Uniwersytet Wrocławski
Wydział Matematyki i Informatyki
Instytut Matematyczny
specjalność: matematyka teoretyczna

Tomasz Rzepecki

Smoothness of bounded invariant equivalence relations

Praca magisterska
napisana pod kierunkiem
dra hab. Krzysztofa Krupińskiego

Wrocław 2014

Oświadczam, że pracę magisterską wykonałem samodzielnie
i zgłaszam ją do oceny.

Data:.................... Podpis autora pracy:.........................

Oświadczam, że praca jest gotowa do oceny przez recenzenta.

Data:.................... Podpis opiekuna pracy:.........................

Uniwersytet Wrocławski
Wydział Matematyki i Informatyki
Instytut Matematyczny
specjalność: matematyka teoretyczna

Tomasz Rzepecki

Gładkość ograniczonych relacji równoważności

Praca magisterska
napisana pod kierunkiem
dra hab. Krzysztofa Krupińskiego

Wrocław 2014

Streszczenie w języku polskim

Tematyka pracy. Tematem pracy są ograniczone, niezmiennicze, borelowskie relacje równoważności. Dokładniej, badamy zależność między typową definiowalnością i gładkością takich relacji – typowo definiowalna relacja równoważności zawsze jest gładka, ale odwrotna implikacja na ogół nie zachodzi.

Głowne wyniki. Główne wyniki pracy są następujące:

(1) Theorem 4.6, techniczne twierdzenie pokazujące, że pewne niezmiennicze relacje równoważności nie są gładkie, udowodnione przez prostą modyfikację dowodu głównego twierdzenia z [KMS13] (podobne wyniki zostały w międzyczasie uzyskane w [KM13] innymi metodami, ale warto wspomnieć, że dowód omawianego twierdzenia został znaleziony przez autora zanim ten drugi preprint został upubliczniony). Jest ono przedstawione również w nieco prostszej formie w następujących wnioskach.

(2) Theorem 5.8, w którym staramy się przeanalizować powiązania między gładkością, typową definiowalnością i pewnymi innymi własnościami ograniczonych, niezmienniczych relacji równoważności przy pewnych dodatkowych założeniach; dowód wykorzystuje wniosek Theorem 4.6 aby pokazać że niektóre z własności są silniejsze od innych, a zestaw (oryginalnych) przykładów pokazuje, że nie ma implikacji odwrotnych.

(3) Theorem 6.2, w którym stosujemy wniosek Theorem 4.6 dla grup definiowalnych (dokładniej, Corollary 4.10) wraz z pewnymi pomysłami z [GK13] oraz [KPS13] w kontekście definiowalnych rozszerzeń grup, aby podać kryterium typowej definiowalności podgrup takich rozszerzeń, co daje w efekcie dowód ważnej hipotezy z [GK13] w przypadku przeliczalnym.

Struktura pracy. W części pierwszej omawiamy tematykę pracy i jej strukturę, wprowadzamy konwencje obowiązujące w dalszej części.

W drugiej części przypominamy bez dowodów znane i podstawowe fakty stanowiące tło dla reszty pracy.

W trzeciej części rozwijamy język ponad to, co mieści się w opublikowanych pracach, tak by zapewnić podstawy formalne dla głównych wyników. Niektóre z wprowadzonych pojęć i udowodnionych wyników stanowią folklor, są znane wśród specjalistów, lub były sugerowane we wcześniejszych pracach, ale niektóre są nowe (np. relacje orbitalne na typach).

W czwartej części wprowadzamy narzędzia potrzebne konkretnie do dowodu Theorem 4.6, zaczerpnięte z [KMS13], przedstawiamy sam dowód (wykorzystując język i niektóre fakty z trzeciej części), a także wyciągamy zeń wnioski, które stosujemy bezpośrednio w dalszej części pracy.

W piątej części pokazujemy szereg przykładów ilustrujących powiązania między rozmaitymi własnościami niezmienniczych relacji równoważności, a także dowodzimy Theorem 5.8. Interpretujemy to twierdzenie jako częściowy wynik w ramach odpowiedzi na uogólnienie hipotezy z [KPS13] na temat złożoności relacji należenia do tego samego typu Lascara.

W szóstej części korzystając z wniosku z Theorem 4.6 dowodzimy Theorem 6.2: pokazujemy, że pewne podgrupy definiowalnych rozszerzeń grup są typowo definiowalne.

Contents

Streszczenie w języku polskim 4
Tematyka pracy 4
Główne wyniki 4
Struktura pracy 4
1. Introduction 5
1.1. Preface 5
1.2. Structure of the paper 6
1.3. Conventions 7
2. Preliminaries 7
2.1. Borel cardinalities of abstract Borel equivalence relations 7
2.2. Strong types 9
2.3. Connected group components 11
2.4. Logic topology 12
3. Framework 12
3.1. Bounded invariant equivalence relations 12
3.2. Remarks on logic topology 16
3.3. Orbital equivalence relations and normal forms 17
3.4. Invariant subgroups as invariant equivalence relations 22
4. The technical theorem 24
4.1. Descriptive-set-theoretic observations 24
4.2. Generic formulas and proper types 25
4.3. Proof of the technical theorem 26
4.4. Variants of Theorem 4.6 28
5. Characterisation of smooth equivalence relations and Borel cardinalities 29
5.1. Counterexamples 30
5.2. Main characterisation theorem 32
5.3. Possible extensions of Theorem 5.8 and related questions 34
5.4. Possible Borel cardinalities 34
6. Applications to definable group extensions 35
6.1. Introduction to extensions by abelian groups 35
6.2. Main theorem for definable group extensions 36
6.3. Corollaries of Theorem 6.2 39
References 41

1. Introduction

1.1. **Preface.** This paper will concern the Borel cardinalities of bounded, invariant equivalence relations. More precisely, the main theme is the connection between type-definability and smoothness of these relations – type-definable equivalence relations are always smooth, as we will see in Fact 3.7; the converse is not true in general.

We establish solid ground for this kind of inquiry, providing proofs of various statements (some of which are folklore) that allow us to express the problem in concrete terms (in particular, we interpret some invariant sets as Borel subsets of type spaces in a consistent manner). The concept of Borel cardinality and smoothness is classical and will be introduced abstractly in section 2., whereas for bounded equivalence relations, we will provide it in section 3., extending the notions introduced in [Cas+01] and [KPS13].

