Definable Sets in Valued Fields

Deirdre Haskell and Dugald Macpherson

Contents

1. Introduction

Let K be an algebraically closed field equipped with a non-trivial valuation to an ordered abelian group Γ. We work with norms, and write Γ multiplicatively, so the ultrametric inequality states $|x + y| \leq \mathrm{Max}\{|x|, |y|\}$. Let k be the residue field. It has been known for a long time that the theory of K has quantifier elimination in the language of rings together with a predicate for the relation $|x| \leq |y|$: indeed, this follows from a model-completeness result of Abraham Robinson [13]. It does not have elimination of imaginaries just with the sort K, since elements of Γ and k are not coded in K. In joint work of Haskell, Hrushovski and Macpherson [4] it is shown that elimination of imaginaries does hold once certain sorts are added, but the necessary sorts are quite complicated. In [5] further ideas from stability theory are developed in the context of algebraically closed valued fields, working with the extra sorts so that elimination of imaginaries holds. The original hope in this project was to prove a conjecture of Holly [7], that elimination of imaginaries holds just with sorts for K, Γ and k, together with a sort for the collection of all *open balls* $B_{<\gamma}(a) = \{x \in K : |x - a| < \gamma\}$, where $a \in K$ and $\gamma \in \Gamma$, and a sort for all *closed balls* $B_{\leq\gamma}(a) = \{x \in K : |x - a| \leq \gamma\}$. This conjecture is false, though for a long time we believed it true. In this survey paper, our goal is to discuss the material of the two papers [4] and [5] in a more informal way. We hope this will serve both to entice the readers to peruse the two longer papers, and to make these papers more accessible. The paper contains no new results.

2. Elimination of Imaginaries

2.1. A template for proving elimination of imaginaries

We begin with an outline of the proof of elimination of imaginaries for an algebraically closed valued field, with respect to a distinguished family of sorts

* Deirdre Haskell: Research partially supported by NSF and NSERC grants.
Dugald Macpherson: Research supported partially by an EPSRC grant.
Both authors were supported by NSF funding for a semester at the Mathematical Sciences Research Institute, Berkeley.

(Theorem 2.1 below). At the level of detail which we give in this section, the outline is extremely general, and could presumably be used to prove elimination of imaginaries in many different contexts. Later in this paper we discuss how a similar proof has been used by T. Mellor to prove elimination of imaginaries for real closed valued fields. The outline here uses a sequence of 'black boxes'; our goal is to motivate the theory which needs to be developed in order to illuminate these black boxes. See Section 2.6 for a discussion of other extensions of these methods. Let K be a structure with a distinguished family of sorts from K^{eq} whose union is G. For any definable subset X of K^n we will write $\ulcorner X \urcorner$ for the element of K^{eq} (unique up to interdefinability over $\emptyset$) which is interdefinable with X. We will call any tuple c from K^{eq} a *code* for X if $\ulcorner X \urcorner \in \mathrm{dcl}(c)$ and $c \in \mathrm{dcl}(\ulcorner X \urcorner)$ (where the definable closure is in K^{eq} unless specified otherwise). Then K has elimination of imaginaries with respect to the sorts of G if and only if every definable set has a code in G. If finite sets of tuples in G are coded in G, then the condition for c to be a code for X can be weakened to $c \in \mathrm{acl}(\ulcorner X \urcorner)$. In general, when we say that K has elimination of imaginaries to the sorts G, we do not have in mind a particular (many-sorted) language in which elimination of imaginaries holds. Given K as a structure, with its class of definable sets, the particular choice of language which defines that same class of sets is not an issue for elimination of imaginaries. In the case of an algebraically closed valued field, the class of definable sets is given by the Robinson language for which the theory eliminates quantifiers. In general, K is a structure over some language L, and there is a canonical extension L^{eq} for K^{eq}. Indeed, each sort S of K^{eq} has the form K^n/E for some $n > 0$ and some 0-definable equivalence relation E on K^n, and L^{eq} consists of L together with sumbols for the sorts and function symbols for the obvious functions $f_S : K^n \to K^n/E$. In the cases in which we are interested, one of the sorts of G is K itself, and we can suppose the language for G is the fragment of L^{eq} consisting of L together with the functions f_S for each sort S of G. With this language, any definable subset of K^n is definable already in the language L. For the particular case of algebraically closed valued fields, we mention in Section 2.2 another language for G in which additionally quantifier elimination holds. Again, no new definable sets are added. We usually assume that K is a large and sufficiently saturated model of its theory, and so G also is a large and sufficiently saturated model, in some unspecified language. The foundation of our proof of elimination of imaginaries is the following simple proposition, which shifts the focus from coding arbitrary definable sets to coding definable functions of one variable.

Proposition 2.1. *Let K be a structure, and $\{G_i : i \in I\}$ be a collection of sorts from K^{eq}, with $G_0 = K$ and $G = \bigcup_{i \in I} G_i$. Assume that for every definable subset U of K, every $i \in I$, and every definable function $f : U \to G_i$, the pair (U, f) is coded by some tuple from G. Then every element of K^{eq} is coded in G.*

Proof – We show by induction that every n-ary relation on K is coded. The case $n = 1$ holds by assumption. Suppose that $X \subset K^{n+1} = K \times K^n$ is definable, and let Y be the projection of X to the first coordinate. For each $a \in Y$, let $X(a) := \{x : (a, x) \in X\}$. By the inductive assumption, each $X(a)$ is coded by some tuple $h(a)$ in G. By compactness, the function h is definable, and Y can be partitioned into finitely many pieces $U_1, \dots, U_k$ such that for each $j = 1, \dots k$, $h|_{U_j}$ is a function to some product of the G_i (and different j correspond to different products). By assumption, each pair $(U_j, h|_{U_j})$ is coded by some tuple c_j in G. Now X is coded by $(c_1, \dots, c_k)$. □

We actually use a slightly more refined version of Proposition 2.1. We identify certain privileged subsets of G called *unary sets* which share with balls in K a useful notion of genericity. These sets play the role of one-types. It should be said that elimination of imaginaries could be proved, perhaps slightly more simply, using Proposition 2.1 rather than the following Proposition 2.2. The second result, however, is more consistent with our general point of view. In particular, some of our results on definable functions on unary sets support the later independence theory.

Proposition 2.2. *Let K be a structure with distinguished sorts $\{G_i : i \in I\}$ with union G, in a language with at least one $\emptyset$-definable symbol. Suppose that K has an $\mathrm{Aut}(K)$-invariant family $\mathcal{U}$ of definable sets with the following property: for every $a \in K^n$ there is a sequence $(a_1, \dots, a_m)$ from K^{eq} such that $\mathrm{dcl}(a) = \mathrm{dcl}(a_1, \dots, a_m)$ and for each $i \leq m$, there is $U \in \mathcal{U}$ such that $a_i \in U$ and U is $a_1 \dots a_{i-1}$-definable. Assume also that whenever g is a definable* *function from a set in $\mathcal{U}$ to G then g is coded in G. Then all definable subsets of K^n are coded in G.*

We describe now the black boxes required in general, and how they work for algebraically closed valued fields. *Black Box 1.* This consists of a description of the unary sets, that is, the sets in the family $\mathcal{U}$ of the last proposition, and some ∞-definable analogues. It must be shown that they satisfy the hypothesis of Proposition 2.2 (see Proposition 2.5 below). It will be seen that Black Boxes 1, 3 and 5 are all really statements about unary sets. *Black Box 2.* Every finite subset of G is coded in G. The proof of this is somewhat intricate, but a constituent of it is given in Proposition 2.6 below. By Proposition 2.2 and this coding of finite sets in G, the following proposition will yield elimination of imaginaries to G.

Proposition 2.3. *Let $U \in \mathcal{U}$ be definable and $f : U \to G$ a definable function. Let $B = \mathrm{acl}(\ulcorner f \urcorner) \cap G$. Then $f \in \mathrm{dcl}(B)$.*

Proof of Proposition 2.3 – Consider

$$\Sigma := \{D \subset U : D, f|_D \text{ both definable over } B\}.$$

If $\bigcup \Sigma = U$, then by compactness, f is B-definable, so we may suppose $\bigcup \Sigma \neq U$. Then there is a complete type p over B whose realisations lie in $U \setminus \bigcup \Sigma$. *Black Box 3.* Then p is the generic type of a unary set $V \in \mathcal{U}$ defined (or ∞-defined) over B. This is an elementary fact about subsets of unary sets, and the corresponding notion of 'generic'. *Black Box 4.* There is a B-definable function g with the same germ on V as f. This is really the core of our whole proof, and requires considerable work. It is sketched in Section 2.5 below, but methods from Sections 2.2–2.4 are also used. The notion of *germ* will be clarified later. Let $X := \{x \in U : f(x) = g(x)\}$. Then $X \cap V$ is non-empty. *Black Box 5.* The definable set X is coded in G. This elementary fact holds because unary sets look very much like balls in K, and by a result of Holly [7], definable subsets of K are canonically expressible as Boolean combinations of balls. Black Box 2 (the coding of finite sets) is also required here. To complete the proof of Proposition 2.3, and hence of elimination of imaginaries to G, observe that X is $B\ulcorner f \urcorner$-definable. Hence X is B-definable, as X is coded in G and $B = \mathrm{acl}(\ulcorner f \urcorner) \cap G$. As p is a complete type over B, it follows that $X \supseteq V$. But as g is B-definable, $f|_X$ is B-definable, so X is an element of Σ, a contradiction. □

Notation. We use letters a, b, x and so on to denote tuples (not just singletons) from an ambient structure. Generally, we work in a large sufficiently saturated model K of the theory of algebraically closed non-trivially valued fields. After the relevant sorts are introduced, G denotes the corresponding multi-sorted structure. In Section 3.2, where the setting is more general, $\mathcal{U}$ denotes the monster model. We use $A, B, C, \ldots$ to denote small subsets of this large model, and M, N to denote small elementary submodels. We often write $a \equiv_C b$ to mean that $\mathrm{tp}(a/C) = \mathrm{tp}(b/C)$. If A, C are sets, we sometimes write $\mathrm{tp}(A/C)$ for the type over C of some (possibly infinite) tuple enumerating A. We often write AB for $A \cup B$, or ABc for $A \cup B \cup \{c\}$. If p is a type over a large saturated model $\mathcal{U}$, say, and A is a parameter set, we write $p|A$ for the restriction of p to A. We commonly write $\Gamma(A)$ for $\mathrm{dcl}(A) \cap \Gamma$, and $k(A)$ for $\mathrm{dcl}(A) \cap k$. We denote by ACVF the incomplete theory of algebraically closed non-trivially valued fields. By the work of Robinson [13], its completions are determined by the characteristics of the field and its residue field.

