

DEFINABLE OPERATIONS ON SETS AND ELIMINATION OF IMAGINARIES

JAN HOLLY

(Communicated by Andreas R. Blass)

ABSTRACT. This paper gives a new and constructive proof of Poizat's theorem that the theory of algebraically closed fields admits elimination of imaginaries. The proof uses ideas of definability for properties and operations on definable sets. In addition, the property of being finite in an algebraically closed field, as well as the property of having a given algebraic dimension are shown to be definable properties.

1. INTRODUCTION

Equivalence relations and their quotient spaces have long interested mathematicians. Taking a logical approach to this study, Poizat in [Po1] introduces elimination of imaginaries with which a theory has in some sense, nice definable quotient spaces. In particular, the theory of algebraically closed fields has this property, and we present a new proof of this fact, in which the functions involved are constructed explicitly and uniformly. The tools developed here should also be useful for attacking similar problems.

To explain the ideas more formally, we begin with some terminology. For a model $\mathcal{A} = (A, \dots)$ of an L -theory T , an L_A -definable set is a set definable with parameters, that is, in the language $L \cup \{a\}_{a \in A}$. Throughout the following, formulas will have parameters from the universe listed explicitly. For example, if $\phi(\vec{a}, \vec{x})$ is an L_A -formula with $\vec{a} = (a_1, \dots, a_m) \in A^m$ and variables $\vec{x} = (x_1, \dots, x_n)$, then $\phi(\vec{t}, \vec{x})$ with variables \vec{t} and \vec{x} is an L -formula. We write $\phi(\vec{a}, A^n)$ for the set $\{\vec{b} \in A^n : \mathcal{A} \models \phi(\vec{a}, \vec{b})\}$, and $\mathcal{D}(A^n)$ for the collection of all L_A -definable subsets of A^n .

Throughout this paper, K is used to denote an arbitrary model of the theory of algebraically closed fields (ACF). The language of ACF is $\{0, 1, +, -, \cdot\}$, and the axioms are those for fields, plus those stating that every nonconstant polynomial has a root.

1.1. Definition. An L -theory T with at least two constant symbols admits *elimination of imaginaries* if for every $\mathcal{A} \models T$ and $n \in \mathbb{Z}^+$, for each definable (without parameters) equivalence relation \sim on A^n , there is a definable

Received by the editors May 13, 1991 and, in revised form, July 29, 1991.

1991 *Mathematics Subject Classification.* Primary 03C10, 03C60; Secondary 12L12, 12Y05.

Key words and phrases. Elimination of imaginaries, algebraically closed field.

©1993 American Mathematical Society
0002-9939/93 \$1.00 + \$.25 per page

function $f: A^n \rightarrow A^p$ (some p) such that

$$\vec{a} \sim \vec{b} \Leftrightarrow f(\vec{a}) = f(\vec{b}) \quad (\vec{a}, \vec{b} \in A^n).$$

The terminology comes from the fact that equivalence classes are sometimes called *imaginaries*. One way for a theory to have elimination of imaginaries is for each definable equivalence relation to have a definable set of representatives. That is, there is a definable set containing exactly one element from each equivalence class. In this case, there is a function $f: A^n \rightarrow A^n$ giving elimination of imaginaries.

Two examples of theories that eliminate imaginaries in this way are arithmetic, as noted by Poizat in [Po1], and the theory of real closed fields, as shown by Van den Dries in [vdD1].

One theory without definable sets of representatives for definable equivalence relations is the theory of algebraically closed fields. There is no way, for example, in a field of characteristic zero to definably choose a square root of two, since there is an automorphism interchanging the two square roots. However, Poizat showed in [Po1] that ACF does admit elimination of imaginaries. His proof uses an equivalent, more abstract definition of elimination of imaginaries, and shows that for every $K \models \text{ACF}$, for each L_K -definable set $S \subset K^n$, there is a set $C \subset K$ such that any automorphism of K preserves S iff it fixes C pointwise.

Presented here is a proof that gives explicitly the desired function f for each definable equivalence relation in ACF; in fact, the construction of the formula for f is independent of the particular field $K \models \text{ACF}$. The same method, given an L_K -definable set S , produces explicitly a finite set C as described above. For algebraically closed fields of characteristic zero, the definable closure of C is the smallest field of definition for S .

The construction uses the ideas of definable properties and operations on sets, which we define in the next section.