There are three main results:

(1) Theorem 4.6, a technical statement showing that some invariant equivalence relations are not smooth, which is proved by a simple modification of the proof of the main result of [KMS13] (very similar results have been since shown in [KM13] using different – though not unrelated – methods, although it should be noted that the latter preprint was circulated after the proof of Theorem 4.6 presented here was found by the author); it is also presented in a more distilled form in the following corollaries;
(2) Theorem 5.8, which attempts to analyse in detail the connection between smoothness, type-definability and some other properties of bounded and invariant equivalence relations, under some additional assumptions; it uses a corollary of Theorem 4.6 to show that some of these properties are stronger than others, and several (original) examples to show that they are not equivalent.
(3) Theorem 6.2, which applies a corollary of Theorem 4.6 for definable groups (more precisely, Corollary 4.10) along with some ideas from [GK13] and [KPS13] in the context of definable group extensions, in order to give a criterion for type-definability of subgroups of such extensions, resulting in a proof of (the countable case of) an important conjecture from [GK13] in Corollary 6.9.

The motivation for this investigation comes from two directions: on one hand, it allows us to use Borel cardinality criteria to show that some objects are type-definable (as shown in Theorem 6.2). On the other hand, we consider an extension of a conjecture posed in [KPS13] about the possible Borel cardinalities of Lascar strong types – as explained in subsection 5.4 – which is in part inspired by Theorem 3.17 from [New03] (which we strengthen in a special case in Corollary 4.8).

It is assumed that the reader is familiar with basic concepts of model theory (e.g. compactness, definable sets, type-definable sets, type spaces, saturated models, indiscernible sequences) and descriptive set theory (e.g. Polish spaces, standard Borel spaces, Borel classes). Less elementary concepts will be introduced.

1.2. **Structure of the paper.** This (first) section contains the preface outlining the goals and motivations of this paper, this subsection detailing the structure of the paper, and an exposition of conventions used later on.

The second section will contain preliminaries, basic, classical facts – providing context for the sequel – divided into theme-based subsections. All it contains are all either well-established facts, or simple observations based upon them. If the reader is familiar with the subject matter, it can be safely skipped and only used as reference for facts used later.

The third section will develop the necessary framework upon which we will base the part that comes after it – the language in which we express the sequel. Many of the concepts introduced there were present or alluded to in some way before (and we will attempt to give credit where it is due in those cases), but some are original (e.g. the concept of orbital and orbital on types invariant equivalence relations). Similarly, most facts are folklore and/or motivated by previous work (which, again, will be attributed when necessary), but in many cases significantly expanded beyond their original form.

The fourth section will contain some lemmas needed for the proof of Theorem 4.6, the proof itself (mimicking the main result of [KMS13], as stated before, adapted to the more general case using the ideas from third section), as well as its immediate corollaries (which are original).

The fifth section will interpret the results of the fourth as a statement that allows us to tell if some equivalence relations are type-definable (esp. in Theorem 5.8), and as a lower bound in a general question of possible Borel cardinalities (cf. Question 7).

It will also discuss the possible extensions of this and showcase, in concrete examples, its limitations. The contents of the fifth section are, for the most part, original research.

Finally, the sixth section will apply Theorem 4.6 (in the flavour of Corollary 4.10) to the context of definable group extensions, obtaining as a corollary the proof of a conjecture from [GK13] (a paper which is also the inspiration for many of the ideas used throughout the section, which will be marked as appropriate).

1.3. **Conventions.** In the following, unless otherwise stated, we assume that we have a fixed complete, countable theory T with infinite models. (The theory may be multi-sorted, and it will, of course, vary in some specific examples.)

We also fix a monster model $\mathfrak{C} \models T$, that is, a model which is κ-saturated and (strongly) κ-homogeneous for κ a regular and sufficiently large cardinal (and whenever we say "small" we mean smaller than this κ). If we assume that there is a sufficiently large and strongly inaccessible cardinal κ, we can take for $\mathfrak{C}$ the saturated model of cardinality κ.

We assume that all parameter sets are contained in $\mathfrak{C}$, every model we consider is an elementary substructure of $\mathfrak{C}$, and every tuple is countable.

Often, we will denote by M an arbitrary, but fixed small (and usually countable) model.

For a small set $A \subseteq \mathfrak{C}$, by A-invariant we mean $\mathrm{Aut}(\mathfrak{C}/A)$-invariant.

For simplicity, whenever we mention definable, type-definable or invariant sets, we mean that they are (unless otherwise stated) $\emptyset$-definable, type-definable or invariant, respectively.

When talking about tuples of elements of $\mathfrak{C}$, we will often say that they are in $\mathfrak{C}$ (as opposed to some product of various sorts of $\mathfrak{C}$), without specifying the length, when it does not matter, or otherwise there is no risk of confusion. Likewise, we will often write $X \subseteq \mathfrak{C}$ when X is a subset of some product of sorts of $\mathfrak{C}$.

If X is some A-invariant set (esp. type-definable over A), we will denote by $S_X(A)$ the set of complete A-types of elements of X, and similarly we will sometimes omit X (or names of sorts in multi-sorted context) in $S_X(A)$, and write simply $S(A)$ instead.

Throughout the paper, formulas and types will be routinely identified with the corresponding subsets of $\mathfrak{C}$, as well as the corresponding subsets of type spaces (or points, in case of complete types). Similarly, invariant sets will be identified with subsets of type spaces and equivalent $L_{\infty,\omega}$ formulas. For example, if $X \subseteq \mathfrak{C}$ is an A-invariant set, then we will identify X with $\bigvee_{i\in I}\bigwedge_{j\in J}\varphi_{i,j}(x,A)$ (where I, J are possibly infinite index sets and $\varphi_{i,j}$ are first order formulas) if we have

$$x \in X \iff \mathfrak{C} \models \bigvee_i \bigwedge_j \varphi_{i,j}(x,A)$$

In this case we also associate with X the subset $X_A = \{\mathrm{tp}(a/A) \mid a \in X\}$ of $S(A)$; when $A = \emptyset$, and there is no risk of confusion, we will sometimes simply write X instead of $X_\emptyset$.

When metrics are mentioned, they are binary functions into $[0,\infty] = \mathbf{R}_{\geq 0} \cup \{\infty\}$ satisfying the usual axioms (coincidence axiom, symmetry and triangle inequality), but in particular, they are allowed to (and usually will) assume ∞.

2. Preliminaries

2.1. **Borel cardinalities of abstract Borel equivalence relations.** In this subsection, we introduce the concept and basic facts about the Borel cardinality in the abstract case, which is a way of expressing the complexity, or the difficulty of classification of some equivalence relations.

Theorem 3.17 will show that some equivalence relations with complicated (not type-definable) classes must necessarily have many classes, whereas Corollary 4.8 will say that with some additional assumptions, they have complicated quotients, which is expressed by the Borel cardinality. For a more comprehensive introduction to the interpretation and occurences of Borel cardinality in various fields of mathematics (as a complexity measure of moduli spaces), see e.g. the survey [Kec99] by Kechris.