2.2. Torsors, modules and the sorts of G

We now return to the situation of an algebraically closed field K with non-trivial valuation $|\cdot| : K \to \Gamma$ to an ordered abelian group Γ. As mentioned earlier, we write Γ multiplicatively. The valuation ring, namely $\{x \in K : |x| \leq 1\}$, is denoted by R, its maximal ideal by $\mathcal{M}$ and the residue field by $k = R/\mathcal{M}$. We emphasise that the value group is stably embedded. In fact, any K-definable subset of Γ^n is definable, with parameters in Γ, in the divisible ordered abelian group $(\Gamma, <, .)$. Likewise, k is an algebraically closed field, and any K-definable subset of k^n is definable in $(k, +, .)$. Thus, k is a strongly minimal set in K^{eq}, and Γ is o-minimal. All these assertions follow from quantifier elimination in a language with sorts K, k, Γ, which itself follows easily from [13]. (The language has symbols for $+, -, ., 0, 1$ on K and on k, for $., <, 1$ on Γ, and also has a valuation map $|\cdot| : K \to \Gamma$ and a residue map res : $K^2 \to k$, with $\mathrm{res}(x, y)$ equal to the residue of xy^{-1} if $|x| \le |y|$, and equal to $0 \in k$ otherwise.) The additional sorts in G are uniformly definable families of *R-torsors* in K^n, that is cosets of R-submodules of K^n. (Formally, an R-torsor is a set together with a regular action of an R-module on it, so is like the affine space of a vector space.) We will see in later sections how general modules arise naturally in the process of coding. It turns out that the following families of modules and torsors suffice in order to code all definable R-modules. Examples are given in [4] which indicate that one could not obtain elimination of imaginaries with a much simpler collection of sorts.

Definition 2.1. For each natural number n, the set S_n consists of the R-sublattices of K^n, that is, the free R-submodules of K^n on n generators, and $\mathcal{S} = \bigcup_{n=1}^{\infty} S_n$. For any $s \in S_n$, we define $\mathrm{red}(s) = s/\mathcal{M}s$ (the reduction of s modulo $\mathcal{M}$). For each n, let $T_n = \bigcup\{s/\mathcal{M}s : s \in S_n\}$ and $\mathcal{T} = \bigcup_{n=1}^{\infty} T_n$. Let $G = K \cup \Gamma \cup k \cup \mathcal{S} \cup \mathcal{T}$. We can now state the main theorem of [4], followed by a more algebraic-looking interpretation of it.

Theorem 2.1. *The theory of algebraically closed valued fields has elimination of imaginaries to the sorts in G.*

Corollary 2.1. *Let $(K, R, +, .)$ be an algebraically closed valued field, with valuation ring R. Then for every imaginary e of K, there is for some n a definable R-submodule of K^n with a code interdefinable with e.*

We discuss the sorts S_n and T_n further. Consider $s \in S_1$, with generator c, say. Then $s = \{rc : r \in R\} = \{x \in K : |x| \le |c|\}$. This description depends only on $\gamma = |c|$, so we could write $s = \gamma R$. Thus the elements of S_1 are precisely the closed balls around 0, and so S_1 can be identified with Γ. Each element of S_1 is isomorphic as an R-module to R, but not canonically so. Similarly, elements of S_n are isomorphic to R^n. Thus $\mathrm{red}(s)$ is isomorphic to k^n (but not canonically), and in particular is a vector space over k. As such, it inherits the good stability-theoretic properties of k, and is a stably embedded, stable structure in K^{eq}. An element of T_1 is of the form $a + \gamma\mathcal{M}$, where $a \in \gamma R$, that is, an open ball around a of radius $|a|$. One might expect that all balls would be required as elements of G; that is, that one would require also closed balls not centred at the origin, and their open sub-balls of the same radius. Part (ii) of the following lemma shows that these torsors can be identified with modules in two dimensions, and hence coded in G by more general theorems to follow.

Lemma 2.1. *(i) Let L be a definable R-submodule of K^n. Then there is a definable subtorsor L' of K^{n-1} and some $\gamma \in \Gamma$ such that $\ulcorner L \urcorner$ is interdefinable over $\emptyset$ with the pair $(\ulcorner L' \urcorner, \gamma)$. (ii) Let L' be a subtorsor of K^{n-1}. Then there is an R-submodule L of K^n such that $\ulcorner L' \urcorner = \ulcorner L \urcorner$.*

The proof of (ii) is a simple linear algebra trick: $\ulcorner L' \urcorner$ is interdefinable with a code for the submodule A of K^n generated by $\{1\} \times L'$, since $A \cap (\{1\} \times K^{n-1}) = \{1\} \times L'$. The proof of (i) requires some work on extensions of homomorphisms on modules. It uses pseudo-convergent sequences and Robinson's model completeness for algebraically closed valued fields. It also uses a standard device: we parse a definable module L of K^n as the graph of a definable homomorphism from an R-submodule A of K to a quotient K^{n-1}/T; here $A := \pi_1(L)$, where $\pi_1 : K^n \to K$ is projection to the first coordinate, and $\ker(\pi_1) \cap L = \{0\} \times T$. The methods here also yield that if $\pi : K^n \to K^m$ is any projection, and $A \in S_n$, then $\pi(A) \in S_m$ (Lemma 2.2.7 of [4]). Linear algebra also gives us another way of talking about the sorts $\mathcal{S}$ and $\mathcal{T}$ which is extremely useful when studying functions from Γ to these sorts. Let $B_n(K) \subset GL_n(K)$ be the group of invertible upper triangular matrices over K, and $B_n(R)$ be the corresponding subgroup of $GL_n(R)$ (where inverses are required to be over R). Let $\mathrm{TB}(K)$ be the set of triangular bases of K^n, that is, bases $(v_1, \ldots, v_n)$ where $v_i \in K^i \times (0)$ (i.e., the last $n-i$ entries of v_i are zero). An element $a = (v_1, \ldots, v_n) \in \mathrm{TB}(K)$ can be identified with an element of $B_n(K)$, with v_i as the i^{th} column. Now $B_n(R)$ acts on $B_n(K) = \mathrm{TB}(K)$ on the right. Two elements M, M' of $\mathrm{TB}(K)$ generate the same R-module in S_n precisely if there is some $N \in GL_n(R)$ with $MN = M'$, and as $M, M' \in B_n(K)$, we must have $N \in GL_n(R) \cap B_n(K) = B_n(R)$. This gives an identification of S_n with the set of orbits of $B_n(R)$ on $\mathrm{TB}(K)$. Equivalently, S_n can be identified with the set of left cosets of $B_n(R)$ in $B_n(K)$. This is a natural way of regarding S_n as a quotient of a power of K by a $\emptyset$-definable equivalence relation. We wish also to treat T_n as a finite union of coset spaces. For each $m = 1, \ldots, n$, let $B_{n,m}(k)$ be the set of elements of $B_n(k)$ whose m^{th} column has a 1 in the m^{th} entry and other entries zero. Let $B_{n,m}(R)$ be the set of matrices in $B_n(R)$ which reduce (coefficientwise) modulo $\mathcal{M}$ to an element of $B_{n,m}(k)$. Let $e \in S_n$, and put $V := \text{red}(e)$. We may put $e = aB_n(R)$ for some $a = (a_1, \ldots, a_n) \in \text{TB}(K)$ (so e is the orbit of a under $B_n(R)$, or the left coset $aB_n(R)$ where a is regarded as a member of $B_n(K)$). There is a natural filtration

$$\{0\} = V_0 < V_1 < \ldots < V_{n-1} < V_n$$

of V, where V_i is the k-subspace of red(e) spanned by $\{\text{red}(a_1), \ldots, \text{red}(a_i)\}$ (here $\text{red}(a_j) = a_j + \mathcal{M}e$). Let TB$(V)$ be the set of triangular bases of V, that is, bases $(v_1, \ldots v_n)$ where $v_i \in V_i \setminus V_{i-1}$. Now $B_n(k)$ acts sharply transitively on TB(V) on the right, with

$$(v_1, \ldots, v_n)(a_{ij}) = (a_{11}v_1, a_{12}v_1 + a_{22}v_2, \ldots, \Sigma_{i=1}^n a_{in}v_i).$$

For each $i = 0, \ldots, n$, put $O_i(V) = V_i \setminus V_{i-1}$ (so $O_0(V) = \{0\}$). It is easily verified that two elements of TB(V) are in the same orbit under $B_{n,m}(k)$ precisely if they agree in the m^{th} entry. Thus, $O_m(V)$ can be identified with $\text{TB}(V)/B_{n,m}(k)$, and $V \setminus \{0\}$ with $\bigcup_{m=1}^n \text{TB}(V)/B_{n,m}(k)$. From the last two paragraphs, it follows that if $M, M' \in \text{TB}(K)$, then they are $B_{n,m}(R)$-conjugate (i.e. there is $N \in B_{n,m}(R)$ with $MN = M'$) precisely if they generate the same lattice, and their images modulo the maximal ideal are $B_{n,m}(k)$-conjugate. This holds precisely if they generate the same lattice, and their m^{th} entries are the same element of T_n. The identification of TB(K) with $B_n(K)$ now yields the following lemma.

