Counting modulo N in pseudo-finite fields

Will Johnson

October 18, 2018

Outline

1. Model-theoretic Euler characteristics
2. Pseudo-finite fields and difference fields
3. $\mathbb{Z}/N$-valued Euler characteristics on pseudo-finite fields.

Definable sets and functions

If M is a structure, $\varphi(\vec{x};\vec{y})$ is a first-order formula, and $\vec{m}$ is a tuple from M with $|\vec{m}| = |\vec{y}|$, then

$$\varphi(M;\vec{m}) := \{\vec{a} \in M^{|\vec{x}|} \ : \ M \models \vec{(}\vec{a};\vec{m})\}$$

Sets of the form $\varphi(M;\vec{m})$ are called *definable sets*.

Definable sets and functions

If M is a structure, $\varphi(\vec{x};\vec{y})$ is a first-order formula, and $\vec{m}$ is a tuple from M with $|\vec{m}| = |\vec{y}|$, then

$$\varphi(M;\vec{m}) := \{\vec{a} \in M^{|\vec{x}|} \ : \ M \models \vec{(}\vec{a};\vec{m})\}$$

Sets of the form $\varphi(M;\vec{m})$ are called *definable sets*.

Definition

If X and Y are definable sets, a function $f : X \to Y$ is a *definable function* if its graph $\Gamma_f \subseteq X \times Y$ is a definable set.

Definable families

Definition

If X and Y are definable sets, a family $\{D_b \ : \ b \in Y\}$ of subsets of X is a *definable family* if the set

$$\coprod_{b \in Y} D_b := \{(a, b) \in X \times Y : a \in D_b\}$$

is a definable subset of $X \times Y$. The definable set $\coprod_{b \in Y} D_b$ is called the *total space* of the family.

Definable functions to abstract sets

Definition

A function $f : X \to R$ from a definable set X to an abstract set R is *definable* if the range $f(X)$ is finite and if $f^{-1}(r)$ is a definable subset of X for every $r \in R$.

Definable functions to abstract sets

Definition

A function $f : X \to R$ from a definable set X to an abstract set R is *definable* if the range $f(X)$ is finite and if $f^{-1}(r)$ is a definable subset of X for every $r \in R$.

Example

In an $\aleph_1$-categorical theory, "Morley rank is definable" means that for every definable family $\{D_a \ : \ a \in X\}$, the function

$$a \mapsto RM(D_a)$$

is a definable function from X to $\mathbb{N}$.

O-minimality

Definition

An ordered structure $(M, <, \ldots)$ is *o-minimal* if every definable subset of M^1 is a finite union of points and intervals.

O-minimality

Definition

An ordered structure $(M, <, \ldots)$ is *o-minimal* if every definable subset of M^1 is a finite union of points and intervals.

Examples:

- The totally ordered set $(\mathbb{Q}, <)$

O-minimality

Definition

An ordered structure $(M, <, \ldots)$ is *o-minimal* if every definable subset of M^1 is a finite union of points and intervals.

Examples:

- The totally ordered set $(\mathbb{Q}, <)$
- The ordered field $(\mathbb{R}, <, +, \cdot)$

O-minimality

Definition

An ordered structure $(M, <, \ldots)$ is *o-minimal* if every definable subset of M^1 is a finite union of points and intervals.

Examples:

- The totally ordered set $(\mathbb{Q}, <)$
- The ordered field $(\mathbb{R}, <, +, \cdot)$

Definable sets are very well-behaved in o-minimal structures, thanks to dimension theory, cell decomposition, etc.

O-minimality

Definition

An ordered structure $(M, <, \ldots)$ is *o-minimal* if every definable subset of M^1 is a finite union of points and intervals.

Examples:

- The totally ordered set $(\mathbb{Q}, <)$
- The ordered field $(\mathbb{R}, <, +, \cdot)$

Definable sets are very well-behaved in o-minimal structures, thanks to dimension theory, cell decomposition, etc.

Definition

An o-minimal structure $(M, <, \ldots)$ is *dense* if the underlying order is dense, i.e.,

$$M \models \forall x, y : (x < y) \implies (\exists z : (x < z) \wedge (z < y))$$

O-minimal Euler characteristic

Let $(M, <, \ldots)$ be a dense o-minimal structure. If X is a definable set, then X can be written as a disjoint union of cells

$$X = C_1 \coprod C_2 \coprod \cdots \coprod C_n$$

O-minimal Euler characteristic

Let $(M, <, \ldots)$ be a dense o-minimal structure. If X is a definable set, then X can be written as a disjoint union of cells

$$X = C_1 \coprod C_2 \coprod \cdots \coprod C_n$$

Each cell C_i has a dimension $\dim(C_i)$. If $M = \mathbb{R}$, then cells of dimension k are homeomorphic to the k-dimensional open ball.

O-minimal Euler characteristic

Let $(M, <, \ldots)$ be a dense o-minimal structure. If X is a definable set, then X can be written as a disjoint union of cells

$$X = C_1 \coprod C_2 \coprod \cdots \coprod C_n$$

Each cell C_i has a dimension $\dim(C_i)$. If $M = \mathbb{R}$, then cells of dimension k are homeomorphic to the k-dimensional open ball.

Definition

The *Euler characteristic* of X is

$$\chi(X) = \sum_{i=1}^{n} (-1)^{\dim(C_i)}$$

O-minimal Euler characteristic

Let $(M, <, \ldots)$ be a dense o-minimal structure. If X is a definable set, then X can be written as a disjoint union of cells

$$X = C_1 \coprod C_2 \coprod \cdots \coprod C_n$$

Each cell C_i has a dimension $\dim(C_i)$. If $M = \mathbb{R}$, then cells of dimension k are homeomorphic to the k-dimensional open ball.

Definition

The *Euler characteristic* of X is

$$\chi(X) = \sum_{i=1}^{n} (-1)^{\dim(C_i)}$$

Proposition (van den Dries)

This doesn't depend on the choice of the cell decomposition.

Properties of Euler characteristic

O-minimal Euler characteristic satisfies the following properties:

1. $\chi(X) = \chi(Y)$ if there is a definable bijection $f : X \to Y$.

Properties of Euler characteristic

O-minimal Euler characteristic satisfies the following properties:

1. $\chi(X) = \chi(Y)$ if there is a definable bijection $f : X \to Y$.
2. $\chi(\emptyset) = 0$
3. $\chi(X \cup Y) = \chi(X) + \chi(Y)$ if $X \cap Y = \emptyset$

Properties of Euler characteristic

O-minimal Euler characteristic satisfies the following properties:

1. $\chi(X) = \chi(Y)$ if there is a definable bijection $f : X \to Y$.
2. $\chi(\emptyset) = 0$
3. $\chi(X \cup Y) = \chi(X) + \chi(Y)$ if $X \cap Y = \emptyset$
4. $\chi(\{a\}) = 1$
5. $\chi(X \times Y) = \chi(X) \times \chi(Y)$

Properties of Euler characteristic

O-minimal Euler characteristic satisfies the following properties:

1. $\chi(X) = \chi(Y)$ if there is a definable bijection $f : X \to Y$.
2. $\chi(\emptyset) = 0$
3. $\chi(X \cup Y) = \chi(X) + \chi(Y)$ if $X \cap Y = \emptyset$
4. $\chi(\{a\}) = 1$
5. $\chi(X \times Y) = \chi(X) \times \chi(Y)$

Additional properties

O-minimal Euler characteristic satisfies the following additional properties:

1. If $\{D_x : x \in X\}$ is a definable family, then the function

$$x \mapsto \chi(D_x)$$

is a definable function from X to $\mathbb{Z}$.

Additional properties

O-minimal Euler characteristic satisfies the following additional properties:

1. If $\{D_x : x \in X\}$ is a definable family, then the function
$$x \mapsto \chi(D_x)$$
is a definable function from X to $\mathbb{Z}$.
2. If $\{D_x : x \in X\}$ is a definable family and $\chi(D_x) = \alpha$ for all x, then
$$\chi\left(\coprod_{x \in X} D_x\right) = \alpha \cdot \chi(X)$$

Additional properties

O-minimal Euler characteristic satisfies the following additional properties:

1. If $\{D_x : x \in X\}$ is a definable family, then the function
$$x \mapsto \chi(D_x)$$
is a definable function from X to $\mathbb{Z}$.
2. If $\{D_x : x \in X\}$ is a definable family and $\chi(D_x) = \alpha$ for all x, then
$$\chi\left(\coprod_{x \in X} D_x\right) = \alpha \cdot \chi(X)$$

Example

If $1 \to G_1 \to G_2 \to G_3 \to 1$ is a short exact sequence of definable groups, then
$$\chi(G_2) = \chi(G_1) \cdot \chi(G_3)$$

Additional properties

O-minimal Euler characteristic satisfies the following additional properties:

1. If $\{D_x : x \in X\}$ is a definable family, then the function
$$x \mapsto \chi(D_x)$$
is a definable function from X to $\mathbb{Z}$.
2. If $\{D_x : x \in X\}$ is a definable family and $\chi(D_x) = \alpha$ for all x, then
$$\chi\left(\coprod_{x \in X} D_x\right) = \alpha \cdot \chi(X)$$

Example

If $1 \to G_1 \to G_2 \to G_3 \to 1$ is a short exact sequence of definable groups, then
$$\chi(G_2) = \chi(G_1) \cdot \chi(G_3)$$

Proof.

Consider the family of cosets of G_1 in G_2, indexed by G_3. □

Euler characteristic in topology

If X is a sufficiently nice topological space, the Euler characteristic $\chi(X)$ is defined as the alternating sum of Betti numbers:

$$\chi(X) = \sum_{i=0}^{\infty} (-1)^i \dim_{\mathbb{C}} H^i(X; \mathbb{C}).$$

Euler characteristic in topology

If X is a sufficiently nice topological space, the Euler characteristic $\chi(X)$ is defined as the alternating sum of Betti numbers:

$$\chi(X) = \sum_{i=0}^{\infty} (-1)^i \dim_{\mathbb{C}} H^i(X; \mathbb{C}).$$

When X is a finite simplicial complex, $\chi(X)$ is given by a more familiar formula

$$\chi(X) = \sum_{i=0}^{\infty} (-1)^i F_i$$

where F_i is the number of i-dimensional simplices.