Recall the notion of a standard Borel space:

Definition. A measure space (X, Σ) is called a standard Borel space if it is isomorphic (as a measure space) to $(P, \mathrm{Bor}(P))$ for some Polish space P, or equivalently, if it is isomorphic (as a measure space) to $(B, \mathrm{Bor}(B))$ for a Borel subset B of some Polish space (cf. Corollary 13.4 in [Kec95]).

Let us introduce the basic notions of the theory of Borel cardinality.

Definition (Borel reduction, Borel reducibility)**.** Suppose X, Y are standard Borel spaces and E, F are Borel equivalence relations on X, Y, respectively. We say that a Borel function $f\colon X \to Y$ is a Borel reduction of E to F if for all $x, x' \in X$ we have

$$x\ E\ x' \iff f(x)\ F\ f(x')$$

If such f exists, we say that E is Borel reducible to F, and denote it by $E \leq_B F$.

Remarks.

- If $f\colon X \to Y$ is a Borel reduction of E to F, and $g\colon Y \to X$ is a Borel section (i.e. $f \circ g = \mathrm{id}_Y$), then g is a Borel reduction of F to E (this is because the condition for Borel reducibility is of the "if and only if" form).
- If E is a Borel equivalence relation on X, and $Y \subseteq X$ is Borel, then the inclusion yields a reduction of $E{\restriction_Y}$ to E (in particular, $E{\restriction_Y} \leq_B E$).

Definition (Borel equivalence, Borel cardinality)**.** If $E \leq_B F$ and $F \leq_B E$, we say that E, F are bireducible or Borel equivalent and denote it by $E \sim_B F$. (Note that it does not, in general, imply that there is a Borel isomorphism taking E to F.) The Borel cardinality of a Borel equivalence relation E is its $\sim_B$-class.

Definition (Smooth equivalence relation)**.** A Borel equivalence relation E is called smooth if it is Borel equivalent to equality on a standard Borel space.

The classification of Borel cardinalities of smooth equivalence relations is rather simple, thanks to the classification of standard Borel spaces.

Remark. The Borel cardinality of a smooth equivalence relation is determined by the number of classes, in particular the Borel cardinalities of smooth equivalence relations are exactly those of

- $\Delta(n)$ with n a positive natural number,
- $\Delta(\mathbf{N})$,
- $\Delta(2^{\mathbf{N}})$,

where $\Delta(X)$ denotes the relation of equality on X (i.e. the diagonal in X^2).

The following is an important equivalence relation for the purpose of telling apart smooth and non-smooth equivalence relations.

Definition. $\mathbf{E}_0$ is defined as the relation on $2^{\mathbf{N}}$ of eventual equality. That is,

$$\begin{aligned}(a_n)_n\ \mathbf{E}_0\ (b_n)_n &\iff (\exists N)(\forall n > N)\, a_n = b_n \\ &\iff \{n \in \mathbf{N} \mid a_n \neq b_n\} \text{ is finite}\end{aligned}$$

The latter condition can also be stated as $\mathbf{E}_0 = \Delta(2)^{\mathbf{N}}/\mathrm{Fin}$ (i.e. $\mathbf{E}_0$ is the countable power of the relation of equality on a two-element set modulo the ideal of finite sets).

Remark. The relation $\mathbf{E}_0$ is in an important class of Borel equivalence relations induced by Borel ideals on $2^{\mathbf{N}}$: for each Borel $I \trianglelefteq 2^{\mathbf{N}}$, we have a Borel equivalence relation E_I on $2^{\mathbf{N}}$:

$$(x_n)_n \, E_I \, (y_n)_n \iff (x_n)_n - (y_n)_n \in I \iff \{n \mid x_n \neq y_n\} \in I$$

$\mathbf{E}_0$ is obtained by taking for I the ideal of finite sets.

The following facts illustrate some important properties of the partial order $\leq_B$ as well as the distinction between smooth and non-smooth equivalence relations.

Fact 2.1 ([Kan08](Theorems 5.5.1, 5.7.1, 5.7.2),[BK96](Theorem 3.4.3))**.**

- *Smooth equivalence relations form an initial segment of Borel cardinalities, i.e. if E is smooth and F is not, then $E \leq_B F$ (Silver dichotomy). (Note that it implies that any non-smooth equivalence relation has at least $\mathfrak{c} = 2^{\aleph_0}$ classes.)*
- $\mathbf{E}_0$ *is the $\leq_B$-least non-smooth equivalence relation, that is, E is not smooth if and only if* $\mathbf{E}_0 \leq_B E$ *(Harrington-Kechris-Louveau dichotomy).*
- *Borel reducibility $\leq_B$ is not a total order, that is, there are equivalence relations E, F such that $E \not\leq_B F$ and $F \not\leq_B E$. (In fact, there are continuum many pairwise incomparable relations, even among those with countable classes.)*
- *If X is a Polish space and $E \subseteq X^2$ is a G_δ equivalence relation on X, then E is smooth.*

2.2. **Strong types.** There are several notions of strong type in model theory. By a strong type we usually mean some canonical choice of equivalence relation on (tuples of elements of) $\mathfrak{C}$ refining the familiar $\equiv$ (i.e. the relation $a \equiv b \iff \mathrm{tp}(a/\emptyset) = \mathrm{tp}(b/\emptyset)$), and invariant, or a single class of such a relation. The most often considered (in order of decreasing coarseness) are the Shelah, Kim-Pillay and Lascar strong types. We will be mostly concerned with the last one. For a general introduction to the three strong types, consult [Cas+01].

The main focus of this paper will be the equivalence relations which are bounded and invariant.

Definition. An invariant equivalence relation E on (an invariant set of tuples in) $\mathfrak{C}$ is said to be bounded if it has a small number of classes, i.e. smaller than the degree of saturation of $\mathfrak{C}$.

Remarks.

- In cases considered here – that is, for countable tuples and countable theories – a bounded invariant equivalence relation can have no more than $\mathfrak{c}$ classes. (This can be seen as a consequence of Proposition 3.3.)
- An invariant equivalence E relation is bounded if and only if there is a (small) cardinal κ such that for any model M, the number of classes of E restricted to M is no greater than κ.
- A definable equivalence relation is bounded if and only if it has finitely many classes (by compactness).

The most important – in this paper – notion of strong type is the following, essentially introduced by Lascar in [Las82].

Definition (Lascar strong type)**.** Suppose a, b are tuples in the same small product of sorts of $\mathfrak{C}$. We say that a and b have the same Lascar strong type, or are Lascar equivalent (which we denote by $a \, E_L \, b$ or $a \equiv_L b$), if one of the (equivalent) conditions listed below is satisfied.

