, it is rather long. One starts by showing (2). The statement that the subgroup generated by the $\overline{X(i)}$ is an algebraic subgroup of G is classical, due to Chevalley. Without loss of generality the family $X(i)$ is closed under inverses. Grosso modo one considers all sets W_α of the form $\overline{X_{\alpha(1)}} \cdots \overline{X_{\alpha(m)}}$, takes one of maximal dimension, and then shows that $W_\alpha \cdot W_\alpha$ is the group H_0 generated by all $\overline{X(i)}$.

To get the result for the family $X(i)$, one then considers the set $U = X_{\alpha(1)} \cdots X_{\alpha(m)}$, and show it has same dimension as H_0. Lemma (8.6) below tells us that the subgroup H generated by U is definable, and since it has the same dimension as H_0, it must be of finite index in $H_0(F)$, and $H_0 = \bar{H}$.

Proof of (1) from (2). The irreducible components of each $\overline{X(i)}$ are defined over F; partitioning each $X(i)$ into finitely many pieces if necessary (and enlarging the family), we may therefore assume that each $\overline{X(i)}$ is irreducible. If $X(i)$ does not contain 1, then replace it by $T(i) = g_i^{-1} X(i)$ for some $g_i \in X(i)$. By (1), the group H generated by the $T(i)$ is definable, and we have $X(i)H = g_i^{-1}H$.

(8.6) Lemma. Let U be a definable subset of $G(F)$, and assume that $\dim(U) = \dim(G)$. Then the subgroup H generated by U is definable, and there is an integer M such that every element of H is the product of at most M elements of U or of U^{-1}.

(8.7) Corollary. Let G be an algebraic group defined over the pseudo-finite field F, and let H be a definable subgroup of G. Suppose that H is definably simple, i.e., that it has no proper non-trivial definable normal subgroup, and that is non-abelian. Then H is simple.

Proof. Without loss of generality, $\bar{H} = G$. Then $G = G^0$, since $G^0(F) \cap H$ is definable and normal in H. Also, $Z(G) \cap H = (1)$. Furthermore, H has no definable subgroup of finite index: if H_1 is such a subgroup, then $\bigcap_{g \in H} H_1^g$ is definable, of finite index in H, and normal.

Let N be a normal subgroup of H. If N is finite, and $1 \neq g \in N$ then $C_H(g)$ is a definable subgroup of H of finite index, and therefore equals H. But then $C_G(g) = \bar{H}$, which contradicts $Z(G) \cap H = 1$. Hence N cannot be finite, and the conjugation class of any $1 \neq g \in N$ is infinite. Pick $1 \neq g \in H$, let $X = g^H = \{h^{-1}gh \mid h \in H\}$ and let $Y = g^{-1}X$. Then $\bar{Y} = g^{-1}\bar{X} = g^{-1}g^G$ is irreducible, and contains 1. Hence, by (8.5), the subgroup U of H generated by Y is definable. It is clearly normal, contained in N and non-trivial: this give us the desired contradiction.

(8.8) Theorem. Let $n > 1$. There is an integer k such that every subgroup of $GL_n(\mathbb{F}_p)$ which is generated by elements of order p is of the form $\langle g_1 \rangle \cdot \ldots \cdot \rangle g_k \langle$ for some $g_1, \dots, g_k \in G$ of order p.

Sketch of proof. If $g \in GL_n(\mathbb{F}_p)$ is of order $p > n$, then

$$\langle g \rangle = \{\exp(t \log(g)) \mid t \in \mathbb{F}_p\},$$

where $\log(g) = (g - I_n) - (g - I_n)^2 + \cdots + (-1)^{n-1}(g - I_n)^{n-1}/(n-1)$ and $\exp(v) = v + v^2/2! + \cdots + v^{n-1}/(n-1)!$.