Example

Euler's formula

$$V - E + F = 2$$

is the statement that the Euler characteristic of the sphere is 2.

The reason for the name "Euler characteristic"

Theorem (van den Dries)

If R is an o-minimal expansion of the real field $(\mathbb{R}, <, +, \cdot)$, and $X \subseteq \mathbb{R}^n$ is a definable and compact set, then the o-minimal Euler characteristic of X agrees with the topological Euler characteristic.

Abstract Euler characteristics

Let M be a structure and R be a (commutative) ring.

Definition

An *R-valued Euler characteristic on M* is a map

$$\chi : \{\text{Definable sets in } M\} \to R$$

satisfying the axioms

1. $\chi(X) = \chi(Y)$ if there is a definable bijection $f : X \to Y$.
2. $\chi(\emptyset) = 0$
3. $\chi(X \cup Y) = \chi(X) + \chi(Y)$ if $X \cap Y = \emptyset$
4. $\chi(\{a\}) = 1$
5. $\chi(X \times Y) = \chi(X) \times \chi(Y)$

Strong and definable Euler characteristics

Definition

An Euler characteristic χ is *definable* if for every definable family $\{D_x : x \in X\}$, the following function is definable:

$$x \mapsto \chi(D_x)$$

Strong and definable Euler characteristics

Definition

An Euler characteristic χ is *definable* if for every definable family $\{D_x : x \in X\}$, the following function is definable:

$$x \mapsto \chi(D_x)$$

Definition

An Euler characteristic χ is *strong* if for every definable family $\{D_x : x \in X\}$, if $\chi(D_x) = \alpha$ for all $x \in X$ then

$$\chi\left(\coprod_{x \in X} D_x\right) = \alpha \cdot \chi(X)$$

Examples

- O-minimal Euler characteristic

Examples

- O-minimal Euler characteristic
- Size, in finite structures:

$$\chi(X) = \#X$$

Examples

- O-minimal Euler characteristic
- Size, in finite structures:

$$\chi(X) = \#X$$

Both these examples are strong and definable.

"Measures"

An R-valued "measure" on a definable set X is a function

$$\mu : \{\text{Definable subsets of } X\} \to R$$

satisfying the axioms:

$$\mu(\emptyset) = 0$$
$$\mu(X \cup Y) = \mu(X) + \mu(Y) \quad \text{if } X \cap Y = \emptyset$$

"Measures"

An R-valued "measure" on a definable set X is a function

$$\mu : \{\text{Definable subsets of } X\} \to R$$

satisfying the axioms:

$$\mu(\emptyset) = 0$$
$$\mu(X \cup Y) = \mu(X) + \mu(Y) \quad \text{if } X \cap Y = \emptyset$$

A measure μ is *definable* if for any definable family $\{D_y \ : \ y \in Y\}$ of subsets of X, the following is a definable function from Y to R:

$$x \mapsto \mu(D_x).$$

"Measures"

An R-valued "measure" on a definable set X is a function

$$\mu : \{\text{Definable subsets of } X\} \to R$$

satisfying the axioms:

$$\mu(\emptyset) = 0$$
$$\mu(X \cup Y) = \mu(X) + \mu(Y) \quad \text{if } X \cap Y = \emptyset$$

A measure μ is *definable* if for any definable family $\{D_y \ : \ y \in Y\}$ of subsets of X, the following is a definable function from Y to R:

$$x \mapsto \mu(D_x).$$

Remark

An R-valued Euler characteristic χ induces an R-valued measure $\chi|_X$ on every definable set X. The Euler characteristic χ is definable iff $\chi|_X$ is definable for every X.

Integration against a "measure"

If μ is an R-valued measure on X, and $f : X \to R$ is definable, there is a well-defined "integral"

$$\int_{x \in X} f(x)\ d\mu(x).$$

Integration against a "measure"

If μ is an R-valued measure on X, and $f : X \to R$ is definable, there is a well-defined "integral"

$$\int_{x \in X} f(x)\ d\mu(x).$$

This can be defined by partitioning X into subsets

$$X = X_1 \coprod \cdots \coprod X_n$$

on which f is constant

$$f(X_i) = \alpha_i,$$

Integration against a "measure"

If μ is an R-valued measure on X, and $f : X \to R$ is definable, there is a well-defined "integral"

$$\int_{x \in X} f(x)\ d\mu(x).$$

This can be defined by partitioning X into subsets

$$X = X_1 \coprod \cdots \coprod X_n$$

on which f is constant

$$f(X_i) = \alpha_i,$$

and then setting

$$\int_{x \in X} f(x)\ d\mu(x) := \sum_{i=1}^{n} \alpha_i \cdot \mu(X_i)$$

Integration against a "measure"

If μ is an R-valued measure on X, and $f : X \to R$ is definable, there is a well-defined "integral"

$$\int_{x \in X} f(x)\ d\mu(x).$$

This can be defined by partitioning X into subsets

$$X = X_1 \coprod \cdots \coprod X_n$$

on which f is constant

$$f(X_i) = \alpha_i,$$

and then setting

$$\int_{x \in X} f(x)\ d\mu(x) := \sum_{i=1}^{n} \alpha_i \cdot \mu(X_i)$$

Remark

This does not depend on the choice of the partition.

Properties of integration

Integration is R-linear:

$$\int_{x \in X} (a \cdot f(x) + g(x))\ d\mu(x) = a \cdot \int_{x \in X} f(x)\ d\mu(x) + \int_{x \in X} g(x) d\mu(x).$$

Properties of integration

Integration is R-linear:

$$\int_{x \in X} (a \cdot f(x) + g(x))\ d\mu(x) = a \cdot \int_{x \in X} f(x)\ d\mu(x) + \int_{x \in X} g(x) d\mu(x).$$

If 1_D is the characteristic function of a definable subset $D \subseteq X$, then

$$\int_{x \in X} 1_D(x)\ d\mu(x) = \mu(D).$$

Properties of integration

Integration is R-linear:

$$\int_{x \in X} (a \cdot f(x) + g(x))\ d\mu(x) = a \cdot \int_{x \in X} f(x)\ d\mu(x) + \int_{x \in X} g(x) d\mu(x).$$

If 1_D is the characteristic function of a definable subset $D \subseteq X$, then

$$\int_{x \in X} 1_D(x)\ d\mu(x) = \mu(D).$$

If μ is definable and $f : X \times Y \to R$ is definable, then

$$\int_{x \in X} f(x, y)\ d\mu(x)$$

is a definable function $Y \to R$. Moreover, this property characterizes definable measures.

Strong Euler characteristics and integration

Let χ be a strong Euler characteristic and $\{D_x : x \in X\}$ be a definable family.
Suppose the following function $X \to R$ is definable:

$$x \mapsto \chi(D_x)$$

Strong Euler characteristics and integration

Let χ be a strong Euler characteristic and $\{D_x : x \in X\}$ be a definable family.
Suppose the following function $X \to R$ is definable:

$$x \mapsto \chi(D_x)$$

Then, for $\mu = \chi|_X$, we have

$$\chi\left(\coprod_{x \in X} D_x\right) = \int_{x \in X} \chi(D_x)\ d\mu(x)$$

Strong Euler characteristics and integration

Let χ be a strong Euler characteristic and $\{D_x : x \in X\}$ be a definable family.
Suppose the following function $X \to R$ is definable:

$$x \mapsto \chi(D_x)$$

Then, for $\mu = \chi|_X$, we have

$$\chi\left(\coprod_{x \in X} D_x\right) = \int_{x \in X} \chi(D_x)\ d\mu(x)$$

In particular, χ of the total space is determined by χ of the fibers and the restriction of χ on the base.

Strong definable Euler characteristics

Let χ be a strong Euler characteristic on a 1-sorted structure M.

Proposition

The following are equivalent:

- χ *is definable*
- $\chi|_{M^1}$ *is definable*

Strong definable Euler characteristics

Let χ be a strong Euler characteristic on a 1-sorted structure M.

Proposition

The following are equivalent:

- χ *is definable*
- $\chi|_{M^1}$ *is definable*

Moreover, when these hold, χ *is determined by* $\chi|_{M^1}$.

Strong definable Euler characteristics

Let χ be a strong Euler characteristic on a 1-sorted structure M.

Proposition

The following are equivalent:

- χ *is definable*
- $\chi|_{M^1}$ *is definable*

Moreover, when these hold, χ *is determined by* $\chi|_{M^1}$.

Proof.

Let $\mu = \chi|_{M^1}$. Assuming μ is definable, one proves by induction that for any definable $X \subseteq M^n$,

$$\chi(D) = \int_{x_1 \in M} \cdots \int_{x_n \in M} 1_D(\vec{x})\ d\mu(x_n) \cdots d\mu(x_1)$$

$\square$

Constructing strong definable Euler characteristics

Proposition

Let M be a 1-sorted structure and μ be a definable measure on M^1. Then μ comes from a strong definable Euler characteristic iff these axioms hold:

Constructing strong definable Euler characteristics

Proposition

Let M be a 1-sorted structure and μ be a definable measure on M^1. Then μ comes from a strong definable Euler characteristic iff these axioms hold:

- $\mu(\{a\}) = 1$ *for every* $a \in M^1$.