(1) there is a sequence $a = a_0, \ldots, a_n = b$ such that for each $i < n$ there is an infinite indiscernible sequence starting with a_i, a_{i+1};
(2) there is a sequence of small models $M_1, \ldots, M_m$ and automorphisms $\sigma_i \in \mathrm{Aut}(\mathfrak{C}/M_i)$ such that $\sigma_m \circ \cdots \circ \sigma_1(a) = b$;
(3) for every bounded, invariant equivalence relation E we have $a\ E\ b$ (in fact, Lascar strong type is bounded and invariant, so this means just that it is the finest bounded and invariant equivalence relation).

For the equivalence of the three definitions of Lascar strong type, see for instance [Cas+01](Definition 1.1, Facts 1.9, 1.11, 1.12 and 1.13).

Lascar equivalence can also be seen as the relation of lying in the same metric component with respect to the Lascar distance.

Definition (Lascar distance)**.** The minimal n for a sequence a_i as in the first item above, or ∞ if such a sequence does not exist, is called the Lascar distance of a and b and denoted by $d_L(a, b)$ (and this is no greater than twice the m for sequence of automorphisms as in the second item by the next fact).

Remarks.

- It is not hard to see that d_L is an $\mathrm{Aut}(\mathfrak{C})$-invariant metric (with values in $\mathbf{N} \cup \{\infty\}$).
- If we consider a graph (simple, unweighted) whose vertices are elements of $\mathfrak{C}$ (or, more generally, tuples in a fixed product of sorts) and two of them are connected by an edge if and only if they are terms in an infinite indiscernible sequence, then Lascar distance is the distance in this graph, and Lascar strong types are the connected components.
- For each n, $d_L(a, b) \leq n$ is a type-definable condition on a, b. (In particular, the (relation of having the same) Lascar strong type is F_σ, in a sense which will be explained later, in Section 3.)

Fact 2.2 ([Cas+01](Fact 1.12))**.** *If a, b have the same type over a model M, then $d_L(a, b) \leq 2$. In particular, any two elements of the same type over a model M have the same Lascar strong type.*

Kim-Pillay types will also be used, and the definition is quite similar to that of Lascar strong types (more specifically, the third of the equivalent definitions).

Definition (Kim-Pillay strong type)**.** Suppose a, b are tuples in the same small product of sorts. We say that a and b have the same Kim-Pillay strong type (denoted by $a\ E_{KP}\ b$ or $a \equiv_{KP} b$) if for any type-definable, bounded equivalence relation E we have $a\ E\ b$.

(Much like in the case of Lascar strong type, Kim-Pillay strong type is bounded and type-definable, so it is the finest bounded, type-definable (*over* $\emptyset$) equivalence relation.)

Remark. Often, we will say "Lascar strong type" when referring to Lascar equivalence $\equiv_L$ (the equivalence relation of having the same Lascar strong type). Similarly, with the relation $\equiv$ of having the same type and the relation $\equiv_{KP}$ of having the same Kim-Pillay strong type, we will say "type", meaning $\equiv$ and "Kim-Pillay type", meaning $\equiv_{KP}$.

In many cases, the two strong types coincide. If they do coincide on all tuples of small length (or, equivalently, the Lascar strong type is type-definable), we say that the theory is G-compact. In particular, it is well-known that every stable (and even every simple) theory is G-compact (see [Kim98], Proposition 13.).

Definition (Strong automorphism groups)**.** We define the Kim-Pillay strong automorphism group

$$\operatorname{Aut} \mathrm{f}_{KP}(\mathfrak{C}) = \{\sigma \in \operatorname{Aut}(\mathfrak{C}) \mid \sigma \text{ preserves each Kim-Pillay strong type setwise}\}$$

and the Lascar strong automorphism group

$$\operatorname{Aut} \mathrm{f}_{L}(\mathfrak{C}) = \{\sigma \in \operatorname{Aut}(\mathfrak{C}) \mid \sigma \text{ preserves each Lascar strong type setwise}\}$$

Fact 2.3 ([Cas+01], Facts 1.4, 1.9)**.** *Both strong automorphism groups are normal in the full automorphism group and each one acts transitively on each class of the respective equivalence relation, so in fact the strong types are their orbit equivalence relations.*

In addition, the latter group is generated by automorphisms over small models.

We will often be dealing with relations defined on sets smaller than an entire sort, so we will need the next fact.

Fact 2.4 ([Cas+01](Fact 1.4 and Corollary 1.5))**.** *The restriction of $\equiv_L$ to an invariant set is the finest among bounded, invariant equivalence relations on this set.*

Similarly, restriction of $\equiv_{KP}$ to a type-definable set in a sort (or product of sorts) S is the finest bounded, type-definable equivalence relation on this set. As a consequence, any bounded type-definable relation E (on a type-definable set) extends to a bounded type-definable relation on the entire S (given by $E \cup \equiv_{KP}$).

2.3. **Connected group components.** Mirroring the strong types are the strong connected components of groups, defined as follows.

Definition. Let B be a small set. Suppose G is a B-type-definable group. Then G^{00}_B is the smallest B-type-definable subgroup of G of bounded (small) index, and $G^{00} = G^{00}_\emptyset$.

Definition. Let B be a small set. Suppose G is a B-invariant group. Then G^{000}_B is the smallest B-invariant subgroup of G of bounded (small) index, and $G^{000} = G^{000}_\emptyset$. (The group G^{000}_B is sometimes denoted by G^∞_B, since it may be equivalently defined as the smallest subgroup of bounded index which is $\mathcal{L}_{\infty,\omega}(B)$-definable.)

Remarks.

- It is not hard to show that these groups always exist. In some cases, both groups G^{00}_B and G^{000}_B remain the same, irrespective of the chosen small set B over which G is invariant.

  Interestingly, it is actually rather common: it is always the case for G definable in NIP theories (see [Gis11]).
- Some authors reserve the names G^{00}, G^{000} for absolute connected components (in the sense that they do not depend on the small parameter set B), but they are beyond the scope of this paper, so instead we will omit the $\emptyset$ to simplify notation, just like we omit $\emptyset$ when talking about e.g. $\emptyset$-type-definable sets.

Fact 2.2 also has an analogue in strong component setting.

Fact 2.5 ([GN08], Proposition 3.4, point 1.)**.** *Suppose that G is a definable group.*

Then the group G^{000} is generated by the set:

$$\{gh^{-1} \mid g, h \in G \text{ and } g \equiv_L h\}$$

In particular, G^{000} is generated by a countable family of type-definable sets of the form (with varying $n \in \mathbf{N}$)

$$\{gh^{-1} \mid g, h \in G \text{ and } d_L(g,h) \leq n\}.$$

and it is the countable union of their compositions (which are type-definable).

The theorem below shows that the relationship between connected components and strong types is not merely superficial, and in fact the connected components are, in a way, a special case of strong types. This idea will be extended later in Proposition 3.27.