Lemma 2.2. *For each $n > 0$, there is a $\emptyset$-definable bijection between T_n and $\cup_{m=1}^n B_n(K)/B_{n,m}(R)$.*

The main purpose of [4] is to prove elimination of imaginaries to the sorts in G. There is a reasonable language in which the many-sorted structure G eliminates quantifiers. In addition to rather obvious symbols connecting up the sorts, and the natural algebraic relations, functions, constants on K, k, Γ, one needs predicates for certain $\emptyset$-definable k-Zariski closed sets in products

$V_1 \times \ldots \times V_m$, where $V_i = \mathrm{red}(s_i)$ (s_i some lattice) so V_i is a k-vector space. Each such V_i also has its k vector space structure, expressible by functions $T_n \times T_n \to T_n$ and $k \times T_n \to T_n$. There is also a notion of a *generic basis* of a lattice $s \in S_n$, and for each formula $\varphi(X)$ (where X is an n^2 tuple of field variables) we need a formula ${}^*\varphi(y)$ (y a lattice variable), so that ${}^*\varphi(s)$ holds precisely if $\varphi(a_1, \ldots, a_n)$ holds where $a_i \in K^n$ for each i and $(a_1, \ldots, a_n)$ is a generic basis of s (this is well-defined). The full details of the language are a little more complicated than this, and can be found in [4, Section 3.1]

2.3. Unary sets

The result of Holly that definable subsets of K are (canonically) Boolean combinations of balls means that definable sets in one variable are much easier to understand than sets in many variables. It is therefore often useful to treat tuples elementwise. In order to handle the sorts S_n and T_n, we develop the notion of *unary set*, which slightly generalises that of a one-variable definable (or ∞-definable) set in K. We first have to define several pieces of terminology. A *definable 1-module* is an R-module (living in K^{eq}) which is definably isomorphic to a quotient of one definable R-submodule of K by another. It will be definably isomorphic to one of $\gamma R/\delta R$, $\gamma R/\delta \mathcal{M}$, $\gamma \mathcal{M}/\delta R$, $\gamma \mathcal{M}/\delta \mathcal{M}$, $K/\delta R$ or $K/\delta \mathcal{M}$, where $\gamma, \delta \in \Gamma$ with $0 \leq \delta \leq \gamma$ (and in fact we may always assume $\gamma = 1$). A *definable 1-torsor* is a definable torsor of a definable 1-module. An *∞-definable* 1-torsor is an intersection of a chain of definable 1-torsors. A *1-torsor* is a definable or ∞-definable 1-torsor. If C is a set of parameters, then a *C-1-torsor* is a definable or ∞-definable 1-torsor for which the parameters come from C; we do not here require that there be any C-definable isomorphism with, say, $\gamma R/\delta R$. We will say that a 1-torsor is *closed* if it is definably isomorphic to a torsor of a module which is a quotient of R; it is *open* if it is definably isomorphic to a torsor of a module which is a quotient of $\mathcal{M}$ or K. In practice, we only ever need to work with 1-torsors which are torsors of R-submodules of quotients of K^n for some n. There is a corresponding notion of *subtorsor* of a 1-torsor, and of the *radius* of a subtorsor. Notice that if $\gamma < |a|$ then the closed ball s around a of radius γ is a closed 1-torsor of the 1-module $\gamma R = \gamma R/0R$. But s is also an element of the 1-module $\gamma' R/\gamma R$, where $|a| = \gamma'$. Thus we can consider a ball both as a 1-torsor and as an element of a 1-torsor, as needed.

Definition 2.2. *A* unary set *is a 1-torsor or an interval* $[0,\alpha)$ *in* Γ*, where* $\alpha \in \Gamma \cup \{\infty\}$*. A* C-unary set *is a unary set (possibly* ∞*-definable) where the parameters may be chosen from* C*. A* unary type *over* C *is the type of an element of a* C*-unary set.*

Definition 2.3. Let $C \subset K^{\mathrm{eq}}$ be a set of parameters. Let U be an $\mathrm{acl}(C)$-unary set and $a \in U$. Then a is *generic in* U *over* C if a lies in no $\mathrm{acl}(C)$-unary proper subset of U (in particular, if U is a 1-torsor, in no proper $\mathrm{acl}(C)$-subtorsor of U).

It is straightforward to prove that if a, b are generic in a unary set over C, then $a \equiv_{\mathrm{acl}(C)} b$; the essential point here is just that definable subsets of K are Boolean combinations of balls, and a corresponding property is inherited by 1-torsors. Thus, we may talk of *the generic type of* U (over C) as the type of an element of U which is generic in U over C. Furthermore, if a is an element of a C-unary set U, and $C = \mathrm{acl}(C) \subset K^{\mathrm{eq}}$, then a realises the generic type over C of a unique unary subset of U; it is given as the intersection of the set of all C-definable unary subsets of U containing a. Without the assumption that $C = \mathrm{acl}(C)$ it might be the case that this intersection V is closed, but a lies in one of finitely many (more than one) conjugate open unary subsets of V of radius $\mathrm{rad}(V)$. In this case a would not realise a generic type of a unary set over C. The following proposition is easy to prove. Its converse also holds if C is a sufficiently saturated model.

Proposition 2.4. *Let* C *be a set of parameters and* p *be the generic type of a* C*-1-torsor* U*. If* U *is definable then* p *is definable.*

We mention one easy but important lemma which ensures that generic types of closed 1-torsors are orthogonal to Γ (in a sense to emerge in Section 3.2) and have good stability properties.

Lemma 2.3. *Let* M *be a model,* $C \subset M$*, and* U *be a* C*-definable closed 1-torsor in* M*. Then if* a *is generic in* U *over* M*, we have* $\mathrm{dcl}(M) \cap \Gamma = \mathrm{dcl}(M \cup \{a\}) \cap \Gamma$.

The idea here is that if a increased Γ, one could define a non-constant function from a strongly minimal set (essentially, the reduction of U by $\mathcal{M}$) to the ordered set Γ. The lemma is false if U is open. For then, as M is a model we may identify U with an open ball, and M will contain a field element b of U; then $|b-a|$ will lie in $\mathrm{dcl}(Ma) \cap \Gamma$ but not in $\mathrm{dcl}(M)$. Finally, the following proposition ties the unary sets back to our sorts $\mathcal{S} \cup \mathcal{T}$. It enables us to use Proposition 2.2 to prove elimination of imaginaries, and is also central to the development of generic independence in Section 3.2.

Proposition 2.5. *Let $s \in G$. Then there is a sequence $(a_1, \ldots, a_m)$ from G such that $\mathrm{dcl}(s) = \mathrm{dcl}(a_1, \ldots, a_m)$, and for each $i = 1, \ldots, m$, a_i is an element of a unary set defined over $\mathrm{dcl}(a_j : j < i)$.*

We call the sequence $(a_1, \ldots, a_m)$ a *unary code* for s. To see the idea of the proof, let $A \in S_n$, and let $\pi_{n-1} : K^n \to K^{n-1}$ be the projection to the first $n-1$ coordinates. Let $B_{n-1} := \pi_{n-1}(A)$, let A_{n-1} be the kernel of A under the projection π^{n-1} of K^n to the last coordinate, and write A_{n-1} as $A'_{n-1} \times \{0\}$. Then $A'_{n-1} \leq B_{n-1}$. Also, let $B_1 := \pi^{n-1}(A)$ and put $\ker(\pi_{n-1}) := \{0\}^{n-1} \times A''_1$. Now each of $B_{n-1}, A'_{n-1}, B_1, A''_1$ is a lattice, so by induction has a unary code. Let Y be the set of lattices which correspond via projections to $B_{n-1}, A'_{n-1}, B_1, A''_1$ in the same way that A does. It suffices to check that Y has the structure of a subset of a 1-torsor over codes for $B_{n-1}, A'_{n-1}, B_1, A''_1$. The point here is that Y is canonically identifiable with a subset of $\mathrm{Hom}_R(B_{n-1}/A'_{n-1}, B_1/A''_1)$, and the latter is isomorphic as an R-module to $R/\alpha R$ for some $\alpha \in \Gamma$ with $\alpha < 1$. We could now get a unary code for $s = \ulcorner A \urcorner$ by concatenating unary codes for B_{n-1}, A'_{n-1}, B_1, A''_1 (which exist by induction on n) followed by the element s of $\mathrm{Hom}_R(B_{n-1}/A'_{n-1}, B_1/A''_1)$. We conclude this section with a proposition which illustrates how modules arise in coding arguments. This is a piece of the argument that finite sets of tuples from G are coded. The proof also is representative of the changes that appear in the move from a field to a valued field. In a pure field, a finite set is coded by the tuple of coefficients of the polynomial which has the elements of the set as its roots. In the valued field, the set of balls is coded by the module of polynomials which have the elements of the balls as their roots.