Constructing strong definable Euler characteristics

Proposition

Let M be a 1-sorted structure and μ be a definable measure on M^1. Then μ comes from a strong definable Euler characteristic iff these axioms hold:

- $\mu(\{a\}) = 1$ *for every* $a \in M^1$.
- *For every definable* $D \subseteq M^2$, *the Fubini property holds:*

$$\int_{x \in M^1} \int_{y \in M^1} 1_D(x, y)\, d\mu(y)\, d\mu(x) = \int_{y \in M^1} \int_{x \in M^1} 1_D(x, y)\, d\mu(x)\, d\mu(y$$

Constructing strong definable Euler characteristics

Proposition

Let M be a 1-sorted structure and μ be a definable measure on M^1. Then μ comes from a strong definable Euler characteristic iff these axioms hold:

- $\mu(\{a\}) = 1$ *for every* $a \in M^1$.
- *For every definable* $D \subseteq M^2$, *the Fubini property holds:*

$$\int_{x \in M^1} \int_{y \in M^1} 1_D(x, y)\, d\mu(y)\, d\mu(x) = \int_{y \in M^1} \int_{x \in M^1} 1_D(x, y)\, d\mu(x)\, d\mu(y$$

This can be used to construct the o-minimal Euler characteristic.

Example: ACF_0

- The field $\mathbb{C}$ can be interpreted in the field $\mathbb{R}$. The o-minimal Euler characteristic $\chi_{\mathbb{R}}$ on $(\mathbb{R}, +, \cdot)$ yields a strong Euler characteristic $\chi_{\mathbb{C}}$ on $(\mathbb{C}, +, \cdot)$.

Example: ACF_0

- The field $\mathbb{C}$ can be interpreted in the field $\mathbb{R}$. The o-minimal Euler characteristic $\chi_{\mathbb{R}}$ on $(\mathbb{R}, +, \cdot)$ yields a strong Euler characteristic $\chi_{\mathbb{C}}$ on $(\mathbb{C}, +, \cdot)$.
- The induced measure on $\mathbb{C}^1$ is

$$\begin{aligned} \mu(X) &= |X| \quad \text{if } X \text{ finite} \\ \mu(\mathbb{C}) &= 1 \quad (\text{because } \mathbb{C} \text{ is a 2-cell}) \\ \mu(\mathbb{C} \setminus X) &= 1 - |X| \quad \text{if } X \text{ is finite} \end{aligned}$$

Example: ACF_0

- The field $\mathbb{C}$ can be interpreted in the field $\mathbb{R}$. The o-minimal Euler characteristic $\chi_{\mathbb{R}}$ on $(\mathbb{R}, +, \cdot)$ yields a strong Euler characteristic $\chi_{\mathbb{C}}$ on $(\mathbb{C}, +, \cdot)$.
- The induced measure on $\mathbb{C}^1$ is

$$\begin{aligned} \mu(X) &= |X| \quad \text{if } X \text{ finite} \\ \mu(\mathbb{C}) &= 1 \quad (\text{because } \mathbb{C} \text{ is a 2-cell}) \\ \mu(\mathbb{C} \setminus X) &= 1 - |X| \quad \text{if } X \text{ is finite} \end{aligned}$$

- This measure is definable, so $\chi_{\mathbb{C}}$ is definable.

Example: ACF_0

- The field $\mathbb{C}$ can be interpreted in the field $\mathbb{R}$. The o-minimal Euler characteristic $\chi_{\mathbb{R}}$ on $(\mathbb{R}, +, \cdot)$ yields a strong Euler characteristic $\chi_{\mathbb{C}}$ on $(\mathbb{C}, +, \cdot)$.
- The induced measure on $\mathbb{C}^1$ is

$$\begin{aligned} \mu(X) &= |X| \quad \text{if } X \text{ finite} \\ \mu(\mathbb{C}) &= 1 \quad (\text{because } \mathbb{C} \text{ is a 2-cell}) \\ \mu(\mathbb{C} \setminus X) &= 1 - |X| \quad \text{if } X \text{ is finite} \end{aligned}$$

- This measure is definable, so $\chi_{\mathbb{C}}$ is definable.
- If $K \models ACF_0$, then the same μ is a definable measure, and

$$d\mu(x)\, d\mu(y) = d\mu(y)\, d\mu(x)$$

because $K \equiv \mathbb{C}$.

Example: ACF_0

- The field $\mathbb{C}$ can be interpreted in the field $\mathbb{R}$. The o-minimal Euler characteristic $\chi_{\mathbb{R}}$ on $(\mathbb{R}, +, \cdot)$ yields a strong Euler characteristic $\chi_{\mathbb{C}}$ on $(\mathbb{C}, +, \cdot)$.
- The induced measure on $\mathbb{C}^1$ is

$$\begin{aligned} \mu(X) &= |X| \quad \text{if } X \text{ finite} \\ \mu(\mathbb{C}) &= 1 \quad (\text{because } \mathbb{C} \text{ is a 2-cell}) \\ \mu(\mathbb{C} \setminus X) &= 1 - |X| \quad \text{if } X \text{ is finite} \end{aligned}$$

- This measure is definable, so $\chi_{\mathbb{C}}$ is definable.
- If $K \models ACF_0$, then the same μ is a definable measure, and

$$d\mu(x)\, d\mu(y) = d\mu(y)\, d\mu(x)$$

  because $K \equiv \mathbb{C}$.
- Therefore, μ gives rise to a $\mathbb{Z}$-valued definable strong Euler characteristic on K.

Semiring-valued Euler characteristics

A (commutative) *semiring* is a structure $(R, +, 0, 1, \cdot)$ such that

- $(R, 0, +)$ is a commutative monoid.
- $(R, 1, \cdot)$ is a commutative monoid.
- For every $a \in R$, the map $x \mapsto x \cdot a$ is a monoid endomorphism of $(R, 0, +)$.

Semiring-valued Euler characteristics

A (commutative) *semiring* is a structure $(R, +, 0, 1, \cdot)$ such that

- $(R, 0, +)$ is a commutative monoid.
- $(R, 1, \cdot)$ is a commutative monoid.
- For every $a \in R$, the map $x \mapsto x \cdot a$ is a monoid endomorphism of $(R, 0, +)$.

Example

The *tropical semiring* is $R = \{0, 1, \omega, \omega^2, \omega^3, \ldots\}$, where, for example

$$\omega^2 + \omega^3 = \omega^3$$
$$\omega^2 \cdot \omega^3 = \omega^5$$

Semiring-valued Euler characteristics

A (commutative) *semiring* is a structure $(R, +, 0, 1, \cdot)$ such that

- $(R, 0, +)$ is a commutative monoid.
- $(R, 1, \cdot)$ is a commutative monoid.
- For every $a \in R$, the map $x \mapsto x \cdot a$ is a monoid endomorphism of $(R, 0, +)$.

Example

The *tropical semiring* is $R = \{0, 1, \omega, \omega^2, \omega^3, \ldots\}$, where, for example

$$\omega^2 + \omega^3 = \omega^3$$
$$\omega^2 \cdot \omega^3 = \omega^5$$

Remark

Semiring-valued Euler characteristics work just as well as ring-valued Euler characteristics!

Dimension as an Euler characteristic

Definition

A theory T is *geometric* if

1. $\exists^\infty$ can be eliminated
2. acl satisfies Steinitz exchange:

$$a \in \mathrm{acl}(\{b\} \cup S) \wedge a \notin \mathrm{acl}(S) \implies b \in \mathrm{acl}(\{a\} \cup S)$$

Dimension as an Euler characteristic

Definition

A theory T is *geometric* if

1. $\exists^\infty$ can be eliminated
2. acl satisfies Steinitz exchange:

$$a \in \mathrm{acl}(\{b\} \cup S) \wedge a \notin \mathrm{acl}(S) \implies b \in \mathrm{acl}(\{a\} \cup S)$$

Examples include ACVF, $\mathbb{Q}_p$, pseudo-finite fields, strongly minimal theories, o-minimal theories. . .

Dimension as an Euler characteristic

Remark

If M is a model of a geometric theory, there is an Euler characteristic taking values in the tropical semiring:

$$\chi(X) = \omega^{\dim(X)}$$

Dimension as an Euler characteristic

Remark

If M is a model of a geometric theory, there is an Euler characteristic taking values in the tropical semiring:

$$\chi(X) = \omega^{\dim(X)}$$

The *existence* of the dimension theory can be proven by checking definability and Fubini for the measure

$$\mu(X) = \begin{cases} 0 & \text{if } X = \emptyset \\ \omega & \text{if } |X| \geq \aleph_0 \\ 1 & \text{otherwise} \end{cases}$$

Zilber degree as an Euler characteristic

Definition

A strongly minimal structure M is *unimodular* if whenever $a \in \mathrm{acl}(Sb)$ and $b \in \mathrm{acl}(Sa)$ but $a, b \notin \mathrm{acl}(S)$,

$$\mathrm{mult}(a/Sb) = \mathrm{mult}(b/Sa)$$

where the multiplicity of an algebraic type is the number of realizations.

Zilber degree as an Euler characteristic

Definition

A strongly minimal structure M is *unimodular* if whenever $a \in \mathrm{acl}(Sb)$ and $b \in \mathrm{acl}(Sa)$ but $a, b \notin \mathrm{acl}(S)$,

$$\mathrm{mult}(a/Sb) = \mathrm{mult}(b/Sa)$$

where the multiplicity of an algebraic type is the number of realizations.

Pseudo-finite and $\aleph_0$-categorical strongly minimal structures are known to be unimodular.

Zilber degree as an Euler characteristic

Remark

In a unimodular strongly minimal structure, there is a semiring-valued Euler characteristic given by

$$\chi(X) = Z(X)\omega^{\dim(X)}$$

where $\dim(X)$ *is the Morley rank of* X *and* $Z(X)$ *is the Zilber degree.*

Zilber degree as an Euler characteristic

Remark

In a unimodular strongly minimal structure, there is a semiring-valued Euler characteristic given by

$$\chi(X) = Z(X)\omega^{\dim(X)}$$

where $\dim(X)$ *is the Morley rank of* X *and* $Z(X)$ *is the Zilber degree.*

The relevant semiring can be constructed by taking the free semiring $\mathbb{N}[\omega]$ on one generator ω, modulo the equivalence relation of having the same leading term.