2. DEFINABLE PROPERTIES AND OPERATIONS

2.1. Definition. Let T be an L -theory and \mathcal{P} be a property of L_A -definable sets $S \subset A^n$ (fixed n), for models \mathcal{A} of T . We call \mathcal{P} a *definable property* if there is an $L(R)$ -sentence $\sigma(R)$ with new n -ary relation symbol R such that for every $\mathcal{A} \models T$ and every set $S \in \mathcal{D}(A^n)$,

$$S \text{ has property } \mathcal{P} \Leftrightarrow (\mathcal{A}, S) \models \sigma(R).$$

In other words, $\sigma(R)$ says, "The set represented by R has property \mathcal{P} ."

As suggestive notation when discussing sets $S \subset A^n$, the new n -ary relation symbol S is sometimes used. (No confusion should arise.)

2.2. Examples. (1) The property of having exactly two elements is definable in every theory (with equality). (2) In the theory of real closed fields, the property of containing an open interval (\Leftrightarrow being infinite) is definable.

A somewhat more surprising example is the following:

2.3. Proposition. *In ACF, the property of being a finite subset of K is definable.*

Proof. Using the strong minimality of ACF (every set $S \in \mathcal{D}(K)$ is either finite or cofinite), the following holds for sets $S \in \mathcal{D}(K)$:

$$S \text{ is finite} \Leftrightarrow S + S \neq K.$$

The property “ $S + S \neq K$ ” can be expressed easily by an $L(S)$ -sentence. \square

2.4. Key Facts. For a definable property \mathcal{P} in an L -theory T :

(1) Given an L -formula $\phi(\vec{x}, \vec{y})$ ($\vec{x} = (x_1, \dots, x_m)$, $\vec{y} = (y_1, \dots, y_n)$), there is an L -formula $\psi(\vec{x})$ such that for each $\mathcal{A} \models T$ and $\vec{a} \in A^m$,

$$\text{The set } \phi(\vec{a}, A^n) \text{ has property } \mathcal{P} \Leftrightarrow \mathcal{A} \models \psi(\vec{a}).$$

(2) Let $\phi(\vec{x})$ be a formula in an expansion L' of L , but still defining an L_A -definable subset $\phi(A^n)$ of A^n . Then there is an L' -sentence expressing, “The set $\phi(A^n)$ has property \mathcal{P} .”

The next result is not used in the remaining sections, but is of independent interest.

2.5. Theorem. In ACF, the property $\mathcal{P}_{n,d}$ for L_K -definable sets $S \subset K^n$, of having dimension d is definable.

Proof. For sets $S \in \mathcal{D}(K)$, $\dim(S) = -\infty$ iff $S = \emptyset$, $\dim(S) = 0$ iff S is finite and nonempty, and $\dim(S) = 1$ iff S is cofinite. Each of these properties is definable since finiteness is a definable property.

Next, let $n > 1$ and use induction. Say $S \in \mathcal{D}(K^n)$, and consider the projection $\pi: S \rightarrow K^{n-1}$ on the first $n - 1$ coordinates. Dimension can be defined inductively (see [vdD2]) as follows:

$$\dim(S) = \max\{\dim(B_0), \dim(B_1) + 1\}$$

where $B_0 = \{\vec{x} \in \pi S : \dim(\pi^{-1}(\vec{x})) = 0\}$ and $B_1 = \{\vec{x} \in \pi S : \dim(\pi^{-1}(\vec{x})) = 1\}$.

The sets B_0 and B_1 are definable in the language $L(S)$, since the properties of having dimensions 0 and 1 in K are definable by the inductive hypothesis. Then using the inductive hypothesis on B_0 and B_1 in K^{n-1} , we can express in the language $L(S)$ that the set S has dimension d by

$$\bigvee_{\max\{b, c+1\}=d} \dim(B_0) = b \ \& \ \dim(B_1) = c. \quad \square$$

Next we consider operations sending L_A -definable sets to L_A -definable sets in models \mathcal{A} of an L -theory T .

2.6. Definition. Let T be an L -theory. Fix positive integers n and p , and suppose that β is an operation with $\beta|_{\mathcal{A}}: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^p)$ for every $\mathcal{A} \models T$. Then β is a *definable operation* if there is an $L(R)$ -formula $\sigma(R, \vec{y})$ with new n -ary relation symbol R and $\vec{y} = (y_1, \dots, y_p)$, such that if $\mathcal{A} \models T$ and $S \in \mathcal{D}(A^n)$, then

$$\beta(S) = \{\vec{a} \in A^p : (\mathcal{A}, S) \models \sigma(R, \vec{a})\}.$$

The formula $\sigma(R, \vec{y})$ is said to *define* β in T .