Theorem 2.6 (see [GN08](Section 3, in particular Propositions 3.3 and 3.4))**.** *If G is a definable group, and we adjoin to $\mathfrak{C}$ a left principal homogeneous space $\mathfrak{X}$ of $G(\mathfrak{C})$ (as a new sort; we might think of it as an "affine copy of G"), along with a binary function symbol for the left action of G on $\mathfrak{X}$, then the Kim-Pillay and Lascar strong types correspond exactly to the orbit equivalence relations of G^{00} and G^{000} acting on $\mathfrak{X}$. Moreover, we have isomorphisms:*

$$\mathrm{Aut}((\mathfrak{C},\mathfrak{X},\cdot)) \cong G(\mathfrak{C}) \rtimes \mathrm{Aut}(\mathfrak{C})$$
$$\mathrm{Aut\,f}_{KP}((\mathfrak{C},\mathfrak{X},\cdot)) \cong G^{00}(\mathfrak{C}) \rtimes \mathrm{Aut\,f}_{KP}(\mathfrak{C})$$
$$\mathrm{Aut\,f}_{L}((\mathfrak{C},\mathfrak{X},\cdot)) \cong G^{000}(\mathfrak{C}) \rtimes \mathrm{Aut\,f}_{L}(\mathfrak{C})$$

Where:

(1) *the semidirect product is induced by the natural action of* $\mathrm{Aut}(\mathfrak{C})$ *on* $G(\mathfrak{C})$*,*
(2) *on* $\mathfrak{C}$*, the action of* $\mathrm{Aut}(\mathfrak{C})$ *is natural, and that of* G *is trivial,*
(3) *on* $\mathfrak{X}$ *we define the action by fixing some* x_0 *and putting – for* $g \in G(\mathfrak{C})$ *and* $\sigma \in \mathrm{Aut}(\mathfrak{C})$ *–* $\sigma_g(h \cdot x_0) = (hg^{-1})x_0$ *and* $\sigma(h \cdot x_0) = \sigma(h) \cdot x_0$*.*

Remark. The isomorphisms are not canonical in general: they depend on the choice of the base point x_0. Since the strong automorphism groups are normal, the resulting subgroups of $\mathrm{Aut}((\mathfrak{C},\mathfrak{X},\cdot))$ do not depend on the choice of x_0.

2.4. **Logic topology.** Logic topology is a useful tool for studying bounded type-definable equivalence relations, though unfortunately it ceases to be effective in more general context of arbitrary bounded invariant equivalence relations. Still, it provides some insight, and we will use it later (in the final section) to prove an important corollary of the main result. It also offers an alternative view on Borel cardinality of some equivalence relations, see Corollary 3.12. For now, we will only introduce the logic topology and cite the fact that it is a compact, Hausdorff topology, and postpone further analysis until later.

Definition. (Logic topology) Whenever we have a bounded type-definable equivalence relation E on a (type-definable) set X, we put on X/E a topology (called logic topology), by declaring the closed sets to be exactly the sets whose preimages are type-definable over some small set.

Fact 2.7 ([Pil04](Lemma 2.5))**.** *If E is a bounded type-definable equivalence relation on a type-definable set X, then X/E is a compact Hausdorff space when given the logic topology.*

Remark. There is a theorem that in o-minimal theories (e.g. o-minimal extensions of real closed fields), if G is a definable group, then the group G/G^{00}, equipped with the logic topology, is actually (isomorphic as a topological group to) a real compact Lie group. For more information, see for instance [Pil04],[Pet07].

3. Framework

3.1. **Bounded invariant equivalence relations.** In this chapter, we extend the theory of Borel cardinality of Lascar strong types as considered in [KPS13] to general invariant and bounded equivalence relations, to provide a uniform way of viewing bounded, invariant equivalence relations as relations on a Polish space. From now on, all bounded equivalence relations are only defined on invariant sets of tuples of at most countable length (within some product of sorts of $\mathfrak{C}$).

Definition (Borel invariant set, Borel class of an invariant set)**.** For any invariant set X, we say that X is Borel if the corresponding subset of $S(\emptyset)$ is, and in this case by Borel class of X we mean the Borel class of the corresponding subset of $S(\emptyset)$ (e.g. we say that X is F_σ if the corresponding set in $S(\emptyset)$ is F_σ, and we might say that X is clopen if the corresponding subset of $S(\emptyset)$ is clopen, i.e. if X is definable).

Similarly if X is A-invariant, we say that it is Borel over A if the corresponding subset of $S(A)$ is (and Borel class is understood analogously).

More generally, if we want to say that a set X is Borel or is in some specific Borel class over a small set without specifying the parameters, we attach a pseudo- prefix, so e.g. we say that X is pseudo-G_δ if X is G_δ over a small set.

Remarks.

- All of the above definitions can be relativised, so e.g. $Y \subseteq X$ is relatively F_σ over A if it is the intersection of a relatively F_σ set with X.
- Notice that if A is countable and X is Borel over A, then $S_X(A)$ – endowed with the σ-algebra generated by formulas over A – is a standard Borel space.
- Since we will use adjective "Borel" and others to refer to subsets of a model, we may confuse it with Borel subsets of a standard Borel space (or just a Polish space). When such a confusion is likely to appear, we may the latter "abstract Borel sets" (though we will only actually use the term "abstract Borel equivalence relation").

We will use this descriptive-set-theoretic lemma several times.

Lemma 3.1 ([Kec95](Exercise 24.20))**.** *Suppose X, Y are compact, Polish spaces and $f\colon X \to Y$ is a continuous, surjective map. Then f has a Borel section, so in particular for any $B \subseteq Y$, $f^{-1}[B]$ is Borel if and only if B is. Moreover, if they are Borel, then the two are of the same Borel class.*

This corollary says that when X is invariant over a small set, we need not specify the parameter set in order to talk about the Borel class of X. It is a generalisation of a well-known fact for sets which are definable or type-definable with parameters.

Corollary 3.2. *Let A, B be any countable sets. Suppose X is A-invariant and B-invariant. Then the Borel class of X over A is the same as the Borel class of X over B (in particular, X is Borel over A if and only if it is Borel over B).*

Proof. Without loss of generality, we can assume that $A \subseteq B$. Then $f\colon S(B) \to S(A)$ is a continuous surjection, and $f^{-1}[X_A] = X_B$, so by Lemma 3.1 we get the result. □

Remark. We can show in the same way that if Δ, Λ are countable sets of formulas with parameters, and X is Δ-invariant (that is, invariant under automorphisms preserving all formulas in Δ) and Λ-invariant, then $X_\Delta \subseteq S(\Delta)$ and $X_\Lambda \subseteq S(\Lambda)$ have the same Borel class, where $S(\Delta)$ and $S(\Lambda)$ are the Stone spaces of the boolean algebras of formulas generated by Δ and Λ, respectively, while X_Δ, X_Λ are defined in the natural manner.