Zilber degree as an Euler characteristic

Remark

In a unimodular strongly minimal structure, there is a semiring-valued Euler characteristic given by

$$\chi(X) = Z(X)\omega^{\dim(X)}$$

where $\dim(X)$ *is the Morley rank of* X *and* $Z(X)$ *is the Zilber degree.*

The relevant semiring can be constructed by taking the free semiring $\mathbb{N}[\omega]$ on one generator ω, modulo the equivalence relation of having the same leading term. For example,

$$2\omega^3 \cdot 3\omega^2 = 6\omega^5$$
$$2\omega^3 + 3\omega^2 = 2\omega^3$$
$$2\omega^3 + 3\omega^3 = 5\omega^3$$

Pseudo-finite counting

Let M be an ultraproduct of finite structures. Let $\mathbb{Z}^*$ be the corresponding ultrapower of $\mathbb{Z}$.

Pseudo-finite counting

Let M be an ultraproduct of finite structures. Let $\mathbb{Z}^*$ be the corresponding ultrapower of $\mathbb{Z}$.
Every definable set X in M has a nonstandard "size" $|X| \in \mathbb{Z}^*$. The map $X \mapsto |X|$ is a $\mathbb{Z}^*$-valued strong Euler characteristic.

Pseudo-finite counting

Let M be an ultraproduct of finite structures. Let $\mathbb{Z}^*$ be the corresponding ultrapower of $\mathbb{Z}$.
Every definable set X in M has a nonstandard "size" $|X| \in \mathbb{Z}^*$. The map $X \mapsto |X|$ is a $\mathbb{Z}^*$-valued strong Euler characteristic.
If $N > 0$, then $\mathbb{Z}^*/N \cong \mathbb{Z}/N$. The map

$$\chi_N(X) := |X| + N\mathbb{Z}^* \in \mathbb{Z}^*/N \cong \mathbb{Z}/N$$

defines a $\mathbb{Z}/N$-valued Euler characteristic χ_N.

Pseudo-finite counting

Remark

The Euler characteristic χ_N is strong.

Pseudo-finite counting

Remark

The Euler characteristic χ_N is strong.

Proof.

1. It is a general fact that if χ is a definable and strong R-valued Euler characteristic, and $\phi : R \to S$ is a ring homomorphism, then $\phi \circ \chi$ is a definable and strong R-valued Euler characteristic.

Pseudo-finite counting

Remark

The Euler characteristic χ_N is strong.

Proof.

1. It is a general fact that if χ is a definable and strong R-valued Euler characteristic, and $\phi : R \to S$ is a ring homomorphism, then $\phi \circ \chi$ is a definable and strong R-valued Euler characteristic.
2. Therefore, χ_N is strong (and definable) on finite structures.

Pseudo-finite counting

Remark

The Euler characteristic χ_N is strong.

Proof.

1. It is a general fact that if χ is a definable and strong R-valued Euler characteristic, and $\phi : R \to S$ is a ring homomorphism, then $\phi \circ \chi$ is a definable and strong R-valued Euler characteristic.
2. Therefore, χ_N is strong (and definable) on finite structures.
3. It is a general fact that an ultraproduct of strong Euler characteristics is strong.

Pseudo-finite counting

Remark

The Euler characteristic χ_N is strong.

Proof.

1. It is a general fact that if χ is a definable and strong R-valued Euler characteristic, and $\phi : R \to S$ is a ring homomorphism, then $\phi \circ \chi$ is a definable and strong R-valued Euler characteristic.
2. Therefore, χ_N is strong (and definable) on finite structures.
3. It is a general fact that an ultraproduct of strong Euler characteristics is strong.
4. Therefore, χ_N is strong on ultraproducts of finite structures.

$\square$

$\hat{\mathbb{Z}}$-valued Euler characteristics

If N divides M, then

$$\chi_N(X) \equiv \chi_M(X) \mod N$$

$\hat{\mathbb{Z}}$-valued Euler characteristics

If N divides M, then

$$\chi_N(X) \equiv \chi_M(X) \mod N$$

Consequently, the χ_N for all N can be assembled into a single $\hat{\chi}(X)$ taking values in the ring

$$\hat{\mathbb{Z}} := \varprojlim_N \mathbb{Z}/N.$$

$\hat{\mathbb{Z}}$-valued Euler characteristics

If N divides M, then

$$\chi_N(X) \equiv \chi_M(X) \mod N$$

Consequently, the χ_N for all N can be assembled into a single $\hat{\chi}(X)$ taking values in the ring

$$\hat{\mathbb{Z}} := \varprojlim_N \mathbb{Z}/N.$$

Let π_N be the natural map $\hat{\mathbb{Z}} \to \mathbb{Z}/N$.

Definition (abuse of notation)

A $\hat{\mathbb{Z}}$-valued Euler characteristic $\hat{\chi}$ is *strong* if $\pi_N \circ \hat{\chi}$ is strong for every N.
A $\hat{\mathbb{Z}}$-valued Euler characteristic $\hat{\chi}$ is *definable* if $\pi_N \circ \hat{\chi}$ is definable for every N.

$\hat{\mathbb{Z}}$-valued Euler characteristics and integration

Remark

For definable and strong $\hat{\mathbb{Z}}$-valued $\hat{\chi}$, the integration formula still holds

$$\hat{\chi}\left(\coprod_{x\in X} D_x\right) = \int_{x\in X} \hat{\chi}(D_x)\ d\hat{\chi}(x),$$

but we need to use p-adic integration.

Where we are headed: Goal #1

If M is an ultraproduct of finite structures, then there is canonically a $\hat{\mathbb{Z}}$-valued Euler characteristic $\hat{\chi}$ that is strong, but not necessarily definable.

Where we are headed: Goal #1

If M is an ultraproduct of finite structures, then there is canonically a $\hat{\mathbb{Z}}$-valued Euler characteristic $\hat{\chi}$ that is strong, but not necessarily definable.

Theorem (W. Johnson)

If K is an ultraproduct of finite fields, then this $\hat{\chi}$ is definable.

$\hat{\chi}$ needn't be definable

If $(M, <)$ is an ultraproduct of finite total orders, then χ_2 isn't definable:

$\hat{\chi}$ needn't be definable

If $(M, <)$ is an ultraproduct of finite total orders, then χ_2 isn't definable:

- Let $-\infty$ be the smallest element of M, and let I_x be the closed interval $[-\infty, x]$.

$\hat{\chi}$ needn't be definable

If $(M, <)$ is an ultraproduct of finite total orders, then χ_2 isn't definable:

- Let $-\infty$ be the smallest element of M, and let I_x be the closed interval $[-\infty, x]$.
- $\chi_2(I_x)$ needs to alternate between 0 and 1, because

$$\chi_2(I_x) = \chi_2(I_{x-1}) + \chi_2(\{x\}) = \chi_2(I_{x-1}) + 1$$

$\hat{\chi}$ needn't be definable

If $(M, <)$ is an ultraproduct of finite total orders, then χ_2 isn't definable:

- Let $-\infty$ be the smallest element of M, and let I_x be the closed interval $[-\infty, x]$.
- $\chi_2(I_x)$ needs to alternate between 0 and 1, because

$$\chi_2(I_x) = \chi_2(I_{x-1}) + \chi_2(\{x\}) = \chi_2(I_{x-1}) + 1$$

- No definable set has this property (use quantifier elimination or consider automorphisms of $\aleph_0$-saturated models).

The theory of finite fields

Let $\mathcal{L}_{rings}$ be the first-order language of rings.

Definition

The (elementary) theory of finite fields is the set T_{ff} of sentences $\varphi \in \mathcal{L}_{rings}$ such that $K \models \varphi$ for every finite field K.

The theory of finite fields

Let $\mathcal{L}_{rings}$ be the first-order language of rings.

Definition

The (elementary) theory of finite fields is the set T_{ff} of sentences $\varphi \in \mathcal{L}_{rings}$ such that $K \models \varphi$ for every finite field K.

Example

- The map $x \mapsto x(x-1)$ isn't injective on any field. On finite fields, it can't be surjective, so

$$\exists y \forall x : y \neq x(x-1)$$

  is in T_{ff}.
- The sentence $\forall x \exists y : x = y^2$ (every element is a square) is true in some finite fields, and false in others. This sentence is not in T_{ff}.

The theory of finite fields

Let $\mathcal{L}_{rings}$ be the first-order language of rings.

Definition

The (elementary) theory of finite fields is the set T_{ff} of sentences $\varphi \in \mathcal{L}_{rings}$ such that $K \models \varphi$ for every finite field K.

Example

- The map $x \mapsto x(x-1)$ isn't injective on any field. On finite fields, it can't be surjective, so

$$\exists y \forall x : y \neq x(x-1)$$

  is in T_{ff}.
- The sentence $\forall x \exists y : x = y^2$ (every element is a square) is true in some finite fields, and false in others. This sentence is not in T_{ff}.

Note: T_{ff} is not complete.

Pseudo-finite fields

Let K be a model of T_{ff}. Either

- K is finite, OR...

Pseudo-finite fields

Let K be a model of T_{ff}. Either

- K is finite, OR...
- K is an infinite field, elementarily equivalent to an ultraproduct of finite fields.

Pseudo-finite fields

Let K be a model of T_{ff}. Either

- K is finite, OR...
- K is an infinite field, elementarily equivalent to an ultraproduct of finite fields.

Fields of the second type are called *pseudo-finite fields*.

Pseudo-finite fields

Let K be a model of T_{ff}. Either

- K is finite, OR...
- K is an infinite field, elementarily equivalent to an ultraproduct of finite fields.

Fields of the second type are called *pseudo-finite fields*.
The following transfer principles hold:

- If $\varphi \in \mathcal{L}_{rings}$ holds in infinitely many finite fields, then φ holds in some pseudo-finite fields.
- If $\varphi \in \mathcal{L}_{rings}$ holds in almost all finite fields, then φ holds in all pseudo-finite fields.

James Ax's work

In *The Elementary Theory of Finite Fields*, Ax wrote down a computable set of axioms and proved that these axioms generate T_{ff}.