If there is no such k, we can find a increasing sequence $p(i)$, $i \in \mathbb{N}$, of prime numbers, and for each $i \in \mathbb{N}$, a subset A_i of elements of $GL_n(\mathbb{F}_{p(i)})$ of order $p(i)$, such that the subgroup G_i of $GL_n(\mathbb{F}_{p(i)})$ generated by A_i cannot be written as $\langle a_1 \rangle \cdot \ldots \cdot \langle a_i \rangle$ for any elements $a_1, \ldots, a_i \in G_i$ of order $p(i)$.

Consider the structure $\mathcal{M}_i = (\mathbb{F}_{p(i)}, +, \cdot, A_i)$ and let $\mathcal{M} = (F, +, \cdot, A)$ a non-principal ultraproduct of the $\mathcal{M}_i$'s. Then F is a pseudo-finite field, and A is a set of unipotent elements of $GL_n(F)$. For each $a \in A$, consider the (definable) subgroup $X(a) = \{\exp(t \log(a)) \mid t \in F\}$. Its Zariski closure is the subgroup (of $GL_n(F^{alg})$) defined by $Y(a) = \{\exp(t \log(a)) \mid t \in F^{alg}\}$, which is a variety. Hence, by Theorem (8.5), the subgroup H of $GL_n(F)$ generated by the $X(a)$, $a \in A$, is definable, and of the form $X(a_1) \cdot \ldots \cdot X(a_k)$ for some k and $a_1, \ldots, a_k \in A$.

There is a formula $\psi(x_1, \ldots, x_k)$ which is satisfied by $a_1, \ldots, a_k$ in $\mathcal{M}$, and which expresses that $X(a_1) \cdot \ldots \cdot X(a_k)$ is a subgroup G of $GL_n(F)$ which contains A. if for each $1 \leq \ell \leq k$, $(a_\ell(i))$ is a sequence representing the element a_i, by Łos' theorem, there is an infinite set J of integers i such that the tuples $(a_1(i), \ldots, a_k(i))$ satisfy ψ. But this contradicts our assumption on the structures $\mathcal{M}_i$ and gives the desired contradiction.

Remark. Observe also the following: let us call G^* the algebraic subgroup of $GL_n(F^{alg}$ generated by the $Y(a)$, $a \in A$. By (8.5) again, we have that $[G^*(F) : G] < \infty$. By compactness, this implies that there is some integer d such that, whenever G is a subgroup of $GL_n(\mathbb{F}_p)$ generated by elements of order p, if G^* denotes the algebraic subgroup generated by all $Y(a)$, $a \in G$ of order p, then $[GL_n(\mathbb{F}_p) : G] \leq d$. Furthermore, it follows (easily) that if $p > d$, then G contains all elements of order p of $G * (\mathbb{F}_p)$. This result was proved by Nori [N].

(8.9) Other results include:

Theorem. Let G be an almost simple algebraic group defined over the pseudo-finite field F (G is *almost simple* if $Z(G)$ is finite and $G/Z(G)$ is a simple (non-abelian) simple algebraic group). Then $G(F)$ has a smallest definable subgroup H of finite index, and $H/Z(H)$ is simple (as an abstract group).

Theorem. Let G be an almost simple algebraic subgroup of GL_n defined over $\mathbb{Z}$. Then there are a finite number of formulas $\psi_1(x, y), \ldots, \psi_m(x, y)$ such that whenever p is a prime and M is a maximal subgroup of $G(\mathbb{F}_p)$ then for some i and tuple a_i in $\mathbb{F}_p$, M is defined by the formula $\psi_i(x, a_i)$.

(8.10) Proposition (folklore). Let R be a domain which is finitely generated as a ring. If R is a field, then R is finite. If R is infinite, then for every $0 \neq a \in R$, there are infinitely many maximal ideals of R which do not contain a.

Proof. The proof is by induction on the number of generators of R, and it suffices to show the second assertion, since it implies the first one (a field has no maximal ideal other than

(0)). So, assume that R is infinite. If $n = 1$, then $R = \mathbb{Z}$, and the assertion is clear. Assume that the result holds for the subring S of R and that $R = S[t]$. There are two cases to consider:

Case 1. t is transcendental over S.