This definition is somewhat self-explanatory, but since we are going to use it quite often, it should be stated explicitly.

Definition. We say that an invariant equivalence relation E on X refines type if for any $a, b \in X$ whenever $a\, E\, b$, then $a \equiv b$ (i.e. $\mathrm{tp}(a/\emptyset) = \mathrm{tp}(b/\emptyset)$). Equivalently, E refines type if $E \subseteq {\equiv}\restriction_X$.

Similarly, we say that E refines Kim-Pillay strong type $\equiv_{KP}$ if $E \subseteq {\equiv_{KP}}\restriction_X$ and likewise we say that Kim-Pillay type refines E if ${\equiv_{KP}}\restriction_X \subseteq E$.

The next definition is very important; it will be used to interpret a bounded, invariant equivalence relation E as an abstract equivalence relation on a Polish space. It is a mild generalisation of E^M_L and E^M_{KP} as introduced in [KPS13].

Definition. Suppose E is a bounded, invariant equivalence relation on an invariant set X, while M is a model.

Then we define $E^M \subseteq S_X(M)^2 \subseteq S(M)^2$ as the relation

$$p \; E^M \; q \iff \text{there are some } a \models p \text{ and } b \models q \text{ such that } a \; E \; b$$

(And the next fact tells us that E-classes are M-invariant, so this is equivalent to saying that for all $a \models p$, $b \models q$ we have $a \; E \; b$, which implies that E^M is an equivalence relation.)

The next fact shows that E^M is well-behaved in the sense explained in parentheses, and the Borel classes of E^M and E are the same (which justifies the definition of Borel class of E at the beginning of this subsection).

Proposition 3.3 (generalisation of Remark 2.2(i) in [KPS13])**.** *Consider a model M, and some bounded, invariant equivalence relation E on an invariant subset X of a product of sorts P.*

Consider the natural restriction map $\pi\colon S_{P^2}(M) \to S_P(M)^2$ (i.e. $\pi(\mathrm{tp}(a,b/M)) = (\mathrm{tp}(a/M), \mathrm{tp}(b/M))$). Then we have the following facts:

- *Each E-class is M-invariant, in particular*

$$a \; E \; b \iff \mathrm{tp}(a,b/M) \in E_M \iff \mathrm{tp}(a/M) \; E^M \; \mathrm{tp}(b/M)$$

  and $\pi^{-1}[E^M] = E_M$.
- *If M is countable, the the Borel class of E^M is the same as that of E_M and the same as that of E (considered as a subset of $S_{P^2}(\emptyset)$).*
- *If M is countable and $Y \subseteq X$ is Borel over M, then the Borel class of the restriction $E^M\restriction_{Y_M}$ is the same as that of $E_M \cap (Y^2)_M$ (and therefore, by Lemma 3.1, independent of the choice of the countable model M over which Y is invariant).*

Proof. For the first bullet, notice that, by Fact 2.4, E is refined by Lascar strong type, which in turn is refined by equivalence over M (by Fact 2.2), and therefore any points equivalent over M are also Lascar equivalent, and therefore E-equivalent.

For the second bullet we use Lemma 3.1:

- $S_P(M)^2$, $S_{P^2}(M)$ are compact Polish spaces, so we apply the lemma to $f = \pi$ and $B = E^M$ (which we can do by the first bullet).
- Secondly, $S_{P^2}(M)$ and $S_{P^2}(\emptyset)$ are Polish, so we apply the lemma with $f = \pi_\emptyset \colon S_{P^2}(M) \to S_{P^2}(\emptyset)$ and $B = E$ (which we can do, since by definition $E_M = \pi_\emptyset^{-1}[E]$).

The last part follows analogously from Lemma 3.1, as $\pi^{-1}[E^M\restriction_{Y_M}] = E_M \cap (Y^2)_M$. □

The next two facts will be frequently used in conjunction with Corollary 3.2 to estimate the Borel class of various sets over a model M.

Corollary 3.4. *If E is a bounded, invariant equivalence relation on X and $Y \subseteq X$ is E-saturated (i.e. containing any E-class intersecting it), then for any model M, Y is M-invariant.*

Proof. Since Y is E-saturated, it is a union of E-classes, each of which is setwise M-invariant (and therefore so is any union of E-classes). □

Corollary 3.5. *If G is an invariant group and H is a subgroup of G containing some invariant subgroup of bounded index (equivalently, we may say that H contains G^{000}), H is invariant over any model M.*

Proof. Immediate from previous corollary with E being the relation of being in the same coset of G^{000}. □

The next proposition establishes a notion of Borel cardinality.

Proposition 3.6 (generalisation of Proposition 2.3 in [KPS13])**.** *For any E which is a bounded, (invariant) Borel equivalence relation on some X invariant in a product P of sorts, and if $Y \subseteq X$ is pseudo-Borel and E-saturated, then the Borel cardinality of restriction of E^M to Y_M does not depend on the choice of the countable model M. (In particular, if $X = Y$ is type-definable without parameters, the Borel cardinality of E^M does not depend on the choice of a countable model M.)*

Proof. Analogous to [KPS13](Proposition 2.3): it is enough to show that if $M \preceq N$ are countable models, then the Borel cardinalities of E^M and E^N coincide. To that end, consider the restriction map $\pi\colon S_P(N) \to S_P(M)$, and a Borel section $s\colon S_P(M) \to S_P(N)$ of π (which we have by Lemma 3.1).

Since Y is Borel over M (by Corollary 3.4 and Corollary 3.2), π and s restrict to Borel maps $\pi\colon S_Y(N) \to S_Y(M)$ and $s\colon S_Y(M) \to S_Y(N)$. On the other hand, by Proposition 3.3 (since E-classes are M-invariant), π is a reduction of $E^N{\restriction}_{Y_N}$ to $E^M{\restriction}_{Y_M}$, and because s is a section of π, it is a reduction of $E^M{\restriction}_{Y_M}$ to $E^N{\restriction}_{Y_N}$. □

We have thus justified the following definition.

Definition. If E is as in the previous proposition, then by Borel cardinality of E we mean the Borel cardinality of E^M for a countable model M. Likewise, we say that E is smooth if E^M is smooth for a countable model M.

In the same manner, if Y is pseudo-Borel and E-saturated, the Borel cardinality of $E{\restriction}_Y$ is the Borel cardinality of $E^M{\restriction}_{Y_M}$ for some countable model M.

Remarks.