Proposition 2.6. *Let $F = \{s_1, \ldots, s_m\}$ be a set of closed balls all of radius γ and distance $\delta > \gamma$ apart. There is a definable R-module J^F such that* $\mathrm{dcl}(\ulcorner F \urcorner) = \mathrm{dcl}(\ulcorner J^F \urcorner, \gamma, \delta)$ *and hence a code for J^F, together with γ, δ, gives a code for F.*

Proof – Let $S = \bigcup_{i=1}^m s_i$ (a subset of K), and let J^F be the set consisting of one variable polynomials

$$\{Q \in K[X] : \deg(Q) \leq m \wedge \forall x \in S(|Q(x)| \leq \delta^{m-1}\gamma)\}.$$

Then J^F is a definable R-submodule of K^{m+1}. Now J^F, γ, δ are clearly definable from $\ulcorner F \urcorner$. We must show that F is recoverable from J^F, γ, δ. For this, it suffices to check that if $Q \in K[X]$ is monic of degree m, then $Q \in J^F$ if and only if Q has a root in each s_i. In one direction, suppose that Q has a root α_i in each s_i. Then $Q(X) = \Pi_{i=1}^m (X - \alpha_i)$. Suppose $x \in S$, with say $x \in s_1$. Then $|x - \alpha_1| \leq \gamma$, and $|x - \alpha_i| = \delta$ for $i = 2, \ldots, m$. Hence $|Q(x)| \leq \delta^{m-1}\gamma$. In the other direction, suppose that $Q \in J^F$ is monic of degree m and has roots $\beta_1, \ldots, \beta_m$ (listing repeated roots according to multiplicity). Then for all $j = 1, \ldots, m$ there is i such that β_i lies at distance less than δ from (all elements of) s_j. For otherwise there is some s_j so that all β_i are at distance at least δ from s_j; then if $x \in s_j$, we have $|Q(x)| \geq \delta^m$, a contradiction. Hence, after relabelling, we may assume that for each i and all $x \in s_i$, $|\beta_i - x| < \delta$. Thus, if $i \neq j$ and $x \in s_j$, we have $|\beta_i - x| = \delta$. Now choose $x \in S$, with $x \in s_i$ say. Then $|Q(x)| = \Pi_{i=1}^m |x - \beta_i| = \delta^{m-1}|\beta_i - x|$. As $Q \in J^F$, this forces $|\beta_i - x| \leq \gamma$, and hence $\beta_i \in s_i$, as required. □

2.4. Coding of modules and the k-internal sets

An algebraically closed valued field K contains its stable residue field as a stably embedded sort in K^{eq}. The residue field has elimination of imaginaries, and we would like to lift this to code definable sets which are in some sense close to the residue field. We begin by defining a multi-sorted structure which consists of an infinite collection of finite dimensional vector spaces over k. In particular, each sort is also stable and stably embedded. This structure plays an important role both in the proof of elimination of imaginaries, and in the later independence theory.

Definition 2.4. For any parameter set C, we denote by $\mathrm{Int}_{k,C}$ a many-sorted structure whose sorts are the k-vector spaces $\mathrm{red}(s)$ where $s \in \mathrm{dcl}(C) \cap \mathcal{S}$. Each sort $\mathrm{red}(s)$ is equipped with its k-vector space structure, along with any other C-definable relations as $\emptyset$-definable relations.

This collection of vector spaces is closed (up to C-definable isomorphism) under direct sums, tensor products, and duals, essentially because of corresponding closure conditions on the collection of C-definable lattices. The restriction of $\mathrm{Int}_{k,C}$ to any finite collection of sorts will have finite Morley rank. Essentially, it is the k-internal structure over C (see below).

Theorem 2.2. *Let $C \subset K^{\mathrm{eq}}$. The structure $\mathrm{Int}_{k,C}$ has elimination of imaginaries.*

The proof of this has two main steps. The first is a proof that if $V = \mathrm{red}(A)$ $(A \in S_n)$ is a sort in $\mathrm{Int}_{k,C}$, then every definable subspace of V is coded in $\mathrm{Int}_{k,C}$. This uses exterior powers to reduce to the case of 1-spaces. For the second step, observe that the notion of Zariski closed set in V^m is independent of the choice of basis. Since any definable subset of V^m is a Boolean combination of Zariski closed sets, it suffices to show that any Zariski closed set Y in V^m is coded by a tuple. We may arrange that $m = 1$. Now let $S(V) = k \oplus V^* \oplus \Sigma_{i=2}^{\infty} \mathrm{Sym}^i(V^*)$, where Sym^i is the i^{th} symmetric power. Then elements of $S(V)$ induce functions $V \to k$ independently of any choice of basis, and Y is determined by the ideal in $S(V)$ which vanishes on Y. This ideal is determined by its intersection U with some $S^m(V) = k \oplus V^* \oplus \Sigma_{i=2}^{m} \mathrm{Sym}^i(V^*)$. Let U' be the pullback of U to $T^m(V) := k \oplus V^* \oplus \Sigma_{i=2}^{m} \otimes^i (V^*)$. By the above closure properties of $\mathrm{Int}_{k,C}$, $T^m(V)$ is a sort in $\mathrm{Int}_{k,C}$. As U' is a subspace of $T^m(V)$ it is coded in $\mathrm{Int}_{k,C}$, and hence so is Y. Of course, $\mathrm{Int}_{k,C}$ consists of elements from k and the sorts $\mathcal{T}$ of our language. The proof of the above theorem also enables us to conclude that the chosen sorts do suffice to code all definable R-submodules of K^n, as stated in the following lemma.

Lemma 2.4. *(i) Every definable R-subtorsor of K^n is coded in G. (ii) If C is any set of parameters, and A is any C-definable R-submodule of K^n, then the elements of $\mathrm{red}(A)$ are coded in $\mathrm{Int}_{k,C}$.*

In fact, $\mathrm{Int}_{k,C}$ has still more coding power. We say that a definable set D is *k-internal* if there is a finite set $F \subset G$ such that $D \subset \mathrm{dcl}(kF)$. Clearly, if $s \in \mathcal{S}$ then $D = \mathrm{red}(s)$ is k-internal; just choose F to be any basis of $\mathrm{red}(s)$ over k. Also, it is straightforward to show that an infinite C-definable k-internal set D is contained in a finite union of sets of the form $\mathrm{red}(s_1)\times\ldots\times\mathrm{red}(s_m)\times F$ where $s_1,\ldots,s_m$ are $\mathrm{acl}(C)$-definable elements of $\mathcal{S}$ and F is a C-definable finite set of tuples from G. Thus, D is almost a subset of $\mathrm{Int}_{k,C}$. The following theorem shows that elements of D are coded in $\mathrm{Int}_{k,C}$ over C, where C is regarded as a set of constants. In particular, finite sets are coded in $\mathrm{Int}_{k,C}$ over C. This does not give immediately that finite sets are coded absolutely, since there remains the problem of coding C.

Theorem 2.3. *Let D be a C-definable k-internal subset of K^{eq}. Then $D \subset \mathrm{dcl}(C \cup \mathrm{Int}_{k,C})$.*

We remark that if $C = \mathrm{acl}(C \cap K)$, then every C-definable lattice s in K^n has a free basis in C, and hence $\mathrm{red}(s)$ is C-definably isomorphic to k^n. In this case, one can work just with k rather than $\mathrm{Int}_{k,C}$.

2.5. Coding of functions

We first recall the notion of germ (in a general setting, not just ACVF). Let $C = \mathrm{acl}(C)$ be a set of parameters, let M be a model containing C, and let p be a type over M which is definable over C, with solution set P. Let f be a C-definable function whose domain contains P. Suppose that $f = f_a$ is defined by the formula $\varphi(x,y,a)$ (so $f_a(x) = y$). If $a, a' \in M$, we say that $f_a, f_{a'}$ *have the same germ on P*, or the same *p-germ*, if the formula $f_a(x) = f_{a'}(x)$ lies in p. By the definability of p, the equivalence relation 'has the same germ' is definable, indeed, definable over any parameter set over which p is definable. Hence, the *germ of f on P* (which is defined to be the equivalence class of φ-definable functions with the same germ) lies in M^{eq}. Furthermore, up to interdefinability this germ is independent of the choice of φ. Since the germ of f is in M^{eq} we can talk of it being coded by a tuple c, as usual. In the above setting, we say that a code c for the germ of f is *strong* if p is definable over c and there is a c-definable function g such that the formula $f(x) = g(x)$ is in

p. In stable theories, codes for germs are always strong, and this is important in group construction arguments. A number of arguments in [4] and [5] use strong codes, usually for functions on types with good stability properties (such as types closely related to $\mathrm{Int}_{k,C}$). In general, functions on definable types in ACVF do not have strong codes. An example [4, Remark 3.3.1] is the germ, on the generic type of the maximal ideal $\mathcal{M}$, of the function $x \mapsto B_{<|x|}(c)$, where c is chosen generically in R. The germ is coded by $B_{<1}(c)$, but there is no $B_{<1}(c)$-definable function with the same germ on $\mathcal{M}$. If t is a definable 1-torsor, and M is a model containing t, then the generic type p of t over M is definable, by Proposition 2.4. If f is a definable function on t, then we will refer to the *germ of f on t* for the p-germ of f. We also may talk of an M-definable function g having the same germ on p as f, even if g is not φ-definable. By this we mean again that the formula $f(x) = g(x)$ is in p. If the type p is not definable, then the equivalence relation 'has the same germ on P' still makes sense, but we avoid talking of the 'germ of f on P', as this is not an interpretable object. By Proposition 2.4, if p is the generic type over a model of a unary set U which is not definable, then p is not definable. Suppose $U = \bigcap_{i \in I} t_i$, where $\{t_i : i \in I\}$ is a strictly decreasing chain of C-definable 1-torsors with no least element. Suppose f, g are definable functions. Then it is easily shown that f and g have the same germ on P if and only if for sufficiently large i, they have the same germ on t_i. We prove a sequence of theorems about coding germs of definable functions on unary sets. The strength of the conclusion depends on which kind of unary set the domain is. For example, the germ of a function defined on an open 1-torsor set will not in general be strongly coded (see the above example), whereas we prove that the germ of a function on a closed 1-torsor is strongly coded. This culminates in Proposition 2.8, which is what is required for Black Box 3.