Since a definable operation is always an operation in every model of T , we abbreviate “For all $\mathcal{A} \models T$, $\beta|_{\mathcal{A}}: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^p)$,” by “ $\beta: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^p)$.”

2.7. Notation. Let $\sigma(R, \vec{y})$ be an L -formula defining an operation $\beta: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^p)$ in a theory T , and let $\phi(\vec{x})$ be a formula in an expansion L' of L , but with the set $\phi(A^n)$ still being L_A -definable. Then replacing $R(\)$

with $\phi(\)$ everywhere in σ , we obtain an L' -formula, denoted $\sigma(\phi(A^n), \vec{y})$, which defines the set $\beta(\phi(A^n))$.

We may also write $\sigma(\psi(\vec{v}, A^n), \vec{y})$ if a formula $\psi(\vec{v}, \vec{x})$ defines a family of subsets of A^n , parameterized by variables $\vec{v} = (v_1, \dots, v_q)$ (some q). Then for fixed $\vec{a} \in A^q$, the formula $\sigma(\psi(\vec{a}, A^n), \vec{y})$ defines the set $\beta(\psi(\vec{a}, A^n))$ as long as the set $\psi(\vec{a}, A^n)$ is L_A -definable.

2.8. Proposition. *Let T be an L -theory with models \mathcal{A} .*

(1) *If $\beta_1: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^p)$ and $\beta_2: \mathcal{D}(A^p) \rightarrow \mathcal{D}(A^q)$ are definable operations, then the composition $\beta_2 \circ \beta_1: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^q)$ is a definable operation.*

(2) *If $f: A^n \rightarrow A^p$ is a definable function (given in all models \mathcal{A} by an L -formula independent of \mathcal{A}), then the operation β with $\beta(S) = f(S)$ for all $S \in \mathcal{D}(A^n)$ is a definable operation.*

(3) *Given a definable equivalence relation on A^n (given in all models \mathcal{A} by an L -formula independent of \mathcal{A}) and a definable operation $\beta: \mathcal{D}(A^n) \rightarrow \mathcal{D}(A^p)$ that sends each equivalence class to a one-element set, write*

$$\beta(\text{eq. class of } \vec{x}) = \{\vec{e}_{\vec{x}}\} \subset A^p, \quad \text{for } \vec{x} \in A^n.$$

Then the function $f: A^n \rightarrow A^p$ defined by $f(\vec{x}) = \vec{e}_{\vec{x}}$ is a definable function.

The proofs of these facts are straightforward.

3. ELIMINATION OF IMAGINARIES

To prove that ACF admits elimination of imaginaries, we show that equivalence classes can be mapped to finite sets, which can then be mapped to single-element sets.

3.1. Lemma. *There is a definable operation β_n sending each nonempty L_K -definable set $S \subset K^n$ (K an algebraically closed field) to a nonempty finite subset of itself.*

Proof. By induction on n . For $n = 1$, the L_K -definable sets are the finite sets and the cofinite sets. Proposition 2.3 allows us to define an operation β_1 by an $L(S)$ -formula $\sigma(S, y)$ saying

$$\begin{aligned} & \text{“}S \text{ is finite and } y \in S, \text{ or} \\ & S \text{ is not finite and } y \in S \cap (\{x : x^2 \in \bar{S}\} \cup \{x : x^3 \in \bar{S}\} \cup \{1\}) \text{”} \end{aligned}$$

where \bar{S} denotes the complement of S .

For finite S , this defines S itself. If S is infinite, then \bar{S} is finite. The set, Q , of square and cube roots of elements in \bar{S} will have a finite nonempty intersection with S unless $S = K$ or $S = K \setminus \{0\}$. In any case, the set $S \cap (Q \cup \{1\})$ is a finite nonempty subset of S .

Note: If only the characteristic zero case is under consideration, the simpler formula “ $y \in S \cap ((\bar{S} + 1) \cup \{1\})$ ” can be used for $\sigma(S, y)$.