- By similar methods, we could show that Borel cardinality (and Borel class) of an invariant equivalence relation is well-defined in an even stronger sense: if we have an extension of f.o. (countable) languages $L \subseteq L'$, $\mathfrak{C}$ is a monster model in both L and L', and E is a bounded, Borel equivalence relation on $\mathfrak{C}$ in the smaller language L, then its Borel cardinality is the same in both signatures.
- We could analogously define projective class for sets invariant over small parameter sets, and it would be similarly independent of the parameters (and would most likely allow an application theory of Borel cardinality for analytic and other projective equivalence relations to bounded and invariant equivalence relations in first order theories).

Fact 3.7. *A bounded, type-definable equivalence relation is smooth. Similarly, if a restriction of a bounded, invariant equivalence relation to a saturated, pseudo-G_δ set Y is relatively type-definable, then the restriction is smooth.*

Proof. If E is type-definable, then so is its domain, and the corresponding subset of $S(M)^2$ is closed (by Proposition 3.3), and in particular G_δ, so by the last point of Fact 2.1 E is smooth. The proof of the second part is analogous: the Borel cardinality of E to Y is the Borel cardinality of $E^M \cap (Y_M)^2$, which is closed in $(Y_M)^2$‘, and thus smooth. □

This last fact (or at least the first part of it) can also be proved in a slightly different way, using the logic topology, which we will do in the following subsection.

3.2. **Remarks on logic topology.** In this subsection, we will find an alternative (and equivalent) definition of Borel cardinality for some bounded equivalence relations.

Proposition 3.8. *Suppose E is a type-definable, bounded equivalence relation on a type-definable set X. Then the quotient map $X \to X/E$ factors through $S_X(M)$, yielding a map $\mathrm{tp}(a/M) \to [a]_E$ (for any model M), which is continuous when X/E is given logic topology.*

Proof. That the quotient map factors is an immediate consequence of the fact that E-classes are M-invariant by Proposition 3.3. Continuity follows from Corollary 3.4 and Corollary 3.2: a closed set $F \subseteq X/E$ corresponds to a pseudo-closed subset of X, which is closed over M because it is E-saturated. □

We will need a topological lemma.

Lemma 3.9. *Suppose X is a compact, zero-dimensional Hausdorff space, Y is a Hausdorff space, while $f\colon X \to Y$ is a continuous, surjective mapping. Then Y has a basis of closed sets consisting of the $f[B]$ for $B \subseteq X$ clopen.*

Proof. If B is clopen, it is compact, so $f[B]$ is closed too, as a compact subset of a Hausdorff space. It remains to show that any closed $D \subseteq Y$ is the intersection of a family of $f[B]$ for varying clopen B.

Choose arbitrary $y \in Y \setminus D$. We intend to find some clopen B such that $f[B]$ contains D but not y. Since Y is Hausdorff, we can find an open set $U \supseteq D$ such that $y \notin U$ (e.g. $Y \setminus \{y\}$).

Now, $f^{-1}[U]$ is an open set, so it is the union of basic clopen sets. But $f^{-1}[D]$ is compact, so it is covered by some finitely many of these, and – since a finite union of clopen sets is clopen – in fact there is a single clopen B with $f^{-1}[D] \subseteq B \subseteq f^{-1}[U]$.

But then – owing to the fact that f is onto – $D = f[f^{-1}[D]] \subseteq f[B] \subseteq f[f^{-1}[U]] = U$, and in particular $f[B]$ is closed and doesn't contain y. □

The above proposition and lemma gives us a somewhat concrete choice of a small basis for logic topology, and shows that it is actually Polish in cases that interest us (this is well-known: see [KN02], Fact 1.3).

Corollary 3.10. *Suppose E is a type-definable, bounded equivalence relation on a type-definable set X.*

Then for any model M, the logic topology on X/E has a basis of closed sets consisting of the quotients of M-definable sets. (In particular if X is contained in some countable product of sorts and M is a countable model, it implies that it is compact, Hausdorff and second-countable, and therefore Polish.)

Proof. Consider the map $\Psi\colon S_X(M) \to X/E$ as in Proposition 3.8. This is a continuous surjection, and X/E is Hausdorff by Fact 2.7, so we can apply the previous lemma and the result follows immediately. □

As a special case, we get the following statement about groups.

Corollary 3.11. *If G is a type-definable group and $H \leq G$ is a type-definable subgroup, then the Logic topology on G/H is compact Polish topology. If H is normal, then G/H is a compact Polish group.*

Proof. Since H is type-definable, the relation of lying in the same coset of H is type-definable, so we can apply the previous fact. It is easy to see that if operations on G are type-definable, then for normal H, the operations on G/H are continuous. □

Finally, the previous discussion allows us to describe the Borel cardinality of some bounded and invariant equivalence relations in a slightly different way.

Corollary 3.12. *If E is a bounded, Borel equivalence relation on a type-definable set X, while F is a type-definable, bounded equivalence relation on X refining E, then E induces naturally a Borel equivalence relation E' on X/F, given by*

$$[x]_F \; E' \; [y]_F \iff x \; E \; y$$

which is bireducible with E^M via the natural map $\Psi\colon S_X(M) \to X/F$.

(In particular, the Borel cardinality of E is the same as the Borel cardinality of E', and if we have $E = F$, E is smooth.)

Proof. X/E is a standard Borel space because it is Polish, by Corollary 3.10.

That E' is a well-defined equivalence relation follows immediately from the fact that F refines E.

By Proposition 3.8 we have $\Psi\colon S_X(M) \to X/F$, which induces a continuous surjection $\Psi^2\colon S_X(M)^2 \to (X/F)^2$, and because E-classes are M-invariant and F-saturated, we have $E^M = (\Psi^2)^{-1}[E']$, so by Lemma 3.1, E' is Borel. It also follows that Ψ is a reduction of E^M to E', and the Borel section of Ψ (obtained via Lemma 3.1) is a reduction in the other direction. □

Remarks.

- The preceding corollary gives us another way to represent as abstract Borel equivalence relations those E which are defined on a type-definable set X and are refined by $\equiv_{KP}\restriction_X$, namely as induced relation E' on $X/{\equiv_{KP}}$.

  This approach has the added benefit of being independent of any additional variables, like the choice of model M. On the other hand, sometimes we want to deal with E finer than $\equiv_{KP}$, and then it is not applicable.
- Similarly, we can show that if $Y \subseteq X$ is pseudo-closed and E-saturated, $E^M\restriction_Y$ is naturally bireducible with $E'\restriction_{Y/F}$.

3.3. **Orbital equivalence relations and normal forms.** In this subsection, we introduce some more specific kinds of invariant equivalence relations, which naturally arise when we interpret the main result.