Theorem 2.4. *Let $U \subset G$ be a closed C-unary set and $f : U \to G$ a definable function. Then (i) the germ of f on U is coded in G, and (ii) the code in G for the germ of f on U is strong.*

Perhaps the idea of the proof of Theorem 2.4(i) is seen most clearly in the case when f is a function $U \to K$, where U is a closed ball. Let p be the generic type of U over C. In this case, write $f = f_c$, let e be a code in K^{eq} for the germ of f on U, and let $B := \mathrm{dcl}(e) \cap G$. Suppose $c \equiv_B c'$, and put $f' := f_{c'}$. We must show that f, f' have the same germ on U. Let M be a model containing c, c', and let a realise $p|M$ (the generic type of U over M). We must show $f(a) = f'(a)$, and for this it suffices to show $af(a) \equiv_M af'(a)$. By Robinson's quantifier elimination, the latter will follow if we know that for each polynomial $F \in (M \cap K)[X,Y]$ and $\gamma \in M \cap \Gamma$, $|F(a, f(a))| = \gamma \leftrightarrow |F(a, f'(a))| = \gamma$ (we allow here $\gamma = 0$). Consider

$$J_f := \{F(X,Y) \in (M \cap K)[X,Y] : \text{ for any } b \models p|M, |F(b, f(b))| \leq 1\},$$

and let J_f^n consist of the polynomials in J_f of total degree at most n. If we identify each member of J_f^n with a tuple of coefficients, then J_f^n becomes an R-module. By the definability of p, J_f^n is $B \cup \mathrm{germ}_p(f)$-definable, and by Lemma 2.4 it is coded in G, so is definable over B. Since $f \equiv_B f'$, it follows that $J_f^n = J_{f'}^n$ for each n. Now suppose $a \models p|M$, and $F(X,Y) \in (M \cap K)[X,Y]$, with $|F(a, f(a))| = \delta > 0$. Since U is closed, by Lemma 2.3 above, $\delta \in M$, so as M is a model there is $d \in M \cap K$ with $|d| = \delta$. Now $d^{-1}F \in J_f^n$, for some n, so $d^{-1}F \in J_{f'}^n$, and hence $|F(a, f'(a))| \leq \delta$. Reversing f and f', we see that $|F(a, f(a))| = |F(a, f'(a))| = \delta$, as required. The proof of part (ii) is similar to several other proofs that codes for definable functions are strong. We will outline one in Section 3.1. The next result is a tool for handling functions on a 1-torsor which is not closed.

Proposition 2.7. *Let $f : \Gamma \to G$ be a definable function. (i) The function f is coded in G. (ii) Let $\gamma_0 \in \Gamma \cup \{\infty\}$. Then the germ of f on the generic type of elements of Γ immediately below γ_0 is coded in G.*

We mention two issues in the proof of (i). First, suppose h is a definable function from Γ to the set of all balls (or more generally, to subtorsors of a 1-torsor). One can show (working piecewise) that the range of f is nested, so is a definable chain of balls, $\{B_\gamma : \gamma \in I\}$ say, for some definable $I \subset \Gamma$. Now in an immediate extension of K there will be a field element a in $\bigcap(B_\gamma : \gamma \in I)$. Hence, by model completeness of algebraically closed valued fields, a can be found in K. It follows that there is a definable function $h : \Gamma \to \Gamma$ such that each B_γ is a ball of radius $h(\gamma)$ around a. The function h takes the form $\gamma \mapsto \delta\gamma^q$ for some rational q and some $\delta \in \Gamma$, since Γ is a stably embedded divisible ordered abelian group. In particular, $\ulcorner h \urcorner$ is coded by δ. The main problems in the proof of (i) arise with functions $\Gamma \to S_n$ and $\Gamma \to T_n$. The idea is to reduce to functions to a 1-torsor, and then quote the last paragraph. To handle the first case, we identify S_n with $B_n(K)/B_n(R)$, and use a sequence of algebraic normal subgroups of $B_n(K)$ with successive one-dimensional quotients to reduce to facts about functions from Γ to a 1-torsor. It can be shown that if $f : \Gamma \to S_n$ is definable, then the domain of f can be partitioned into finitely many intervals I, on each of which f has the form $\gamma \mapsto bh(\gamma)B_n(R)$, where b is a unitriangular element of $B_n(K)$ and $h(\gamma) \in D_n(K)/D_n(R)$; here $D_n(K)$ consists of the diagonal matrices in $B_n(K)$, and $D_n(R) = D_n(K) \cap B_n(R)$, so $D_n(K)/D_n(R)$ is naturally identifiable with Γ^n. The function h is easily coded in G, but some work is still needed, partly because b is not determined by f, and partly to show the intervals I can be taken to be $\ulcorner f \urcorner$-definable. Similarly, functions to T_n are treated as functions $\Gamma \to B_n(K)/B_{n,m}(R)$ for some m. The last two results very easily yield that the germ of a function on an open 1-torsor is coded in G. However, to obtain the next proposition (Black Box 3) some work is still needed. The main cases are when U is an open 1-torsor or the intersection of a chain of subtorsors of a 1-torsor.

Proposition 2.8. *Let U be a unary set, f be a definable function to G with domain containing U, and $B \subset G$ with $B = \mathrm{acl}_G(B\ulcorner f \urcorner)$. Suppose that U is (∞)-definable over B. Then there is a B-definable function g with the same germ on U as f.*

Notice that this proposition does not imply that f is strongly coded on U, as the function g is defined from a code for f, not a code for the germ of f. The conclusion of the proposition is already given by the strong coding in the above results in the cases when U is a closed 1-torsor or a subset of Γ. When U is an open 1-torsor or an ∞-definable 1-torsor, the method of the proof is to approximate U from within by closed 1-torsors and use the strong coding of f on each of these given by Theorem 2.4 to build a sequence of functions to approximate the required function g. Both parts of Proposition 2.7 are used.

2.6. Elimination of imaginaries for other valued structures.

In his PhD thesis [12], Mellor has proved an analogous elimination of imaginaries for real closed valued fields. The structure here is a real closed field K, together with a predicate for a convex non-trivial valuation ring R. By a theorem of Cherlin and Dickmann [2], the theory of such structures is complete and has quantifier elimination in a 1-sorted language with a predicate for the relation $|x| \leq |y|$. The value group Γ and residue field k are stably embedded o-minimal structures (a divisible abelian group and a real closed field respectively) and the whole field is weakly o-minimal, by [3]; that is, any definable subset of K is a finite union of convex sets. Mellor proves elimination of imaginaries with the same collection G of sorts: K, Γ, k, S_n, where elements of S_n are R-lattices in K^n, and the corresponding T_n. The proof has a similar structure, though he uses Proposition 2.1 rather than Proposition 2.2, and always works with balls rather than unary sets. Black Box 2 (coding of finite sets) is very easy with the ordering, and for Black Box 3, he uses a natural notion of *left generic type* of a ball (a notion which makes sense in any weakly o-minimal structure). The left generic type of a ball is definable. A lot of work is required to prove the analogue of Black Box 4 for left generic types of balls (or chains of balls). Essentially, given a function $K \rightarrow G$ defined on a closed ball U, Mellor finds a related function f^* on a closed ball U^*, with f^* and U^* both definable in the algebraically closed valued field K^* (a degree 2 extension of K which is identified with K^2). The germ of f^* on U^* has a strong code, by Theorem 2.4, and this code is interdefinable in K with a tuple c from G. Over c, there is a definable (in the *real* closed valued field K) function g^* with the same germ as f^* on U^*, and from g^* it is possible to define a function g which is definable from the germ of f on U, and agrees with f left generically on U. The argument is intricate, and involves also an induction on n (where f is a function to the set of torsors in K^n). This gives the real closed analogue of Theorem 2.4, and the rest of the proof of Black Box 4 proceeds much as in Section 2.5 above. Black Box 5 works for real closed valued fields, by results of Holly [7]. Hrushovski has circulated a proof of elimination of imaginaries for $\mathbf{Q}_p$, with the field sort and for each n a sort S_n for all rank n lattices over $\mathbf{Z}_p$. The proof is based on a proposition which obtains elimination of imaginaries for a theory T from the corresponding statement for a model completion of the universal part of T, possibly in a smaller language. The result for $\mathbf{Q}_p$ is thus deduced from elimination of imaginaries for ACVF. The proof that the hypotheses of the proposition hold uses many of the ideas from [4] and [5] such as strong codes of germs of functions. In particular, it is shown that in $\mathbf{Q}_p$ many types (in the many-sorted structure) have invariant extensions (see Definition 3.1 below). In [6], the authors developed the notion of *C-minimality*, introduced in [11] and more recently taken further by Simonetta (see for example [14]). Formally, a C-relation is a ternary relation on a set satisfying certain axioms suggested by the combinatorial behaviour of the set of maximal chains in a semilinearly ordered set. The notion of C-minimality is suggested by o-minimality, so a structure M with a C-relation is *C-minimal* if, for any $N \equiv M$, any definable subset of N is quantifier-free definable just from the relation C and equality. A valued field has a C-relation given by $C(x; y, z) \Leftrightarrow v(x - y) < v(y - z)$. The affine group of the field preserves C, rather as addition and multiplication by positive elements preserves the ordering in an ordered field. The main theorem of [6] was that any C-minimal valued field is algebraically closed. It is known [10] that the rigid analytic expansions of algebraically closed valued fields introduced by Lipshitz in [9] are C-minimal. It would be interesting to try to prove elimination of imaginaries (with the same collection of sorts) for the Lipshitz structure, or for other C-minimal expansions of an algebraically closed valued field. In such a structure, the residue field will still be strongly minimal, the value group will be o-minimal, and the basic genericity behaviour of unary sets (see Black Box 3) will be unchanged.

3. Independence theories

It is well-known that, in a stable theory, any notion of independence satisfying a certain list of desirable properties is equivalent to non-forking. In a theory such as ACVF, which is neither stable nor simple, the classical definition of non-forking gives no information in general. However, it is possible that other definitions of non-forking may be adapted to an algebraically closed valued field, and provide a workable notion. In [5] we study four different possible definitions of non-forking in an algebraically closed valued field, three of which we review here. In particular, we examine under what circumstances they are equivalent, and what algebraic information they give about a type. In our study of types in an algebraically closed valued field, we found that two properties play a key role. The first is the existence and possible uniqueness of an invariant type.