For $n > 1$, define β_n by considering the projection $\pi: S \rightarrow K^{n-1}$ of a set $S \in \mathcal{D}(K^n)$ to the first $n - 1$ coordinates. By the inductive hypothesis, a nonempty finite subset $\{\vec{b}_1, \dots, \vec{b}_k\}$ of $\pi(S)$ is chosen by β_{n-1} . Then for each $i = 1, \dots, k$, project the fiber $\pi^{-1}(\vec{b}_i)$ to the last coordinate, and use β_1

to choose a finite subset $\{c_{i1}, \dots, c_{i m(i)}\}$. So the set

$$\bigcup_{i=1}^k \{(\vec{b}_i, c_{i1}), \dots, (\vec{b}_i, c_{i m(i)})\} \subset K^n$$

is a nonempty finite subset of S .

Formally, for $i = 1, \dots, n - 1$, let $\sigma_i(S_i, \vec{y}_i)$ define the operation β_i . Then considering sets $S_n \in \mathcal{D}(K^n)$,

$$\begin{aligned} (x_1, \dots, x_{n-1}, z) \in \beta_n(S_n) &\Leftrightarrow \vec{x} \in \beta_{n-1}(\pi(S_n)) \ \& \ z \in \beta_1(\{y : (\vec{x}, y) \in S_n\}) \\ &\Leftrightarrow \sigma_{n-1}(\exists y S_n(K^{n-1}, y), \vec{x}) \ \& \ \sigma_1(S_n(\vec{x}, K), z). \end{aligned}$$

Let $\sigma_n(S_n, \vec{x}, z)$ be this $L(S_n)$ -formula. \square

3.2. Lemma. *For fixed n and k , there is a definable operation sending distinct k -element sets $\{\vec{a}_1, \dots, \vec{a}_k\} \subset K^n$ to distinct one-element sets in K^p (fixed p).*

Proof. This is related to Lemma 5 in [Po1].

Case $n = 1$. To a k -element set $\{a_1, \dots, a_k\} \subset K$, associate the vector $(b_1, \dots, b_k) \in K^k$ of coefficients of the polynomial

$$(z - a_1)(z - a_2) \cdots (z - a_k) = z^k + b_1 z^{k-1} + \cdots + b_{k-1} z + b_k.$$

So for $n = 1$, we have $p = k$. The formula $\sigma(S, \vec{y})$ defining the operation says,

“There are exactly k members, a_1, \dots, a_k , of S , and

$$y_1 = \sum_{i=1}^k a_i \ \& \ y_2 = \sum_{1 \leq i < j \leq k} a_i a_j \ \& \ \cdots \ \& \ y_k = \prod_{i=1}^k a_i.”$$

Unique factorization in $K[z]$ guarantees that distinct k -element sets map to distinct sequences of coefficients. (“Sequence of coefficients” = “point in K^p .”) Notice that the order of a_1, \dots, a_k is unimportant, so the operation is well defined.

Case $n > 1$. To each k -element set

$$\{(a_{11}, \dots, a_{1n}), \dots, (a_{k1}, \dots, a_{kn})\} \subset K^n$$

associate the vector, in some predetermined order, of coefficients of the polynomial

$$\begin{aligned} &(z - a_{11}x_1 - \cdots - a_{1n}x_n) \cdots (z - a_{k1}x_1 - \cdots - a_{kn}x_n) \\ &= \sum_{i_0 + \dots + i_n = k} b_{i_0, \dots, i_n} z^{i_0} x_1^{i_1} \cdots x_n^{i_n}. \end{aligned}$$

Just as in the $n = 1$ case, there is an $L(S)$ -formula defining this operation, and the operation is well defined, mapping distinct finite sets to distinct sequences. \square

Using the usual tricks to pad vectors of length $< q$ to get vectors of length q , we obtain:

3.3. Corollary. For fixed n and k , there is an integer q and a definable operation $\alpha_{n,k}$ sending distinct finite sets of $j \leq k$ elements $\{\vec{a}_1, \dots, \vec{a}_j\} \subset K^n$ to distinct one-element sets in K^q .

3.4. Lemma. Let formula $\phi(\vec{x}, \vec{y})$, with “parameters” \vec{x} , define a family $\{\phi(\vec{a}, K^n) : \vec{a} \in K^m\}$ of subsets of K^n , and let β_n be the operation described in Lemma 3.1. Then there is an integer k such that $\#(\beta_n(\phi(\vec{a}, K^n))) \leq k$ for every $\vec{a} \in K^m$.