Then R is a polynomial ring over S. Assume first that S is finite (so that it is a field). Then $S[t]$ contains infinitely many irreducible polynomials, and each irreducible polynomial generates a maximal ideal. A non-zero element of R has only finitely many irreducible factors, and this gives the result.

If S is infinite and $0 \neq a(t) \in R$, then S contains infinitely many elements c such that $a(c) \neq 0$. If $a(c) \neq 0$, then by induction hypothesis, there are infinitely many maximal ideals P of S such that $a(c) \notin P$. Then $(P, t - c)$ is a maximal ideal of R which does not contain $a(t)$.

Case 2. t is algebraic over S.

If S is finite, then so is R and we are done. Assume that R is infinite, let $f(X) \in S[X]$ be a polynomial vanishing at t and of minimal degree. If b is its leading coefficient, then t is integral over $S[1/b]$, and therefore, if P is a maximal ideal of S not containing b, then P extends to a maximal ideal of R (not containing b): indeed, the ideal P extends uniquely to a maximal ideal P' of $S[1/b]$; the image of the monic polynomial $f(X)/b$ in the ring $S[1/b][X]$ is therefore non-constant, and has a non-constant irreducible factor.

Let $0 \neq a \in R$, and let $g(X) \in S[X]$ be a polynomial vanishing at a and of minimal degree. Since $a \neq 0$, the constant term c of the polynomial $g(X)$ is non-zero. Let P be a maximal ideal of R which does not contain bc (by induction hypothesis, there are infinitely many of those). By the discussion above, P generates a proper ideal of S. If P' is a maximal ideal of R which contains P, then $P' \cap S = P$. This finishes the proof.

(8.11) Let R be a domain, generated by $a_1, \dots, a_n$ as a ring. Let $X = \mathcal{M}ax(R)$ the set of maximal ideals of R. We define a topology on X as follows: a basis of open sets are the open sets $\mathcal{O}_a = \{P \in X \mid a \notin P\}$.

Proposition. Let R be as above, and assume that R is infinite.

(1) If $P \in X$, then R/P is finite.

(2) Let $m \in \mathbb{N}$. The set $\{P \in X \mid |R/P| \leq m\}$ is finite.

(3) If $a, b \in R$ are non-zero, then $\mathcal{O}_a \cap calo_b = \mathcal{O}_{ab}$. Hence a finite intersection of basic open sets is a basic open set. The space X is compact, and every singleton is closed.

(4) Let $0 \neq a \in R$. Then $\mathcal{O}_a$ is infinite.

Proof. (1) The quotient R/P is a field, which is generated by the images $a_1/P, \dots, a_n/P$ of $a_1, \dots, a_n$ in R/P. By (8.10), R/P must be finite.

(4) is also clear by (8.10).

(2) If $P, P' \in X$, we have $P = P'$ if and only if $a_1/P, \dots, a_n/P$ and $a_1/P', \dots, a_n/P'$ satisfy the same equations with coefficients in $\mathbb{Z}$. Hence, $P = P'$ if and only if there is an isomorphism between the field R/P and R/P' which sends a_i/P to a_i/P' for $1 \leq i \leq n$. But there are only finitely many isomorphism types of fields with n elements named, and having size $\leq m$.

(3) If P is a maximal ideal of R, then P is prime. Hence, if $a, b \in R$, then $ab \notin P \iff a \notin P$ and $b \notin P$. This shows the first assertion.

Let $P \in X$. If $Q \in X$, $Q \neq P$, there is $b_Q \in Q \setminus P$. Then $\{P\} = \bigcap_{Q \in X, Q \neq P} \mathcal{O}_{b_Q}$.