Definition 3.1. *Let $C = \mathrm{acl}(C)$ and let p be a type over C. A type q over a model $M \supset C$ is an* invariant *extension of p (or* $\mathrm{Aut}(M/C)$-invariant *extension of p) if $q|C = p$ and $\mathrm{Aut}(M/C)$ fixes q.*

Note that if $C \subset M$ then any type over M which is C-definable is $\mathrm{Aut}(M/C)$-invariant. When dealing with invariant extensions, we have in mind that M has a large automorphism group; but appeals to the automorphism group could be avoided, if we instead said that q is an invariant extension of p if every partial elementary map $M \to M$ is elementary over any realisation of q. In a stable theory with elimination of imaginaries, any non-forking extension to M of a type over an algebraically closed set will be invariant. The paper [8] gives a number of examples of situations where all types have invariant extensions (there, such types are called *strongly determined types*). For example, in the random graph, or in any weakly o-minimal structure, every type has an invariant extension. However, we caution that, as pointed out by Hrushovski, Lemma 2.2 of [8] appears to be false. Hence the claimed proof in Theorem 2.10 that every type in a C-minimal structure has an invariant extension is invalid. We do not know if this holds even for all C-minimal expansions of algebraically closed valued fields. A second key property is the strong coding of germs of definable functions. As shown in [4, Remark 3.3.1], a code for the germ of a definable function f on a type in a stable theory is always strong. The existence of a function defined just from the germ of f (and the parameters defining the type) is a powerful tool.

3.1. Stable domination

Our first definition of independence can be given for any complete theory which eliminates imaginaries. Let $\mathcal{U}$ be a sufficiently saturated model.

For any small set C of parameters, write St_C for the multi-sorted structure $\langle D_i, R_{ij}\rangle_{i\in I, j\in J_i}$ whose sorts D_i are the C-definable, stable, stably embedded subsets of $\mathcal{U}$. For each sort D_i, all the C-definable relations R_{ij} are included as $\emptyset$-definable sets. Notice that C itself is included in St_C as each point is a sort. The structure St_C is stable, and we use the symbol $\downarrow$ for non-forking (for subsets of St_C) in the usual sense of stability theory. For any $A \subset \mathcal{U}$, we define $\mathrm{St}_C(A) = \mathrm{dcl}(CA) \cap \mathrm{St}_C$. We will write $\mathrm{St}_C(A) = A^{\mathrm{st}}$ when there is no ambiguity about the base. Below, by $\mathrm{tp}(B/A)$ we mean the type over A of a (possibly infinite) *enumeration* of B, supposed fixed in the context.

Definition 3.2. *Define* $A \downarrow_C^{\mathrm{dom}} B$ *if* $A^{\mathrm{st}} \downarrow_C B^{\mathrm{st}}$ *and* $\mathrm{tp}(B/CA^{\mathrm{st}}) \vdash \mathrm{tp}(B/CA)$. *We say that* $\mathrm{tp}(A/C)$ *is* stably dominated *if, whenever* $B \subset \mathcal{U}$ *and* $A^{\mathrm{st}} \downarrow_C B^{\mathrm{st}}$, *we have* $A \downarrow_C^{\mathrm{dom}} B$.

An application of Beth's theorem yields that if $\mathrm{tp}(A/C)$ is stably dominated, then it has an extension to $\mathcal{U}$ which is definable over C, and hence $\mathrm{Aut}(\mathcal{U}/C)$-invariant. Invariant extensions of a type over an eq-algebraically closed set in a stable theory are unique. This yields the following.

Proposition 3.1. *Suppose that* $C \supset \mathrm{St}_C \cap \mathrm{acl}(C)$ *and* $\mathrm{tp}(A/C)$ *is stably dominated. Then* $\mathrm{tp}(A/C)$ *has a unique* $\mathrm{Aut}(M/C)$*-invariant extension to any model* M, *and if* A' *realises this extension then* $A' \downarrow_C^{\mathrm{dom}} M$.

If p is a stably dominated type over $C = \mathrm{acl}(C)$ and $A \supset C$, we can thus write $p|A$ for the restriction to A of the unique invariant extension of p to $\mathcal{U}$. A much more difficult result is that germs of definable functions on stably dominated types have strong codes (in the sense of Section 2.5). For simplicity we state it under the assumption $C = \mathrm{acl}(C)$.

Theorem 3.1. *Let* p *be a stably dominated type over* $C = \mathrm{acl}(C)$, *and let* f *be a definable function whose domain contains the set* P *of realisations of* p. *Then, working over the parameter set* C, *the* p*-germ of* f *has a strong code. Furthermore, suppose that* $f(a) \in \mathrm{St}_{Ca}$ *for all* $a \in P$. *Then the code for the* p*-germ of* f *is in* St_C.

All of our proofs of strong codes for definable functions are similar in outline. Write e for the p-germ of f, as an element of $\mathcal{U} = \mathcal{U}^{\mathrm{eq}}$. Assume f is definable by a formula with parameters b, and let $q = \mathrm{tp}(b/Ce)$. For any $b' \models q$ let f' be the function defined by the same formula as f with parameters b'. The main step is to show:

$$\text{if } a \models p|Cf \text{ and } a \models p|Cf' \text{ then } f(a) = f'(a). \quad (*)$$

This is stronger than saying that e is a code for the germ of f – the latter implies only that if $a \models p|Cff'$ then $f(a) = f'(a)$. If the whole structure is stable then $(*)$ is easy: just choose $b'' \models q$ with $b'' \downarrow_C bb'a$, put $f'' := f_{b''}$, and observe that $a \downarrow_C ff''$ and $a \downarrow_C f'f''$, so $f(a) = f''(a) = f'(a)$. In our case, extra argument is needed to prove $(*)$, since the existence of such a b'' is not evident. The condition $(*)$ for all $f' \in \mathrm{tp}(f/Ce)$ yields that $f(a) \in \mathrm{dcl}(Cea)$. Hence, there is an $\mathrm{Aut}(\mathcal{U}/Ce)$-invariant function h defined on $p|Ce$ such that $h(a) = f_{b'}(a)$ for some (in fact, any) $b' \models q$ with $a \models p|Cb'$. Thus, there is a formula $\varphi(x,y)$ over Ce with $\varphi(a, f(a))$, and we may adapt φ to ensure that $\varphi(x,y)$ is the graph of a function g. Then g is Ce-definable, and agrees with h on $p|Ce$, as required. The final assertion in the theorem is used repeatedly. It is proved by a stability-theoretic argument with weight: one first finds $a_1, \ldots, a_n$ realising p such that $e \in \mathrm{dcl}(a_1, \ldots, a_n, f(a_1), \ldots, f(a_n))$. These results on strong codes can sometimes be proved without stable domination. One needs that p is definable, that $f(a) \in \mathrm{St}_{Ca}$ for all $a \models p$, and also the technical condition (BS): if a is a finite tuple then any chain of sets of the form $B = \mathrm{dcl}(B) \cap \mathrm{St}_C$ between St_C and $\mathrm{St}_C(a)$ has length at most $|T|^+$, where T is the ambient theory. The latter holds for algebraically closed valued fields, by a Morley rank argument. Using strong codes, it is possible to establish several basic facts about stable domination.

Proposition 3.2. *(i) Suppose $C \subset B$, $\mathrm{tp}(A/C)$ is stably dominated, and $A \downarrow_C^{\mathrm{dom}} B$. Then $\mathrm{tp}(A/B)$ is stably dominated, $\mathrm{St}_C(AB) = \mathrm{St}_C(A^{\mathrm{st}}B^{\mathrm{st}})$, and if $\mathrm{tp}(B/C)$ is stably dominated then $\mathrm{St}_B(A) = \mathrm{dcl}(BA^{st}) \cap B^{\mathrm{st}}$. (ii) If $\mathrm{tp}(A/C)$ and $\mathrm{tp}(B/CA)$ are stably dominated, then so is $\mathrm{tp}(AB/C)$.*

When the theory is that of algebraically closed valued fields, the structure St_C is essentially $\mathrm{Int}_{k,C}$; that is, every element of St_C is coded by a tuple of $\mathrm{Int}_{k,C}$. Thus, a type should be dominated by its stable part if there is no interaction with the obviously unstable part, namely the value group. The following theorem formalises this.

Theorem 3.2. *Let $T \models ACVF$. Let p be an invariant extension of* $\mathrm{tp}(a/C)$ *to* $\mathcal{U}$*, and suppose that for any model $M \supset C$ and any $a' \models p$, $\Gamma(Ma') = \Gamma(M)$. Then* $\mathrm{tp}(a/C)$ *is stably dominated.*

3.2. Generic independence, orthogonality and domination

We now return to ACVF. The notion which we below call *generic independence* has serious flaws; nevertheless it is the easiest to work with, and also has properties which make it seem the most fundamental of the definitions of non-forking that we examine here. The motivation for the definition comes from the following situation in an algebraically closed field. A point a in an irreducible variety V defined over parameters C is *generic* if a does not lie in any lower-dimensional variety defined over C. Also $a \downarrow_C B$ if a remains generic in V over the additional parameters B. In an algebraically closed valued field, algebraic closure is the same as in the pure field, and hence has the exchange property and gives a notion of dimension and independence. However, this independence (field-theoretic algebraic independence) is insensitive to the valuation. The basic building blocks for definable sets are the unary sets, and these do not have a good dimension theory, since there are uniformly definable chains of unary sets. Therefore, when defining *generic* we refer simply to a proper unary subset, instead of a lower-dimensional set. This is the motivation for the definition of generic in Definition 2.3 and hence the following definition of $\downarrow^g$. It is a potentially useful notion in any C-minimal structure.

Definition 3.3. Let $C \subset B \subset G$. Suppose a is generic over C in the C-unary set U. Then $a \downarrow^g_C B$ if a is generic in U over B.