Proof. By induction on n . For $n = 1$, the operation β_1 sends finite sets to themselves and cofinite sets S to sets of cardinality $\leq 5(\#\overline{S}) + 1$. The strong minimality of ACF implies there is an integer r such that for all $\vec{a} \in K^m$,

$$\phi(\vec{a}, K) \text{ finite} \Rightarrow \#(\phi(\vec{a}, K)) \leq r,$$

$$\phi(\vec{a}, K) \text{ cofinite} \Rightarrow \#(\neg\phi(\vec{a}, K)) \leq r,$$

so let $k = 5r + 1$.

For $n > 1$, the set $\beta_n(\phi(\vec{a}, K^n))$ is the union over all $(b_1, \dots, b_{n-1}) \in \beta_{n-1}(\exists y_n \phi(\vec{a}, K^{n-1}, y_n))$ of the sets $\beta_1(\phi(\vec{a}, b_1, \dots, b_{n-1}, K))$. Let integer s be a cardinality bound from the inductive hypothesis applied to the formula $\exists y_n \phi(\vec{x}, y_1, \dots, y_{n-1}, y_n)$ with parameters \vec{x} , and let r be such a bound for the formula $\phi(\vec{x}, y_1, \dots, y_{n-1}, y_n)$ with all $\vec{x}, y_1, \dots, y_{n-1}$ considered as the parameters. Then by the definition of β_n , the integer $k = r \cdot s$ is a bound on $\#(\beta_n(\phi(\vec{a}, K^n)))$. \square

3.5. Corollary. Given a definable equivalence relation on K^n , there is an integer k such that each equivalence class E has $\#(\beta_n(E)) \leq k$, where β_n is the operation described in Lemma 3.1.

3.6. Theorem. ACF admits elimination of imaginaries (Poizat’s Theorem). Moreover, in an algebraically closed field K , from a formula $\theta(\vec{x}, \vec{y})$ defining an equivalence relation \sim on K^n , one can construct a formula defining a function $f: K^n \rightarrow K^q$ (some q) such that $\vec{a} \sim \vec{b} \Leftrightarrow f(\vec{a}) = f(\vec{b})$. In fact, the formula defining f is independent of the particular field K .

Proof. Given the definable equivalence relation \sim on K^n defined by θ , let β_n be the operation in Lemma 3.1 and let k be a bound (independent of K) on the size of the images under β_n of the equivalence classes. Then letting $\alpha_{n,k}$ be the operation of Corollary 3.3, the composition $\alpha_{n,k} \circ \beta_n: \mathcal{D}(K^n) \rightarrow \mathcal{D}(K^q)$ sends each equivalence class to a distinct one-element set in K^q . Using Proposition 2.8(3), there is a function $f: K^n \rightarrow K^q$ such that $\alpha_{n,k} \circ \beta_n(\text{eq. class of } \vec{x}) = \{f(\vec{x})\}$ for all $\vec{x} \in K^n$. So $f(\vec{a}) = f(\vec{b})$ iff \vec{a} and \vec{b} are in the same equivalence class.

Since all the functions and operations mentioned are actually constructed in the proofs of the previous results, this gives a formula for f . \square

4. CONSTRUCTIONS

Construction of representatives for points in projective space. The function f in Theorem 3.6 gives actual points in K^q “representing” the classes of an equivalence relation, so one might ask how the function f behaves when forming a particular quotient space. One example of a widely used quotient space is \mathbb{P}^n , the projective space over a field, say the field of complex numbers.

Question. How are the points of \mathbb{P}^n represented in some \mathbb{C}^q by the function f ?

Answer. First remember that the space \mathbb{P}^n is $(\mathbb{C}^{n+1} \setminus \{0\}) / \sim$ where the equivalence relation \sim is given by

$$(x_0, \dots, x_n) \sim (y_0, \dots, y_n) \Leftrightarrow \exists \lambda \neq 0 \quad (x_0, \dots, x_n) = (\lambda y_0, \dots, \lambda y_n).$$

Use $(x_0 : x_1 : \dots : x_n)$ to denote the equivalence class of (x_1, \dots, x_n) , and for the moment include the point $0 \in \mathbb{C}^{n+1}$ as having its own equivalence class.