James Ax's work

In *The Elementary Theory of Finite Fields*, Ax wrote down a computable set of axioms and proved that these axioms generate T_{ff}.
From this, Ax proved:

Corollary

T_{ff} is computable: there is an algorithm which takes $\varphi \in \mathcal{L}_{rings}$ and determines whether φ holds in all finite fields.

Proof.

James Ax's work

In *The Elementary Theory of Finite Fields*, Ax wrote down a computable set of axioms and proved that these axioms generate T_{ff}.
From this, Ax proved:

Corollary

T_{ff} is computable: there is an algorithm which takes $\varphi \in \mathcal{L}_{rings}$ and determines whether φ holds in all finite fields.

Proof.

T_{ff} is computably enumerable (CE) because it is generated by a computable set of axioms.
T_{ff} is co-CE because one can inspect individual finite fields one by one to look for counterexamples. □

Ingredients in Ax's proof

Let T_{Ax} denote Ax's axioms for the theory of finite fields. The proof that $T_{Ax} \vdash T_{ff}$ uses three ingredients:

1. Weil's Riemann Hypothesis for function fields: used to show that finite and pseudo-finite fields satisfy T_{Ax}.

Ingredients in Ax's proof

Let T_{Ax} denote Ax's axioms for the theory of finite fields. The proof that $T_{Ax} \vdash T_{ff}$ uses three ingredients:

1. Weil's Riemann Hypothesis for function fields: used to show that finite and pseudo-finite fields satisfy T_{Ax}.
2. A model theoretic argument: used to characterize elementary equivalence between fields satisfying T_{Ax}.

Ingredients in Ax's proof

Let T_{Ax} denote Ax's axioms for the theory of finite fields. The proof that $T_{Ax} \vdash T_{ff}$ uses three ingredients:

1. Weil's Riemann Hypothesis for function fields: used to show that finite and pseudo-finite fields satisfy T_{Ax}.
2. A model theoretic argument: used to characterize elementary equivalence between fields satisfying T_{Ax}.
3. The Chebotarev density theorem: used to show that every model of T_{Ax} is elementarily equivalent to a pseudo-finite field.

Ingredients in Ax's proof

Let T_{Ax} denote Ax's axioms for the theory of finite fields. The proof that $T_{Ax} \vdash T_{ff}$ uses three ingredients:

1. Weil's Riemann Hypothesis for function fields: used to show that finite and pseudo-finite fields satisfy T_{Ax}.
2. A model theoretic argument: used to characterize elementary equivalence between fields satisfying T_{Ax}.
3. The Chebotarev density theorem: used to show that every model of T_{Ax} is elementarily equivalent to a pseudo-finite field.

1 and 3 are moderately deep results from algebraic geometry and number theory, respectively.

Ax's characterization of pseudo-finite fields

From the axioms that generate T_{ff}, Ax gave the following algebraic characterization of pseudo-finite fields:

Corollary

A field K is pseudo-finite exactly if it satisfies the following algebraic conditions:

1. *K is perfect.*

Ax's characterization of pseudo-finite fields

From the axioms that generate T_{ff}, Ax gave the following algebraic characterization of pseudo-finite fields:

Corollary

A field K is pseudo-finite exactly if it satisfies the following algebraic conditions:

1. *K is perfect.*
2. *For every $n \geq 1$, there is exactly one algebraic field extension of degree n.*

Ax's characterization of pseudo-finite fields

From the axioms that generate T_{ff}, Ax gave the following algebraic characterization of pseudo-finite fields:

Corollary

A field K is pseudo-finite exactly if it satisfies the following algebraic conditions:

1. *K is perfect.*
2. *For every $n \geq 1$, there is exactly one algebraic field extension of degree n.*
3. *K is pseudo-algebraically closed (PAC)*

Two definitions of PAC

For model theorists:

Definition

A field K is *perfect and PAC* if, when we embed K into a monster model $\mathbb{M}$ of ACF, every stationary type over K is finitely satisfiable in K.

Two definitions of PAC

For model theorists:

Definition

A field K is *perfect and PAC* if, when we embed K into a monster model $\mathbb{M}$ of ACF, every stationary type over K is finitely satisfiable in K.

Remark

It suffices to check the stationary types of Morley rank 1.

Two definitions of PAC

For model theorists:

Definition

A field K is *perfect and PAC* if, when we embed K into a monster model $\mathbb{M}$ of ACF, every stationary type over K is finitely satisfiable in K.

Remark

It suffices to check the stationary types of Morley rank 1.

For algebraic geometers:

Definition

A field K is *PAC* if for every geometrically integral finite type K-scheme V, the set $V(K)$ of K-rational points is non-empty.

Two definitions of PAC

For model theorists:

Definition

A field K is *perfect and PAC* if, when we embed K into a monster model $\mathbb{M}$ of ACF, every stationary type over K is finitely satisfiable in K.

Remark

It suffices to check the stationary types of Morley rank 1.

For algebraic geometers:

Definition

A field K is *PAC* if for every geometrically integral finite type K-scheme V, the set $V(K)$ of K-rational points is non-empty.

Remark

It suffices to check 1-dimensional V (i.e., algebraic curves).

Surprisingly pseudo-finite fields

- If σ is chosen randomly in $\mathrm{Gal}(\mathbb{Q}) := \mathrm{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$ with respect to Haar measure, then the field of elements fixed by σ is pseudo-finite with probability 1.

Surprisingly pseudo-finite fields

- If σ is chosen randomly in $\mathrm{Gal}(\mathbb{Q}) := \mathrm{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$ with respect to Haar measure, then the field of elements fixed by σ is pseudo-finite with probability 1.
- If $(Z, +, \cdot)$ is a model of Peano arithmetic, and (p) is a principal prime ideal, then $Z/(p)$ is a pseudo-finite field.

Surprisingly pseudo-finite fields

- If σ is chosen randomly in $\mathrm{Gal}(\mathbb{Q}) := \mathrm{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$ with respect to Haar measure, then the field of elements fixed by σ is pseudo-finite with probability 1.
- If $(Z, +, \cdot)$ is a model of Peano arithmetic, and (p) is a principal prime ideal, then $Z/(p)$ is a pseudo-finite field.
- Most subfields of $\mathbb{F}_p^{alg}$ are pseudo-finite. For example, the subfield generated by

$$\mathbb{F}_{p^2} \cup \mathbb{F}_{p^3} \cup \mathbb{F}_{p^5} \cup \cdots$$

is pseudo-finite.

Surprisingly pseudo-finite fields

- If σ is chosen randomly in $\mathrm{Gal}(\mathbb{Q}) := \mathrm{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$ with respect to Haar measure, then the field of elements fixed by σ is pseudo-finite with probability 1.
- If $(Z, +, \cdot)$ is a model of Peano arithmetic, and (p) is a principal prime ideal, then $Z/(p)$ is a pseudo-finite field.
- Most subfields of $\mathbb{F}_p^{alg}$ are pseudo-finite. For example, the subfield generated by
$$\mathbb{F}_{p^2} \cup \mathbb{F}_{p^3} \cup \mathbb{F}_{p^5} \cup \cdots$$
is pseudo-finite.
- If $\mathbb{A}_{\mathbb{Q}}$ is the number theorist's ring of adeles and $\mathfrak{m}$ is any maximal ideal other than the "obvious" ones, then quotient $\mathbb{A}_{\mathbb{Q}}/\mathfrak{m}$ is a pseudo-finite field.

Where we are headed: Goal #2

Theorem (W. Johnson)

Every pseudo-finite field K admits a definable and strong $\hat{\mathbb{Z}}$-valued Euler characteristic $\hat{\chi}$.

Where we are headed: Goal #2

Theorem (W. Johnson)

Every pseudo-finite field K admits a definable and strong $\hat{\mathbb{Z}}$-valued Euler characteristic $\hat{\chi}$.

This $\hat{\chi}$ isn't uniquely determined, unfortunately. It depends on a choice of an element $\sigma \in \mathrm{Aut}(K^{alg}/K)$.

Strong Euler characteristics needn't exist

Proposition

If K is an algebraically closed field of positive characteristic, then K admits no R-valued strong Euler characteristics for any nontrivial ring R.

The proof involves examining a couple explicit maps from K to K and proving that $1 = 0$ in R.

Where we are headed: Goal #3

Let $\mathcal{L}^+_{rings}$ be the first-order language of rings expanded with new quantifiers $\mu^n_k x$, where

$$\mu^n_k x : P(x) \iff$$

the number of x such that $P(x)$ holds is congruent to k mod n

Where we are headed: Goal #3

Let $\mathcal{L}^+_{rings}$ be the first-order language of rings expanded with new quantifiers $\mu^n_k x$, where

$$\mu^n_k x : P(x) \iff$$

$$\text{the number of } x \text{ such that } P(x) \text{ holds is congruent to } k \text{ mod } n$$

For example,

- $\mu^2_0 x$ means "there exists an even number of x such that..."
- $\mu^2_1 x$ means "there exists an odd number of x such that..."

Where we are headed: Goal #3

Let $\mathcal{L}^+_{rings}$ be the first-order language of rings expanded with new quantifiers $\mu^n_k x$, where

$$\mu^n_k x : P(x) \iff$$

$$\text{the number of } x \text{ such that } P(x) \text{ holds is congruent to } k \text{ mod } n$$

For example,

- $\mu^2_0 x$ means "there exists an even number of x such that..."
- $\mu^2_1 x$ means "there exists an odd number of x such that..."

These quantifiers make sense in any finite structure.

Where we are headed: Goal #3

Let $\mathcal{L}^+_{rings}$ be the first-order language of rings expanded with new quantifiers $\mu^n_k x$, where

$$\mu^n_k x : P(x) \iff$$

the number of x such that $P(x)$ holds is congruent to k mod n

For example,

- $\mu^2_0 x$ means "there exists an even number of x such that..."
- $\mu^2_1 x$ means "there exists an odd number of x such that..."

These quantifiers make sense in any finite structure.