Let U_i, $i \in I$, be a family of open sets such that $X = \bigcup_i U_i$. By (1), every U_i is a union of basic open sets, and we may therefore assume that every U_i is a basic open set, i.e., of the form $\mathcal{O}_{a_i}$ for some $0 \neq a_i \in R$. Then $X = \bigcup_i \mathcal{O}_{a_i}$ is equivalent to $\emptyset = \bigcap_i (X \setminus \mathcal{O}_{a_i})$, i.e., no element of X contains all a_i's, i.e., the ideal generated by all a_i's contains 1. Hence there are $b_1, \dots, b_m \in R$, $i(1), \dots, i(m) \in I$ such that $1 = \sum_{j=1}^m b_i a_{i(j)}$. Then $X = \bigcup_{j=1}^m \mathcal{O}_{a_{i(j)}}$.

Note X is in general not Hausdorff.

(8.12) Exercise. Let R be an infinite domain, generated by $a_1, \dots, a_n$ as a ring. Let $(b_m)_{m \in \mathbb{N}}$ be an enumeration of the elements of R. For each m choose a maximal ideal $P(m)$ of R such that $b_0 \cdots b_m \notin P(m)$. Show that if $\mathcal{U}$ is any non-principal ultrafilter on $\mathbb{N}$, then R embeds naturally into $\big(\prod_{m \in \mathbb{N}} R/P(m)\big)/\mathcal{U}$.

(8.13) (8.12) can be used to show results on "reduction mod p" for almost all p. For instance, the following result appears in [MVW], and is proved in [HP2] using model-theoretic techniques:

Theorem. Let G be an algebraic subgroup of GL_n which is defined over $\mathbb{Q}$, and is almost simple and simply connected (*almost simple* means that $Z(G)$ is finite, and $G/Z(G)$ is simple non-abelian; *simply connected* means that there is no proper isogeny from any algebraic group onto G). Let Γ be a finitely generated subgroup of $G(\mathbb{Q})$ which is Zariski dense in G. The homomorphism $\pi_p : \mathbb{Z} \to \mathbb{F}_p$, "reduction mod p", then gives us for almost all p an algebraic subgroup G_p of GL_n (which is also almost simple and simply connected), and sends Γ to a subgroup Γ_p of $G_p(\mathbb{F}_p)$.

Then for almost all p, $\pi_p(\Gamma) = G_p(\mathbb{F}_p)$.

(8.14) The Frobenius. For each $q = p^n \in \mathcal{Q}$, consider the field $\mathbb{F}_p^{alg}$ together with the automorphism $\sigma_q : x \mapsto x^q$.

Let $\mathcal{U}$ be a non-principal ultrafilter on $\mathcal{Q}$, and consider the difference field

$$(K, \sigma^*) = \prod_{q \in \mathcal{Q}} (\mathbb{F}_p^{alg}, \sigma_q)/\mathcal{U}.$$

Hrushovski [H] and Macintyre [M] showed that this difference field is *generic*, i.e., that every finite system of difference equations over K (i.e., polynomial equations in $X, \sigma^*(X), \dots,$, where X is a tuple of variables) which has a solution in some difference field extending (K, σ^*) already has a solution in K^*. To do that, Hurshovski shows a Lang-Weil-type estimate of the number of solutions of certain systems in the difference fields $(\mathbb{F}_p^{alg}, si_q)$. By linearizing the system of equations, and applying a trick, one reduces genericity question to the following question: given varieties U and Vdefined over K^* and of the same dimension, with $V \subset U \times U^{\sigma^*}$ (U^{σ^*} denotes the variety obtained by applying σ^* to the defining equations of U), and such that the projection maps $V \to U$ and $V \to U^{\sigma^*}$ are dominant, show that there is some $a \in K^*$ such that $(a, \sigma^*(a)) \in V$. The result of Hrushovski, which implies the genericity of the difference fields (K^*, σ^*), is the following:

Theorem. Fix d and n. There is a constant $C > 0$ such that, for any $q \in \mathcal{Q}$, whenever U and V are varieties of dimension d defined over $\mathbb{F}_p^{alg}$ by polynomials in $\mathbb{F}_p^{alg}[X_1, \dots, X_n]_{\leq d}$

and $\mathbb{F}_p^{alg}[X_1 \dots, X_n, Y_1, \dots, Y_n]_{\leq} d$ respectively, and are such that $V \subset U \times U^{\sigma_q}$ and the projection maps $V \to U$ and $V \to U^{\sigma_q}$ are dominant, then

$$\big||\{a \in {\mathbb{F}_p^{alg}}^n \mid (a, \sigma_q(a)) \in V\}| - cq^d\big| < Cq^{d-1/2},$$

where $c = [\mathbb{F}_p^{alg}(V) : \mathbb{F}_p^{alg}(U)]/[\mathbb{F}_p^{alg}(V) : \mathbb{F}_p^{alg}(U)]_{\text{insep}}$.