Consider the case $n = 1$. Each class is sent to a nonempty finite subset of itself by the operation β_2 of Lemma 3.1 as follows:

To begin, project to the first coordinate to get a subset of \mathbb{C} . Considering first the classes $(x_0 : x_1)$ with $x_0 \neq 0$, this always gives the set $\mathbb{C} \setminus \{0\}$. Apply β_1 to get $\beta_1(\mathbb{C} \setminus \{0\}) = \{1\}$. Then in the fiber above $\{1\}$, each class has a single point $(1, \frac{x_1}{x_0})$, so β_2 sends such a class to the finite subset consisting of that single point.

Next consider the equivalence class $E = (0 : x_1)$ with $x_1 \neq 0$. The projection to the first coordinate gives $\{0\} \subset \mathbb{C}$, and β_1 sends this to itself. The fiber above $\{0\}$ is the entire equivalence class E , so projecting now to the second coordinate gives $\mathbb{C} \setminus \{0\}$, which β_1 sends to $\{1\}$. This means that β_2 chooses $\{(0, 1)\}$ as the finite subset of E . The final equivalence class, $(0 : 0)$, clearly maps to itself under the operation β_2 .

Normally at this stage, the map $\alpha_{n,k}$ is needed in order to send finite sets to single-element sets, but in this case the sets already consist of a single point each. We can stop here and obtain the function f from the operation β_2 above. In summary, points in \mathbb{P}^1 are represented by points in \mathbb{C}^2 as follows:

$$\begin{aligned} \text{for } x_0 \neq 0, \quad & (x_0 : x_1) \mapsto (1, x_1/x_0); \\ \text{for } x_1 \neq 0, \quad & (0 : x_1) \mapsto (0, 1). \end{aligned}$$

Noticing that this is the usual way to canonically represent points in projective space, we ask whether the same thing happens for \mathbb{P}^n with $n > 1$. It is not hard to see that the process works just as well, giving representatives in \mathbb{C}^{n+1} ,

$$\begin{aligned} \text{for } x_0 \neq 0, \quad & (x_0 : x_1 : \dots : x_n) \mapsto \left(1, \frac{x_1}{x_0}, \frac{x_2}{x_0}, \dots, \frac{x_n}{x_0}\right); \\ \text{for } x_1 \neq 0, \quad & (0 : x_1 : \dots : x_n) \mapsto \left(0, 1, \frac{x_2}{x_1}, \dots, \frac{x_n}{x_1}\right); \\ & \vdots \\ & (0 : \dots : 0 : x_n) \mapsto (0, \dots, 0, 1). \end{aligned}$$

A final remark on this example: the points chosen by β_{n+1} to represent the elements of projective space *are actually members* of their respective equivalence classes. Of course, we cannot expect this to happen for all definable equivalence relations.

Construction of the smallest field of definition of a set. Poizat’s presentation in [Po1] shows that if an L -theory T admits elimination of imaginaries, then for every L_A -definable set $S \subset A^n$ ($\mathcal{A} \models T$), there is an L_A -formula $\phi(\vec{c}, \vec{x})$ with parameters c_1, \dots, c_m defining S , such that for no other tuple \vec{d} of

parameters does $\phi(\vec{d}, \vec{x})$ define S . (In fact, this is an equivalent definition of elimination of imaginaries.) Let $C = \{c_1, \dots, c_m\}$. Then an automorphism δ of \mathcal{A} fixes the set S iff δ fixes C pointwise. In addition, the definable closure, F , of C is the smallest definably closed set of definition of S . (F is unique in this sense.)

In ACF, when considering a set $S \in \mathcal{D}(K^n)$ where K has characteristic zero, this set F is the smallest field of definition of S . That is, S can be defined by a formula that uses parameters only from F , and F is the smallest subfield of K for which this is true. Using Theorem 3.6, we can construct this subfield for a set $S \in \mathcal{D}(K^n)$ as follows.

Given an L_K -formula $\psi(\vec{a}, \vec{x})$ defining S , define an equivalence relation \sim on K^m by $\theta(\vec{v}, \vec{w}) := \forall \vec{x}(\psi(\vec{v}, \vec{x}) \leftrightarrow \psi(\vec{w}, \vec{x}))$. With the function f in Theorem 3.6 (here, $f: K^m \rightarrow K^q$),

$$f(\vec{a}) = f(\vec{b}) \Leftrightarrow \vec{a} \sim \vec{b} \\ \Leftrightarrow \psi(\vec{a}, \vec{x}) \text{ and } \psi(\vec{b}, \vec{x}) \text{ define the same set.}$$

Letting $\vec{c} = f(\vec{a}) \in K^q$, the formula

$$\phi(\vec{c}, \vec{x}) := \exists \vec{v} (f(\vec{v}) = \vec{c} \ \& \ \psi(\vec{v}, \vec{x}))$$

defines S , and for no other tuple \vec{d} does $\phi(\vec{d}, \vec{x})$ define S . The definable closure of $\{c_1, \dots, c_q\}$ is then the smallest field of definition of S .