Main Theorem (W. Johnson)

There is an algorithm which takes a sentence $\varphi \in \mathcal{L}^+_{rings}$ and determines whether φ holds in every finite field.

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares)

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares) is equivalent to

$$(1 + 1 \neq 0) \wedge (\exists x : x^2 = -1)$$

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares) is equivalent to

$$(1 + 1 \neq 0) \wedge (\exists x : x^2 = -1)$$

Example

$\mu_1^5 x : x = x$ (size of the field is congruent to 1 mod 5)

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares) is equivalent to

$$(1 + 1 \neq 0) \wedge (\exists x : x^2 = -1)$$

Example

$\mu_1^5 x : x = x$ (size of the field is congruent to 1 mod 5) is equivalent to

$$(5 \neq 0) \wedge (\exists x : x \neq 1 \wedge x^5 = 1)$$

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares) is equivalent to

$$(1 + 1 \neq 0) \wedge (\exists x : x^2 = -1)$$

Example

$\mu_1^5 x : x = x$ (size of the field is congruent to 1 mod 5) is equivalent to

$$(5 \neq 0) \wedge (\exists x : x \neq 1 \wedge x^5 = 1)$$

Non-example

$\mu_3^5 x : x = x$ *(size of the field is congruent to 3 mod 5)*

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares) is equivalent to

$$(1 + 1 \neq 0) \wedge (\exists x : x^2 = -1)$$

Example

$\mu_1^5 x : x = x$ (size of the field is congruent to 1 mod 5) is equivalent to

$$(5 \neq 0) \wedge (\exists x : x \neq 1 \wedge x^5 = 1)$$

Non-example

$\mu_3^5 x : x = x$ *(size of the field is congruent to 3 mod 5) is not equivalent to any sentence in* $\mathcal{L}_{rings}$.

Eliminating the μ_k^n, or not

Reduce to Ax's decidability result by eliminating the μ_k^n quantifiers.

Example

$\mu_1^2 x \exists y : x = y^2$ (there are an odd number of perfect squares) is equivalent to

$$(1 + 1 \neq 0) \wedge (\exists x : x^2 = -1)$$

Example

$\mu_1^5 x : x = x$ (size of the field is congruent to 1 mod 5) is equivalent to

$$(5 \neq 0) \wedge (\exists x : x \neq 1 \wedge x^5 = 1)$$

Non-example

$\mu_3^5 x : x = x$ *(size of the field is congruent to 3 mod 5) is not equivalent to any sentence in* $\mathcal{L}_{rings}$.

We need an expanded structure.

Difference fields

Definition

A *difference field* (K, σ) is a field K and an endomorphism $\sigma : K \to K$.

Difference fields

Definition

A *difference field* (K, σ) is a field K and an endomorphism $\sigma : K \to K$.

Why?

“Difference” fields

Differentiation satisfies these rules:

$$\partial(f \cdot g) = f\partial g + g\partial f$$
$$\partial(f + g) = \partial f + \partial g$$

"Difference" fields

Differentiation satisfies these rules:

$$\partial(f \cdot g) = f\partial g + g\partial f$$
$$\partial(f + g) = \partial f + \partial g$$

Definition

A *differential field* is a field K and an operator $\partial : K \to K$ satisfying the product and sum rules for differentiation.

"Difference" fields

The difference operator $(\delta f)(x) = f(x+1) - f(x)$ satisfies analogous rules:

$$\delta(f \cdot g) = f\delta g + g\delta f + (\delta f)(\delta g)$$
$$\delta(f + g) = \delta f + \delta g$$

"Difference" fields

The difference operator $(\delta f)(x) = f(x+1) - f(x)$ satisfies analogous rules:

$$\delta(f \cdot g) = f\delta g + g\delta f + (\delta f)(\delta g)$$
$$\delta(f + g) = \delta f + \delta g$$

A *difference field* should be a field K and an operator $\delta : K \to K$ satisfying the product and sum rules for differences.

Remark

The axioms of δ are an obfuscated way of saying that $x \mapsto x + \delta x$ is a field endomorphism.

The theory of Frobenius difference fields

In characteristic p, the following identities hold:

$$(x+y)^p = x^p + y^p$$
$$(xy)^p = x^p \cdot y^p$$

In particular, $x \mapsto x^p$ is a field endomorphism. More generally, $x \mapsto x^q$ is a field endomorphism for any power $q = p^k$.

The theory of Frobenius difference fields

In characteristic p, the following identities hold:

$$(x+y)^p = x^p + y^p$$
$$(xy)^p = x^p \cdot y^p$$

In particular, $x \mapsto x^p$ is a field endomorphism. More generally, $x \mapsto x^q$ is a field endomorphism for any power $q = p^k$.

Definition

For $q = p^k$, the qth Frobenius difference field is $\mathbb{F}_p^{alg}$ with σ given by $\sigma(x) = x^q$.

The theory of Frobenius difference fields

In characteristic p, the following identities hold:

$$(x+y)^p = x^p + y^p$$
$$(xy)^p = x^p \cdot y^p$$

In particular, $x \mapsto x^p$ is a field endomorphism. More generally, $x \mapsto x^q$ is a field endomorphism for any power $q = p^k$.

Definition

For $q = p^k$, the qth Frobenius difference field is $\mathbb{F}_p^{alg}$ with σ given by $\sigma(x) = x^q$.

Definition

The *(elementary) theory of Frobenius difference fields* is the set T_{frob} of φ holding in every Frobenius difference field.

Hrushovski's work

In *The Elementary Theory of the Frobenius Automorphism* (149 pages), Hrushovski wrote down a computable set of axioms and (probably) proved that these axioms generate T_{frob}.

Hrushovski's work

In *The Elementary Theory of the Frobenius Automorphism* (149 pages), Hrushovski wrote down a computable set of axioms and (probably) proved that these axioms generate T_{frob}.
From this, Hrushovski proved:

Corollary

T_{frob} is computable: there is an algorithm which takes φ in the language of difference fields and determines whether φ holds in all finite fields.

Hrushovski's work

In *The Elementary Theory of the Frobenius Automorphism* (149 pages), Hrushovski wrote down a computable set of axioms and (probably) proved that these axioms generate T_{frob}.
From this, Hrushovski proved:

Corollary

T_{frob} is computable: there is an algorithm which takes φ in the language of difference fields and determines whether φ holds in all finite fields.

Corollary

The models of T_{frob} are the Frobenius difference fields plus the models of ACFA.

ACFA

Proposition

The theory of difference fields has a model companion, called ACFA.

ACFA

Proposition

The theory of difference fields has a model companion, called ACFA.

This means:

- The models of ACFA are exactly the existentially closed difference fields.
- Every difference field can be embedded into a model of ACFA.
- If $K \leq L$ is an inclusion of models of ACFA, then $K \preceq L$.

ACFA and pseudo-finite fields

Definition

If (K, σ) is a difference field, the *fixed field* K^σ is

$$K^\sigma := \{x \in K \mid \sigma(x) = x\}$$

ACFA and pseudo-finite fields

Definition

If (K, σ) is a difference field, the *fixed field* K^σ is

$$K^\sigma := \{x \in K \mid \sigma(x) = x\}$$

ACFA turns out to be closely related to pseudo-finite fields.

Proposition

- *If* $(K, \sigma) \models$ *ACFA, then* K^σ *is pseudo-finite.*

ACFA and pseudo-finite fields

Definition

If (K, σ) is a difference field, the *fixed field* K^σ is

$$K^\sigma := \{x \in K \mid \sigma(x) = x\}$$

ACFA turns out to be closely related to pseudo-finite fields.

Proposition

- *If $(K, \sigma) \models ACFA$, then K^σ is pseudo-finite.*
- *Up to elementary equivalence, all pseudo-finite fields arise this way.*

ACFA and pseudo-finite fields

Definition

If (K, σ) is a difference field, the *fixed field* K^σ is

$$K^\sigma := \{x \in K \mid \sigma(x) = x\}$$

ACFA turns out to be closely related to pseudo-finite fields.

Proposition

- *If* $(K, \sigma) \models$ *ACFA, then* K^σ *is pseudo-finite.*
- *Up to elementary equivalence, all pseudo-finite fields arise this way.*

Likewise,

Proposition

- *If* (K, σ) *is a Frobenius difference field, then* K^σ *is a finite field.*
- *All finite fields arise this way.*

ACFA and pseudo-finite fields

Definition

If (K, σ) is a difference field, the *fixed field* K^σ is

$$K^\sigma := \{x \in K \mid \sigma(x) = x\}$$

ACFA turns out to be closely related to pseudo-finite fields.

Proposition

- *If $(K, \sigma) \models$ ACFA, then K^σ is pseudo-finite.*
- *Up to elementary equivalence, all pseudo-finite fields arise this way.*

Likewise,

Proposition

- *If (K, σ) is a Frobenius difference field, then K^σ is a finite field.*
- *All finite fields arise this way.*

Consequently, Hrushovski's work *generalizes* Ax's work.

Periodic difference fields

We need a (much easier) variant of Hrushovski's work on ACFA and Frobenius difference fields.

Periodic difference fields

We need a (much easier) variant of Hrushovski's work on ACFA and Frobenius difference fields.

Definition

A difference field (K, σ) is *periodic* if every $x \in K$ has finite orbit under σ, i.e., for every $x \in K$ there is some $n \geq 1$ such that $\sigma^n x = x$.

Periodic difference fields

We need a (much easier) variant of Hrushovski's work on ACFA and Frobenius difference fields.

Definition

A difference field (K, σ) is *periodic* if every $x \in K$ has finite orbit under σ, i.e., for every $x \in K$ there is some $n \geq 1$ such that $\sigma^n x = x$.

Example

- If $\sigma \in \mathrm{Aut}(\mathbb{Q}^{alg}/\mathbb{Q})$, then $(\mathbb{Q}^{alg}, \sigma)$ is periodic.
- Frobenius difference fields are periodic.

The multi-sorted language

Problem: periodic difference fields are not an elementary class.