The lack of a good dimension theory makes it hard to lift this definition directly to tuples, as one can in the pure algebraically closed field case. Instead we give a sequential definition. Below, we say that *a is an* acl*-generating sequence* for A over C if $A = \mathrm{acl}(Ca)$, and A is *finitely* acl*-generated over C* if such a exists.

Definition 3.4. Let A, B, C be sets, with $A \subseteq G$. For $a = (a_1, \dots, a_n) \in G$, and $U = (U_1, \dots, U_n)$, define $a \downarrow^g_C B$ *via U* ('a is generically independent from

B over C via U') to hold if for each $i \le n$, U_i is an $\mathrm{acl}(Ca_1 \ldots a_{i-1})$-unary set, and a_i is a generic element of U_i over $\mathrm{acl}(BCa_1 \ldots a_{i-1})$. We usually omit reference to U. The sequence a is said to be *unary* over C if $a \downarrow^g_C C$. We shall say $A \downarrow^g_C B$ via a, U if a is an acl-generating sequence for A and $a \downarrow^g_C B$ via U. We say $A \downarrow^g_C B$ if $A \downarrow^g_C B$ via some a, U. Finally, we say $A \downarrow^g_C B$ *via any generating sequence* if for any acl-generating sequence a from A we have $a \downarrow^g_C B$ via U for some U.

A version of the above definition would also make sense in any o-minimal (or weakly o-minimal) structure M: say $a \downarrow^g_C B$ if, for any $b \in B \cup C$ with $b < a$, there is $c \in \mathrm{dcl}(C)$ with $b \le c$ and $c < a$; then put $a_1 \ldots a_n \downarrow^g_C B$ if, for each $i = 1 \ldots n$, $a_i \downarrow^g_{C\{a_j : j<i\}} B$. In this case, the independence relation is clearly dependent on the order of the tuple. As is to be expected from a notion of independence defined sequentially, $a \downarrow^{\mathrm{g}}_C b$ in general is not symmetric and depends on the order of the tuple a. On the positive side, g-independence has good transitivity properties 'on the right': if $C \subset D \subset E$, then $a \downarrow^g_C E$ if and only if $a \downarrow^g_C D$ and $a \downarrow^g_D E$. Also, it is easy to prove by induction on the length of a tuple the existence of generically independent extensions of types. For given a finite sequence a and a set C, by Proposition 2.5 there is a finite unary sequence a' with $\mathrm{dcl}(Ca) = \mathrm{dcl}(Ca')$. Then if $B \supset C$, there is $a'' \equiv_C a'$ with $a'' \downarrow^g_C B$ (and likewise there is $B' \equiv_C B$ with $a' \downarrow^g_C B'$). It is also possible to prove uniqueness: if $C = \mathrm{acl}(C)$, a, a' are tuples with $a \equiv_C a'$ and $a \downarrow^g_C B$ and $a' \downarrow^g_C B$, then $a \equiv_{CB} a'$. The proof is inductive on n where $a = (a_1, \ldots, a_n)$, but has a subtlety to handle the case when $a_{i+1} \in \mathrm{acl}(Ca_1 \ldots a_i)$. This case is treated in Lemma 2.10 of [5], which rests on Lemma 3.4.6 of [4] (an analogue of Black Box 4). The above uniqueness gives the following result on existence of invariant extensions.

Theorem 3.3. *Let $C = \mathrm{acl}(C)$ and let A be finitely* acl*-generated by a unary sequence a over C. Let M be a model containing C. There is $A' = \mathrm{acl}(Ca')$ with $A'a' \equiv_C Aa$ such that $A' \downarrow^{\mathrm{g}}_C M$ via a', and $\mathrm{tp}(A'/M)$ is invariant under* $\mathrm{Aut}(M/C)$.

Generic independence can also be used to show there is a rich supply of definable types, in the following sense. The proof is a rather elementary induction on the length of a unary code.

Theorem 3.4. *Let $C = \mathrm{acl}(C)$ be a subset of a model M. Let Σ be the set of n-types over C which have an extension to a type over M which is C-definable. Then Σ is dense in the Stone space $S_n(C)$.*

In Proposition 3.1, we asserted that a stably dominated type has a unique invariant extension to a model. Hence, for stably dominated types, $\downarrow^{\mathrm{g}}$ and $\downarrow^{\mathrm{dom}}$ are equivalent. In order to work with both notions, we need a more algebraic condition for a type to be stably dominated. This is suggested to us both by the hypothesis of Theorem 3.2, and by the intuition that a type in an algebraically closed valued field should be governed by its relationship to the residue field and value group.

Definition 3.5. If $C = \mathrm{acl}(C)$ and $a = (a_1, \ldots, a_n) \in G$ is unary over C, we say that $\mathrm{tp}(a/C)$ is *orthogonal* to Γ (written $\mathrm{tp}(a/C) \perp \Gamma$) if, for any model M with $a \downarrow^{g}_{C} M$ and $C \subseteq \mathrm{dcl}(M)$, we have $\mathrm{dcl}(M) \cap \Gamma = \mathrm{dcl}(Ma) \cap \Gamma$. We define $\mathrm{tp}(A/C) \perp \Gamma$ to hold if $A = \mathrm{acl}(Ca)$ for some a with $\mathrm{tp}(a/C) \perp \Gamma$.

If a is generic in a unary set U over C, then $\mathrm{tp}(a/C) \perp \Gamma$ if and only if U is closed. The right-to-left direction follows essentially from Lemma 2.3. For the left-to-right direction, observe that if U is not closed then there is a model $M \supset C$ containing an element b of U, and argue as in the remarks following 2.3. Theorem 3.2, along with a partial converse, yield the following.

Theorem 3.5. *Let $C = \mathrm{acl}(C)$, and let a be a tuple from G. Then $\mathrm{tp}(a/C)$ is stably dominated if and only if it is orthogonal to Γ.*

From Theorem 3.2 (and a little bit more) we have that if $\mathrm{tp}(a/C)$ is orthogonal to Γ, then $\mathrm{tp}(\mathrm{acl}(Ca)/C)$ is stably dominated. It follows from this, Proposition 3.1 and Theorem 3.3 that the definition of orthogonality for a set A is independent of the acl-generating sequence over C: that is, if $\mathrm{acl}(Ca) = \mathrm{acl}(Ca')$, then $\mathrm{tp}(a/C) \perp \Gamma$ if and only if $\mathrm{tp}(a'/C) \perp \Gamma$. ¿From the above remarks and the symmetry of $\downarrow^{\mathrm{dom}}$, we obtain the following proposition. In particular, we see that the assumption of orthogonality to Γ (on either the left or the right of $\downarrow$) is enough to guarantee that $\downarrow^{\mathrm{g}}$ is symmetric and independent of the choice of generating sequence.

Proposition 3.3. *Suppose $C = \mathrm{acl}(C)$, and A, B are finitely* acl*-generated over C. Suppose also* $\mathrm{tp}(A/C) \perp \Gamma$. *Then the following are equivalent. (i)* $\mathrm{St}_C(A) \downarrow_C \mathrm{St}_C(B)$. *(ii) $A \downarrow_C^{\mathrm{dom}} B$. (iii) $A \downarrow_C^{g} B$. (iv) $B \downarrow_C^{g} A$.*

Using generic independence, we generalise the notion of domination by the stable part of a set, to domination by any distinguished subset.

Definition 3.6. Suppose $A' \subset A$ and both are finitely acl-generated. We say that *A is dominated by A' over C* if, for any $B \subset G$, if $A' \downarrow_C^{g} B$ via some generating sequence, then $\mathrm{tp}(B/\mathrm{acl}(CA')) \vdash \mathrm{tp}(B/CA)$.

Our strongest domination results are obtained under the extra assumption that the base is a *maximal* field, that is, has no proper immediate extensions. Under this condition, we do not need any assumption of orthogonality to get strong conclusions about domination of a field by its residue field and value group.

Theorem 3.6. *Suppose that $C < L$ are algebraically closed valued fields with C maximal and L of finite transcendence degree over C. Then (i) L is dominated by $k(L) \cup \Gamma(L)$ over C, (ii) L is dominated by $k(L)$ over $C \cup \Gamma(L)$.*

The algebraic ingredients here are variants of the following result of [1].

Proposition 3.4. *Let $C < A$ be an extension of non-trivially valued fields, with C maximal, and let V be a finite dimensional C-vector subspace of A. Then there is a basis $\{v_1, \ldots, v_k\}$ of V such that, for any $c_1, \ldots, c_k \in C$,* $|\Sigma_{i=1}^{k} c_i v_i| = \mathrm{Max}\,\{|c_i v_i| : 1 \leq i \leq k\}$.

Over a maximal field there is also a nice criterion for orthogonality to Γ (in the field sort). It implies in particular that if C is maximal, then $\mathrm{tp}(a/C) \perp \Gamma$ if and only if $\Gamma(C) = \Gamma(Ca)$. This is false without maximality: if $a \in K$ lies in a proper immediate extension of a field C, then $\Gamma(C) = \Gamma(Ca)$ but a is not generic in a closed ball over C, so $\mathrm{tp}(a/C) \not\perp \Gamma$.

Proposition 3.5. *Let C be an algebraically closed valued field, and F a maximal immediate extension of C, and $a \in K^n$. Then the following are equivalent. (i)* $\mathrm{tp}(a/C) \perp \Gamma$*; (ii)* $\mathrm{tp}(a/C) \vdash \mathrm{tp}(a/F)$*, and* $\Gamma(Ca) = \Gamma(C)$.

A major theme in all this work is the way in which instability phenomena come from the value group. We mention a further result along these lines, proved via the orthogonality to Γ theory. The proof uses the stability of $\mathrm{Int}_{k,C}$.