Remark. In such a construction, quantifier elimination is used in analyzing definable sets. A constructive proof of the quantifier elimination can be found in [KK].

5. STRONGLY MINIMAL STRUCTURES

Notice that Lemma 3.1 is actually stronger than necessary. It gives a *single* definable operation that sends *every* nonempty L_K -definable set $S \subset K^n$ to a finite subset. It would be sufficient to have *for each* definable equivalence relation an operation sending the equivalence classes to finite subsets.

Since the equivalence classes of a definable equivalence relation form a uniformly definable family of sets (obtained by substituting different parameters into a fixed formula), such a result can be generalized, as was pointed out by Anand Pillay in conversation.

5.1. Proposition. *Let \mathcal{A} be strongly minimal with the algebraic closure of \emptyset ($\text{acl}(\emptyset)$) being infinite. Given a uniformly definable family \mathcal{F} of sets $S \subset A^n$ (some n), there is a definable operation β that sends each nonempty set $S \in \mathcal{F}$ to a nonempty finite subset of itself.*

Proof. By induction on n .

For $n = 1$, by strong minimality there is an integer k such that

$$S \in \mathcal{F} \text{ finite} \Rightarrow \#(S) \leq k, \\ S \in \mathcal{F} \text{ cofinite} \Rightarrow \#(\bar{S}) \leq k,$$

so the property of being a finite member of \mathcal{F} is a definable property, as is that of being a cofinite member of \mathcal{F} . Since $\text{acl}(\emptyset)$ is infinite, there is a finite definable set C with $\#(C) > k$. Let $\beta(S) = S$ for finite $S \in \mathcal{F}$, and $\beta(S) = S \cap C$ for cofinite $S \in \mathcal{F}$.

For $n > 1$, the operation is defined by considering projections and using induction, as in the proof of Lemma 3.1. One must also check that the projections and other sets involved belong to uniformly definable families in order to use the inductive hypothesis. The details are left to the reader. \square

Also notice in Lemma 3.1 that for an L_K -definable set S , the finite subset constructed is contained in every algebraically closed set over which S is definable. A theory T is said to have *weak elimination of imaginaries* if every L_A -definable set $S \subset A^n$ ($\mathcal{A} \models T$) has a smallest algebraically closed set of definition (cf. [Po2, p. 422]). In fact, elimination of imaginaries automatically gives weak elimination of imaginaries.

For strongly minimal theories in which every algebraically closed set is a model ($\text{acl}(\emptyset)$ being infinite implies this), the existence of smallest algebraically closed sets of definition follows from Remark 2.4 in [Pi]. Namely, take an algebraically closed set of definition of smallest dimension. Note that this argument shows the existence of such a set, but does not give an actual construction. For ACF, the method of Lemma 3.1 allows one to construct these sets.

ACKNOWLEDGMENT

I would like to thank Lou van den Dries for suggesting this problem, in particular, for the idea of associating definable sets with finite subsets. I would also like to thank Alexandre Borovik for his question about projective space.

REFERENCES

- [vdD1] Lou van den Dries, *Algebraic theories with definable Skolem functions*, J. Symbolic Logic **49** (1984), 625–629.
- [vdD2] ———, *Dimension of definable sets, algebraic boundedness and henselian fields*, Ann. Pure Appl. Logic **45** (1989), 189–209.
- [KK] Kreisel and Krivine, *Elements of mathematical logic*, North-Holland, Amsterdam, 1971.
- [Pi] Anand Pillay, *Some remarks on definable equivalence relations in O -minimal structures*, J. Symbolic Logic **51** (1986), 709–714.
- [Po1] Bruno Poizat, *Une théorie de Galois imaginaire*, J. Symbolic Logic **48** (1983), 1151–1170.
- [Po2] ———, *Cours de théorie des modèles*, Nur al-Mantiq wal-Ma'rifah, Paris, 1985.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, URBANA, ILLINOIS 61801
E-mail address: hutson@symcom.math.uiuc.edu