The multi-sorted language

Problem: periodic difference fields are not an elementary class.
Solution: work in a multi-sorted structure $(K_1, K_2, K_3, \ldots)$, where

- K_i is the fixed field of σ^i.
- Each K_i has the structure of a difference field.
- There are inclusion maps $\iota_{n,m} : K_n \to K_m$ whenever $n|m$.

The multi-sorted language

Problem: periodic difference fields are not an elementary class.
Solution: work in a multi-sorted structure $(K_1, K_2, K_3, \ldots)$, where

- K_i is the fixed field of σ^i.
- Each K_i has the structure of a difference field.
- There are inclusion maps $\iota_{n,m} : K_n \to K_m$ whenever $n|m$.

Remark

- *In this language, periodic difference fields are now an elementary class.*

The multi-sorted language

Problem: periodic difference fields are not an elementary class.
Solution: work in a multi-sorted structure $(K_1, K_2, K_3, \ldots)$, where

- K_i is the fixed field of σ^i.
- Each K_i has the structure of a difference field.
- There are inclusion maps $\iota_{n,m} : K_n \to K_m$ whenever $n|m$.

Remark

- *In this language, periodic difference fields are now an elementary class.*
- *In this language, Frobenius difference fields are essentially finite: every definable set is finite.*

The multi-sorted language

Problem: periodic difference fields are not an elementary class.
Solution: work in a multi-sorted structure $(K_1, K_2, K_3, \ldots)$, where

- K_i is the fixed field of σ^i.
- Each K_i has the structure of a difference field.
- There are inclusion maps $\iota_{n,m} : K_n \to K_m$ whenever $n|m$.

Remark

- *In this language, periodic difference fields are now an elementary class.*
- *In this language, Frobenius difference fields are essentially finite: every definable set is finite.*

To highlight the change of language, we will refer to periodic difference fields as *periodic fields* henceforth.

Hrushovski's work: miniature version

Let T_{frob} denote the (elementary) theory of Frobenius *periodic* fields.

Hrushovski's work: miniature version

Let T_{frob} denote the (elementary) theory of Frobenius *periodic* fields.

1. There is a computable set of axioms which generate T_{frob}.
2. Therefore, T_{frob} is computable.
3. The theory of periodic difference fields has a model companion "ACPF."
4. The models of T_{frob} are the Frobenius periodic fields and the models of ACPF.

All of these facts were probably initially discovered by Hrushovski, prior to his work on ACFA.

ACPF and pseudo-finite fields

ACPF is closely linked to pseudo-finite fields:

Proposition

1. *If (K, σ) is a model of ACPF, then K^σ is a pseudo-finite field.*
2. *All pseudo-finite fields arise this way.*

ACPF and pseudo-finite fields

ACPF is closely linked to pseudo-finite fields:

Proposition

1. *If (K, σ) is a model of ACPF, then K^σ is a pseudo-finite field.*
2. *All pseudo-finite fields arise this way.*

More explicitly, if K is pseudo-finite then $\mathrm{Aut}(K^{alg}/K)$ is isomorphic to $\hat{\mathbb{Z}}$. If σ is any *topological generator* of $\hat{\mathbb{Z}}$, then (K^{alg}, σ) is a model of ACPF with fixed field K.

ACPF and pseudo-finite fields

ACPF is closely linked to pseudo-finite fields:

Proposition

1. *If (K, σ) is a model of ACPF, then K^σ is a pseudo-finite field.*
2. *All pseudo-finite fields arise this way.*

More explicitly, if K is pseudo-finite then $\mathrm{Aut}(K^{alg}/K)$ is isomorphic to $\hat{\mathbb{Z}}$. If σ is any *topological generator* of $\hat{\mathbb{Z}}$, then (K^{alg}, σ) is a model of ACPF with fixed field K.
Consequently, the previous slide generalizes Ax's work.

Today's main theorem

Main Theorem (W. Johnson)

Let (K, σ) be a model of ACPF. Then there is a canonical $\hat{\mathbb{Z}}$ 0-definable Euler characteristic $\hat{\chi}$ on (K, σ), which agrees with the pseudo-finite $\hat{\chi}$ if (K, σ) is an ultraproduct of Frobenius periodic fields.

Today's main theorem

Main Theorem (W. Johnson)

Let (K, σ) be a model of ACPF. Then there is a canonical $\hat{\mathbb{Z}}$ 0-definable Euler characteristic $\hat{\chi}$ on (K, σ), which agrees with the pseudo-finite $\hat{\chi}$ if (K, σ) is an ultraproduct of Frobenius periodic fields.

Theorem

The μ_k^n quantifiers can be uniformly eliminated across Frobenius periodic fields.

Example: counting points in the fixed field

Earlier, we considered the statement that the number of elements in the field is 3 mod 5:

$$\mu_3^5 x : x = x,$$

Example: counting points in the fixed field

Earlier, we considered the statement that the number of elements in the field is 3 mod 5:

$$\mu_3^5 x : x = x,$$

Interpreting the finite field as a fixed field of a Frobenius periodic field, the analogous statement for Frobenius periodic fields is

$$\mu_3^5 x \in K_1 : x = x.$$

Recall K_1 is the sort for the fixed field K^σ.

Example: counting points in the fixed field

Let $\{1, \omega, \omega^2, \omega^3, \omega^4\}$ be the 5th roots of unity. In the qth Frobenius periodic field, $\sigma(x) = x^q$, and $q = |K_1|$.

Example: counting points in the fixed field

Let $\{1, \omega, \omega^2, \omega^3, \omega^4\}$ be the 5th roots of unity. In the qth Frobenius periodic field, $\sigma(x) = x^q$, and $q = |K_1|$. Therefore:

- If $|K_1| \equiv 1 \mod 5$, then σ fixes ω
- If $|K_1| \equiv 2 \mod 5$, then σ maps

$$\omega \mapsto \omega^2 \mapsto \omega^4 \mapsto \omega^8 = \omega^3 \mapsto \omega^6 = \omega$$

- If $|K_1| \equiv 3 \mod 5$, then σ maps

$$\omega \mapsto \omega^3 \mapsto \omega^9 = \omega^4 \mapsto \omega^{12} = \omega^2 \mapsto \omega^6 = \omega$$

- If $|K_1| \equiv 4 \mod 5$, then σ maps

$$\omega \mapsto \omega^4 \mapsto \omega^{16} = \omega$$

Example: counting points in the fixed field

Consequently,

$$\left(\mu_3^5 x \in K_1 : x = x\right) \iff$$

$$\left(\exists x \in K_4 : x^5 = 1 \wedge x \neq 1 \wedge \sigma(x) = x^3\right)$$

Example: counting points in the fixed field

Consequently,

$$\left(\mu_3^5 x \in K_1 : x = x\right) \iff$$

$$\left(\exists x \in K_4 : x^5 = 1 \wedge x \neq 1 \wedge \sigma(x) = x^3\right)$$

In general, $\chi_N(K_1)$ is determined by examining how σ acts on primitive Nth roots of unity

Example: counting points in the fixed field

Consequently,

$$\left(\mu_3^5 x \in K_1 : x = x\right) \iff$$

$$\left(\exists x \in K_4 : x^5 = 1 \wedge x \neq 1 \wedge \sigma(x) = x^3\right)$$

In general, $\chi_N(K_1)$ is determined by examining how σ acts on primitive Nth roots of unity

...unless N is divisible by the field characteristic.

Bad characteristic

In characteristic 2, what is $|K_1|$ mod 256?

Bad characteristic

In characteristic 2, what is $|K_1|$ mod 256?

- It's 2 if $|K_1| = 2$.
- It's 4 if $|K_1| = 4$.
- ...
- It's 128 if $|K_1| = 128$.
- It's 0 otherwise.

Bad characteristic

In characteristic 2, what is $|K_1|$ mod 256?

- It's 2 if $|K_1| = 2$.
- It's 4 if $|K_1| = 4$.
- ...
- It's 128 if $|K_1| = 128$.
- It's 0 otherwise.

So $\chi_{256}(K_1)$ is definable.

Bad characteristic

In characteristic 2, what is $|K_1|$ mod 256?

- It's 2 if $|K_1| = 2$.
- It's 4 if $|K_1| = 4$.
- ...
- It's 128 if $|K_1| = 128$.
- It's 0 otherwise.

So $\chi_{256}(K_1)$ is definable.
In characteristic 2, what is $|K_1|$ mod N?

Bad characteristic

In characteristic 2, what is $|K_1|$ mod 256?

- It's 2 if $|K_1| = 2$.
- It's 4 if $|K_1| = 4$.
- ...
- It's 128 if $|K_1| = 128$.
- It's 0 otherwise.

So $\chi_{256}(K_1)$ is definable.
In characteristic 2, what is $|K_1|$ mod N?

1. Write $N = N_0 \cdot 2^k$ with N_0 odd.
2. Use the approach above to determine $\chi_{2^k}(K_1)$.
3. Use roots of unity to determine $\chi_{N_0}(K_1)$.

Bad characteristic

In characteristic 2, what is $|K_1|$ mod 256?

- It's 2 if $|K_1| = 2$.
- It's 4 if $|K_1| = 4$.
- ...
- It's 128 if $|K_1| = 128$.
- It's 0 otherwise.

So $\chi_{256}(K_1)$ is definable.
In characteristic 2, what is $|K_1|$ mod N?

1. Write $N = N_0 \cdot 2^k$ with N_0 odd.
2. Use the approach above to determine $\chi_{2^k}(K_1)$.
3. Use roots of unity to determine $\chi_{N_0}(K_1)$.
4. $\mathbb{Z}/N \cong (\mathbb{Z}/N_0) \times (\mathbb{Z}/2^k)$.

Genus 0 curves

Let C be the definable set

$$C = \{(x, y) \in (K_1)^2 \mid x^2 - y^2 = 1\}$$

Question: how can we determine $\chi_N(C)$?