Proposition 3.6. *Let $C = \mathrm{acl}(C)$ and let $(a_i : i \in \omega)$ be an indiscernible sequence over $C \cup \Gamma(K)$. Then $\{a_i : i \in \omega\}$ is an indiscernible set over $C \cup \Gamma(K)$.*

One other body of results deserves mention. Suppose that C is a set of parameters. It is often useful to add elements to C to obtain a model D, in the sorts of G, of the theory of algebraically closed valued fields (possibly with trivial valuation). For example, if s is a C-definable lattice, then s has a D-definable basis, so is D-definably isomorphic to R^n. If $s \in \mathrm{dcl}(C)$ is a lattice, we say that s is *resolved* in $D \supset C$ if D contains an R-basis of s. If $t \in \mathrm{red}(s)$, then t is *resolved* in D if s is resolved in D and $D \cap K^n$ contains an element of t. We say C is *resolved* if $C = \mathrm{acl}(C)$ and all elements of C are resolved in C. If $C \subset D$ and D is resolved we say that D is a *resolution* of C. Observe that if D is a resolution of C then $D = \mathrm{acl}(D \cap K)$. Such a resolution D may not be an elementary substructure of G, for the valuation on D may be trivial. Results on resolutions are used in the proof of Theorem 3.2, and elsewhere. It is rather easy to show that any set $C \subset G$ has a resolution D such that $\Gamma(C) = \Gamma(D)$. Likewise, C has a resolution D such that $k(D) = k(\mathrm{acl}(C))$. What is not so clear, however, is that the two can be done simultaneously. We say that D is a *prime resolution* of C if D is a resolution of C and D embeds over C into any resolution of C. By the above remarks, if D is a prime resolution of C then $k(D) = k(\mathrm{acl}(C))$ and $\Gamma(D) = \Gamma(C)$.

Theorem 3.7. *Let $C \subset K$, and let E be a finite subset of G. Then $C \cup E$ has a prime resolution D, and D is unique up to isomorphism over $C \cup E$. In addition, D is atomic and minimal over $C \cup E$.*

3.3. J-independence

In a (pure) algebraically closed field, there is an ideal of polynomials naturally associated to a type:

$$I(\mathrm{tp}(a/C)) = \{f(X) \in K[X] : \mathrm{tp}(a/C) \vdash f(x) = 0\} .$$

If a is a generic point of a variety V, then $I(\mathrm{tp}(a/C))$ is precisely the ideal of polynomials which vanish on V, and hence $a \downarrow_C B$ if and only if $I(\mathrm{tp}(a/C)) = I(\mathrm{tp}(a/BC))$. This point of view on non-forking emphasises ties with algebraic geometry. A property of non-forking such as transitivity on the right is immediate, whereas symmetry poses a problem even to formulate. In the case of an algebraically closed valued field, where we need to take both the valuation and the different sorts of variables into account, we define, instead of an ideal, an R-module of polynomials $J(\mathrm{tp}(a/C))$, and then give the analogous definition of J-independence.

Definition 3.7. Let $p = p(x, v)$ be a partial type with $x = (x_1, \ldots, x_\ell)$ field variables and $v = (v_1, \ldots, v_m)$ a tuple of torsor variables with v_i ranging through $S_{n_i} \cup T_{n_i}$. Let Y be a tuple of field variables of length $n_1 + \ldots + n_m$. Define

$$J(p) = \{f(X, Y) \in K[X, Y] : p(x, v) \vdash (\forall y_1 \in v_1) \ldots (\forall y_m \in v_m)|f(x, y)| < 1\}.$$

Here, y_i is a n_i-tuple of field variables for each $i = 1, \ldots, m$. We shall often write $J(a/B)$ for $J(\mathrm{tp}(a/B))$. Note that $J(p)$ can be regarded as a collection of type-definable R-submodules of $K[X]$ (for increasingly long tuples X). Next, for a tuple a in G, and B, C lying in arbitrary sorts, define $a \downarrow_C^J B$ to hold if and only if $J(a/C) = J(a/BC)$. Finally, if $A \subseteq G$, then $A \downarrow_C^J B$ holds if $a \downarrow_C^J B$ for all tuples a from A.

To compare J-independence and g-independence requires understanding the behaviour of polynomial functions on generic points of unary sets. What we find is that, as parameters are added, the J-module will change if and only if new elements of the value group become definable. That is, the norm of a polynomial at a generic point of the type bounds its norm throughout the type, provided the type is orthogonal to Γ. In particular, we prove a maximum modulus principle (Theorem 3.8 below); it says something like that the maximum norm of a polynomial on a closed domain is taken at the boundary. We shall let $\mathrm{tp}^+(A/C)$ denote the set of all inequalities $|f(x)| \leq \gamma$ or $|f(x)| < \gamma$ which lie in $\mathrm{tp}(A/C)$, where f is a polynomial over $\mathbf{Z}$; here, A consists of field elements and x is a tuple of type variables corresponding to an enumeration of A. We write $A \downarrow^m_C B$ if $\mathrm{tp}(A/C) \cup \mathrm{tp}(B/C) \vdash \mathrm{tp}^+(AB/C)$. The latter can be seen as a kind of orthogonality condition, or a form of independence.

Theorem 3.8. *Let $C = \mathrm{acl}((C \cap K) \cup H)$, where $H \subset \bigcup_{n>0} S_n$, and let A, B be valued fields with $C \subset \mathrm{dcl}(A) \cap \mathrm{dcl}(B)$ and $\Gamma(C) = \Gamma(A)$. Suppose $A \downarrow^g_C B$. Then $A \downarrow^m_C B$.*

As an exercise with the definitions, we note the following lemma.

Lemma 3.1. *Suppose that $C \leq B$ are valued fields with C algebraically closed, and $a \in K^n$ with $a \downarrow^J_C B$. Then $a \downarrow^m_C B$.*

Proof – Suppose that b is a tuple from B and $\mathrm{tp}^+(ab/C) \vdash |f(a,b)| \leq \gamma$, where $f(X,Y) \in \mathbf{Z}[X,Y]$. Let $\delta \in \Gamma(K)$ with $\delta > \gamma$. Then there is $d \in K$ with $|d| = \delta$. Now if $a' \equiv_B a$ then $|f(a',b)| < \delta$, so $d^{-1}f(x,b) \in J(a/B)$. Hence, as $a \downarrow^J_C B$, $d^{-1}f(x,b) \in J(a/C)$. It follows that if $a' \equiv_C a$ then $|f(a',b)| < \delta$, so $|f(a',b)| \leq \gamma$. The proof with $<$ in place of $\leq$ is similar. □

Using Theorem 3.8, it is not hard to show that if $C = \mathrm{acl}(C) \subset A \cap B$ and $\mathrm{tp}(A/C) \perp \Gamma$, then $A \downarrow^g_C B \Rightarrow A \downarrow^J_C B$. For example, suppose C, A, B are fields, and $f(x,e) \in J(A/B)$. We must show $f(x,e) \in J(A/C)$. We may translate A over B to arrange that $A \downarrow^g_B e$, and hence by transitivity that $A \downarrow^g_C Be$. Theorem 3.8 now gives that $\mathrm{tp}(A/C) \cup \mathrm{tp}(Be/C) \vdash \mathrm{tp}^+(ABe/C)$, and hence $\mathrm{tp}(A/C) \vdash |f(x,e)| < 1$, as required. In fact, one can prove the following.

Theorem 3.9. *Assume C, A, B are algebraically closed sets in G, and that $C \subset A \cap B$, with either $\mathrm{tp}(A/C) \perp \Gamma$ or $\mathrm{tp}(B/C) \perp \Gamma$. Then $A \downarrow^g_C B$ if and only if $A \downarrow^J_C B$. If instead C, A, B are algebraically closed fields (so in the sort K), then each condition is also equivalent to the condition $A \downarrow^m_C B$.*

References

[1] Baur, W.: Die Theorie der Paare reell abgeschlossner Körper, *Logic and algorithmic* (in honour of E. Specker), Monographie No. **30** de L'Enseignement Mathématique, Université de Genève, Geneva (1982), 25–34.

[2] Cherlin, G. and Dickmann, M.A.: Real closed rings II: model theory, *Ann. Pure Appl. Logic* **25** (1983), 213–231.

[3] Dickmann, M.A.: Elimination of quantifiers for ordered valuation rings, *J. Symb. Logic* **52** (1987), 116–128.

[4] Haskell, D., Hrushovski, E. and Macpherson, H.D.: Definable sets in algebraically closed valued fields. Part 1: elimination of imaginaries. Preprint.

[5] Haskell, D., Hrushovski, E. and Macpherson, H.D.: Definable sets in algebraically closed valued fields. Part II: an independence theory. In preparation.

[6] Haskell, D. and Macpherson, H.D.: Cell decompositiions of C-minimal structures, *Ann. Pure Appl. Logic* **66** (1994), 113–162.

[7] Holly, J.E.: Canonical forms for definable subsets of algebraically closed and real closed valued fields, *J. Symb. Logic* **60** (1995), 843–860.

[8] Ivanov, A.A. and Macpherson, H.D.: Strongly determined types, *Ann. Pure Appl. Logic* **99** (1999), 197–230.

[9] Lipshitz, L.: Rigid subanalytic sets, *Amer. J. Math.* **115** (1993), 77–108.

[10] Lipshitz, L. and Robinson, Z.: One-dimensional fibers of rigid subanalytic sets, *J. Symb. Logic* **63** (1998), 83–88.

[11] Macpherson, H.D. and Steinhorn, C.: On variants of o-minimality, *Ann. Pure Appl. Logic* **79** (1996), 165–209.

[12] Mellor, T.: Imaginaries in real closed valued fields, Ph.D. thesis, University of Leeds (2002).

[13] Robinson, A.: *Complete theories*, North-Holland, Amsterdam (1956).

[14] Simonetta, P.: Abelian C-minimal groups, *Ann. Pure Appl. Logic* **110** (2001), 1–22.