Generic difference fields have very nice model-theoretic properties, similar to the ones exhibited by pseudo-finite fields. Observe that $Fix(si)$, the subfield of elements fixed by σ, is then pseudo-finite. Generic difference fields can be thought of as "universal models" for difference fields. As for pseudo-finite fields, the elementary theory of a generic difference field is entirely determined by the behaviour of the automorphism on the algebraic closure of the prime field. Again, one almost obtains quantifier-elimination; every definable set is the finite-to-one projection of a set defoned by difference equations. The theory of generic difference field is decidable, and any of its completions is supersimple and eliminates imaginaries. Furthermore, Ryten and Tomasic have shown that Hrushovski's result allows one to define a finitely additive measure on definable subsets of generic difference fields, similar to the one defined for pseudo-finite fields.

Bibliography

[A] J. Ax, The elementary theory of finite fields, Annals of Math. 88 (1968), 239 – 271.

[CK] C.C. Chang, H.J. Keisler, Model theory, North-Holland, Amsterdam 1977.

[CDM] Z. Chatzidakis, L. van den Dries, A. Macintyre, Definable sets over finite fields, J. reine u. ang. Math. 427 (1992), 107 – 135.

[DS] L. van den Dries, K. Schmidt, Bounds in the theory of polynomials rings over fields. A non-standard approach. Invent. Math. 76 (1984), 77 – 91.

[FHJ1] M. Fried, D. Haran, M. Jarden, Galois stratification over Frobenius fields, Adv. in Math. 51 (1984), 1 – 35.

[FHJ2] M. Fried, D. Haran, M. Jarden, Effective counting of the points of definable sets over finite fields, Israel J. Math. 85 (1994), 103 – 133.

[FJ] M. Fried, M. Jarden, Field Arithmetic, Ergebnisse 11, Springer Berlin-Heidelberg 1986.

[FS] M. Fried, G. Sacerdote, Solving diophantine problems over all residue class fields of a number field and all finite fields, Annals of Math. 104 (1976), 203 – 233.

[He] G. Hermann, Die Frage der endlich vielen Schritte in der Theorie der Polynomideale, Math. Ann. 95 (1926), no. 1, 736 – 788.

[H] E. Hrushovski, The first-order theory of the Frobenius, preprint, available at ArXiv: `http://front.math.ucdavis.edu/math.LO/0406514`.

[HP1] E. Hrushovski, A. Pillay, Groups definable in local fields and pseudo-finite fields, Israel J. of Math. 85 (1994), 203 – 262.

[HP2] E. Hrushovski, A. Pillay, Definable subgroups of algebraic groups over finite fields, J. reine angew. Math. 462 (1995), 69 – 91.

[L1] S. Lang, Introduction to algebraic geometry, Addison-Wesley Pub. Co., Menlo Park 1973.

[L2] S. Lang, Algebra, Addison-Wesley Pub. Co., Menlo Park 1984.

[M] A. Macintyre, Nonstandard Frobenius, in preparation.

[MVW] C.R. Matthews, L.N. Vaserstein and B. Weisfeiler, Congruence properties of Zariski dense subgroups, Proc. London ¡ath. Soc. XLVIII (1984), 385 – 576.

[N] M.V. Nori, On subgroups of $GL_n(\mathbb{F}_p)$, Inventiones Math. 88 (1987), 257 – 275.

[S] A. Seidenberg, Constructions in algebra, Trans. Amer. Math. Soc. 197 (1974), 273 – 313.