Genus 0 curves

Let C be the definable set

$$C = \{(x, y) \in (K_1)^2 \mid x^2 - y^2 = 1\}$$

Question: how can we determine $\chi_N(C)$?
Answer: C is a *genus 0 curve*, so it is birationally equivalent to the affine line. In this case, the function

$$(x, y) \mapsto x - y$$

puts C into definable bijection with K_1, except for errors at finitely many points.

Genus 0 curves

Let C be the definable set

$$C = \{(x, y) \in (K_1)^2 \mid x^2 - y^2 = 1\}$$

Question: how can we determine $\chi_N(C)$?
Answer: C is a *genus 0 curve*, so it is birationally equivalent to the affine line. In this case, the function

$$(x, y) \mapsto x - y$$

puts C into definable bijection with K_1, except for errors at finitely many points.
Therefore, $\chi_N(C)$ is determined by $\chi_N(K_1)$.

Elliptic curves

If E is an elliptic curve (curve of genus 1), such as

$$y^2 = x^3 - x,$$

then there is a natural group structure on the points of E.

Elliptic curves

If E is an elliptic curve (curve of genus 1), such as

$$y^2 = x^3 - x,$$

then there is a natural group structure on the points of E.

Proposition (Folk theorem)

When N is prime to the characteristic, the mod N value of $|E(\mathbb{F}_q)|$ is determined by how the qth-power Frobenius automorphism acts on the Nth roots of unity and the N-torsion in the group $E(\mathbb{F}_q^{alg})$.

Elliptic curves

If E is an elliptic curve (curve of genus 1), such as

$$y^2 = x^3 - x,$$

then there is a natural group structure on the points of E.

Proposition (Folk theorem)

When N is prime to the characteristic, the mod N value of $|E(\mathbb{F}_q)|$ is determined by how the qth-power Frobenius automorphism acts on the Nth roots of unity and the N-torsion in the group $E(\mathbb{F}_q^{alg})$.

This allows us to prove uniform definability of $\chi_N(E(K_1))$ by breaking into cases according to how σ acts on the N-torsion in $E(K)$ and $K^\times$.

Elliptic curves

If E is an elliptic curve (curve of genus 1), such as

$$y^2 = x^3 - x,$$

then there is a natural group structure on the points of E.

Proposition (Folk theorem)

When N is prime to the characteristic, the mod N value of $|E(\mathbb{F}_q)|$ is determined by how the qth-power Frobenius automorphism acts on the Nth roots of unity and the N-torsion in the group $E(\mathbb{F}_q^{alg})$.

This allows us to prove uniform definability of $\chi_N(E(K_1))$ by breaking into cases according to how σ acts on the N-torsion in $E(K)$ and $K^\times$.

Remark

In the case of bad characteristic $p|N$, this approach still works, but is much more complicated and involves non-reduced finite group schemes.

Higher genus curves

1. If C is an algebraic curve of genus > 1, then there is a canonically associated definable group J, called the *Jacobian*.

Higher genus curves

1. If C is an algebraic curve of genus > 1, then there is a canonically associated definable group J, called the *Jacobian*.
2. $\chi_N(C)$ is determined by examining how σ acts on the N-torsion in the multiplicative group and the Jacobian.
3. For N prime to the characteristic, this again follows from folk theorems in arithmetic geometry.
4. For $N = p^k$, one needs to carefully analyze non-reduced finite group schemes.

Example: the set of squares

Let D be the definable set of perfect squares in K_1, i.e.,

$$D = \{x^2 \mid x \in K_1\}$$

Example: the set of squares

Let D be the definable set of perfect squares in K_1, i.e.,

$$D = \{x^2 \mid x \in K_1\}$$

Then

$$\hat{\chi}(D) = \begin{cases} \hat{\chi}(K_1) & \text{in characteristic 2} \\ \frac{\hat{\chi}(K_1)+1}{2} & \text{otherwise} \end{cases}$$

Example: the set of squares

Let D be the definable set of perfect squares in K_1, i.e.,

$$D = \{x^2 \mid x \in K_1\}$$

Then

$$\hat{\chi}(D) = \begin{cases} \hat{\chi}(K_1) & \text{in characteristic 2} \\ \frac{\hat{\chi}(K_1)+1}{2} & \text{otherwise} \end{cases}$$

Note the utility of working with $\hat{\chi}$ rather than one of the individual χ_N: division by 2 wouldn't be well defined in $\mathbb{Z}/4$.

General 1-dimensional definable sets

General definable subsets of K_1 are like the set of squares:

Proposition

If $X \subseteq K_1$ is definable, then there is an algebraic curve C defined over K_1, and an algebraic map $f : C \to K_1$ defined over K_1 such that $X = f(C(K_1))$.

General 1-dimensional definable sets

General definable subsets of K_1 are like the set of squares:

Proposition

If $X \subseteq K_1$ is definable, then there is an algebraic curve C defined over K_1, and an algebraic map $f : C \to K_1$ defined over K_1 such that $X = f(C(K_1))$.

This fact follows on general model-theoretic grounds by combining

- Model completeness of ACPF
- The ability to amalgamate periodic difference fields over bases that are algebraically closed.

General 1-dimensional definable sets

General definable subsets of K_1 are like the set of squares:

Proposition

If $X \subseteq K_1$ is definable, then there is an algebraic curve C defined over K_1, and an algebraic map $f : C \to K_1$ defined over K_1 such that $X = f(C(K_1))$.

This fact follows on general model-theoretic grounds by combining

- Model completeness of ACPF
- The ability to amalgamate periodic difference fields over bases that are algebraically closed.

Example

If D is the set of squares in K_1, then take $C = K_1$ and $f : K_1 \to K_1$ given by $f(x) = x^2$.

Proposition

If $X \subseteq K_1$ is definable, then there is an algebraic curve C defined over K_1, and an algebraic map $f : C \to K_1$ defined over K_1 such that $X = f(C(K_1))$.

Proposition

In the situation above, $\hat{\chi}(X)$ is determined in an explicit way by the values

$$\hat{\chi}(C), \hat{\chi}(C \times_{K_1} C), \hat{\chi}(C \times_{K_1} C \times_{K_1} C), \ldots$$

Endgame: Beth implicit definability

The overall proof strategy is to define the χ_N uniformly across Frobenius periodic fields and models of ACPF by explicitly writing down axioms.

Endgame: Beth implicit definability

The overall proof strategy is to define the χ_N uniformly across Frobenius periodic fields and models of ACPF by explicitly writing down axioms. These axioms are:

1. Each χ_N is a strong $\mathbb{Z}/N$-valued Euler characteristic.
2. If $N|M$, then $\chi_N(X) \equiv \chi_M(X) \mod N$.

Endgame: Beth implicit definability

The overall proof strategy is to define the χ_N uniformly across Frobenius periodic fields and models of ACPF by explicitly writing down axioms. These axioms are:

1. Each χ_N is a strong $\mathbb{Z}/N$-valued Euler characteristic.
2. If $N|M$, then $\chi_N(X) \equiv \chi_M(X) \mod N$.
3. If C is an algebraic curve, then $\chi_N(C)$ is given by an explicit black box from arithmetic geometry, involving N-torsion in the Jacobian and the multiplicative group.

Endgame: Beth implicit definability

After writing down axioms, one then shows

1. The axioms are true in Frobenius periodic fields, hence in ultraproducts of Frobenius periodic fields.
2. In models of ACPF, the axioms determine at most one $\hat{\chi}|K_1$.

Endgame: Beth implicit definability

After writing down axioms, one then shows

1. The axioms are true in Frobenius periodic fields, hence in ultraproducts of Frobenius periodic fields.
2. In models of ACPF, the axioms determine at most one $\hat{\chi}|K_1$.
3. By Beth implicit definability, it follows that $\hat{\mu} := \hat{\chi}|K_1$ is 0-definable (when it exists).

Endgame: Beth implicit definability

After writing down axioms, one then shows

1. The axioms are true in Frobenius periodic fields, hence in ultraproducts of Frobenius periodic fields.
2. In models of ACPF, the axioms determine at most one $\hat{\chi}|K_1$.
3. By Beth implicit definability, it follows that $\hat{\mu} := \hat{\chi}|K_1$ is 0-definable (when it exists).
4. Because $\hat{\chi}$ is strong, it follows that $\hat{\chi}$ must be definable, if it exists.

Endgame: Beth implicit definability

After writing down axioms, one then shows

1. The axioms are true in Frobenius periodic fields, hence in ultraproducts of Frobenius periodic fields.
2. In models of ACPF, the axioms determine at most one $\hat{\chi}|K_1$.
3. By Beth implicit definability, it follows that $\hat{\mu} := \hat{\chi}|K_1$ is 0-definable (when it exists).
4. Because $\hat{\chi}$ is strong, it follows that $\hat{\chi}$ must be definable, if it exists.
5. If K is a model of ACPF, then $K \equiv L$ for some ultraproduct L of Frobenius periodic fields.
6. The Euler characteristic $\hat{\chi}$ exists and is 0-definable in L, and therefore exists and is 0-definable in K.

Directions for future research

Can we prove definability of χ_N in these other pseudo-finite settings?

- Finite-rank sets in ACFA
- Non-standard integers modulo an infinite integer.

Directions for future research

Can we prove definability of χ_N in these other pseudo-finite settings?

- Finite-rank sets in ACFA
- Non-standard integers modulo an infinite integer.

Are there any connections to p-adic L-functions?

References

- Ax, James. "The elementary theory of finite fields." 1968.
- Hrushovski, Ehud. "The elementary theory of the Frobenius automorphism." Preprint.
- Krajicek, Jan and Thomas Scanlon. "Combinatorics with definable sets: Euler characteristics and Grothendieck Rings." 2000.
- van den Dries, Lou. *Tame topology and o-minimal structures.* 1998.

Acknowledgments

I would like to thank the organizers for inviting me.

Acknowledgments

I would like to thank the organizers for inviting me.
Parts of this work were carried out while I was funded by the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1106400. Any opinions, findings, and conclusions or recommendations expressed in this talk are those of the author and do not necessarily reflect the views of the NSF.

Questions?