Definition (Normal form)**.** If $\Phi_n(x, y)$ is a sequence of (partial) types such that $\Phi_0(x, y) = \{x = y\}$ and which is increasing (i.e. for all n, $\Phi_n(x, y) \vdash \Phi_{n+1}(x, y)$), then we say that $\bigvee_{n\in\mathbf{N}} \Phi_n(x, y)$ is a normal form for an invariant equivalence relation E on an an (invariant) set X if we have for any $a, b \in X$ an equivalence $a \; E \; b \iff \mathfrak{C} \models \bigvee_{n\in\mathbf{N}} \Phi_n(a, b)$, and if the binary function $d = d_\Phi\colon \mathfrak{C}^2 \to \mathbf{N} \cup \{\infty\}$ defined as

$$d(a, b) = \min\{n \in \mathbf{N} \mid \mathfrak{C} \models \Phi_n(a, b)\}$$

(where $\min\emptyset = \infty$) is an invariant metric with possibly infinite values, that is, it satisfies the axioms of coincidence, symmetry and triangle inequality. In this case we say that d induces E on X.

Remark. When talking about normal forms, we will sometimes implicitly assume that $\Phi_0(x, y) = \{x = y\}$ without stating it outright, as it will never be anything else.

Example 3.13. The prototypical example of a normal form is $\bigvee_n d_L(x, y) \leq n$, inducing $\equiv_L$, and d_L is the associated metric.

Remark. The Lascar distance, by its very definition, has the nice property that it is “geodesic” in the sense that if two points a, b are at distance n, then there is a sequence of points $a = a_0, a_1, \ldots, a_n = b$ such that each pair of successive points is at distance 1. The metrics obtained from normal forms usually will not have this property (notice that existence of such a “geodesic” metric for E is equivalent to E being the transitive closure of a type-definable relation).

Example 3.14. If $\Phi_n(x,y)$ is an increasing sequence of type-definable equivalence relations, then $\bigvee_n \Phi_n(x,y)$ is trivially a normal form. In particular, if $E = \Phi(x,y)$ is type-definable, then we can put – for all $n > 0$ – $\Phi_n(x,y) = \Phi(x,y)$, yielding a somewhat degenerate normal form for E.

Definition. If we have an invariant equivalence relation E on a set X with a normal form $\bigvee_{n\in\mathbf{N}} \Phi_n(x,y)$, corresponding to a metric d, and $Y \subseteq X$ is some set, then the diameter of Y is the supremum of d-distances between points in Y.

Fact 3.15. *If E is as above, and X is (the set of realisations of) a single complete type, then all E-classes have the same diameter (because the* $\operatorname{Aut}(\mathfrak{C})$ *acts transitively on X in this case, and the diameter is invariant under automorphisms).*

The following proposition is the essential step in adapting the techniques of [KMS13] to prove Theorem 4.6.

Proposition 3.16. *Suppose E is a relatively F_σ (over $\emptyset$), bounded equivalence relation on an invariant (over $\emptyset$) set X. Then E has a normal form $\bigvee_n \Phi_n$ such that $\Phi_1(x,y)$ holds for any x,y which are terms of an infinite indiscernible sequence. (This implies that for any a,b, if $d_L(a,b) \leq n$, then $\models \Phi_n(a,b)$, so that $d \leq d_L$. It also shows that every relatively F_σ equivalence relation has a normal form.)*

Proof. Because E is bounded, the Lascar strong type restricted to X is a refinement of E (by Fact 2.4), and hence $E \cup (\equiv_L\restriction_X) = E$. In addition, since E is relatively F_σ, we can find types $\Phi_n(x,y)$ such that $x \mathrel{E} y \iff \bigvee_n \Phi_n(x,y)$.

Consider the sequence $\Phi'_n(x,y)$ of types, defined recursively:

(1) $\Phi'_0(x,y) = \{x = y\}$
(2) $\Phi'_1(x,y) = \Phi_1(x,y) \vee \Phi_1(y,x) \vee x = y \vee d_L(x,y) \leq 1$
(3) $\Phi'_{n+1}(x,y) = \Phi_{n+1}(x,y) \vee \Phi_{n+1}(y,x) \vee (\exists z)(\Phi'_n(x,z) \wedge \Phi'_n(z,y))$

It is easy to see that $\bigvee \Phi'_n$ is a normal form and represents the smallest equivalence relation containing E and $\equiv_L$ (as a set of pairs), which is just E, and $d_L(x,y) \leq 1$ (i.e. the statement that x,y are in an infinite indiscernible sequence) implies $\Phi'_1(x,y)$ by the definition.

The statement in the parentheses follows from the fact that $d_L(a,b) \leq n$ is defined as the n-fold composition of $d_L(a,b) \leq 1$. □

The theorem of Newelski we will see shortly is a motivating example for the study of Borel cardinality: it can be interpreted as saying that some equivalence relations have Borel cardinality of at least $\Delta(2^{\mathbf{N}})$. We will see later in Corollary 4.8 that for E which are orbital (a concept which we will define soon), we can strengthen this result to replace $\Delta(2^{\mathbf{N}})$ with $\mathbf{E}_0$, and this is optimal in the sense explained in a remark after Corollary 4.8.

Theorem 3.17 ([New03](Corollary 1.12))**.** *Assume $x \mathrel{E} y$ is an equivalence relation refining $\equiv$, with normal form $\bigvee_{n\in\mathbf{N}} \Phi_n$. Assume $p \in S(\emptyset)$ and $Y \subseteq p(\mathfrak{C})$ is pseudo-closed and E-saturated. Then either E is equivalent on Y to some $\Phi_n(x,y)$ (and therefore E is relatively type-definable on Y), or $|Y/E| \geq 2^{\aleph_0}$.*

Remark. Newelski uses a slightly more stringent definition of a normal form (which we may enforce in all interesting cases without any significant loss of generality), i.e. that d satisfies not only triangle inequality, but also

$$d(a,b), d(b,c) \leq n \implies d(a,c) \leq n+1$$

but the definition used here is sufficient for the previous theorem, and in addition, it has the added benefit of being satisfied by the Lascar distance d_L, and the author feels that it is more natural in general.

The following corollary allows us some freedom with regards to the normal form, allowing us to replace – in some cases – any normal form with one chosen as in Proposition 3.16, without loss of generality.

Corollary 3.18. *Suppose E is an relatively F_σ equivalence relation finer than $\equiv$. Then for any class C of E, the following are equivalent:*

(1) *C is pseudo-closed,*
(2) *C has finite diameter with respect to each normal form of E (i.e. it has finite diameter with respect to the metric induced by each normal form, as introduced previously),*
(3) *C has finite diameter with respect to some normal form of E.*

Proof. Assume that C is pseudo-closed. Setting $Y = C$ in Theorem 3.17 we immediately get that C has finite diameter with respect to any normal form of E. Implication from the second condition to third follows from the fact that E has a normal form by the previous proposition, and the implication from third to first is trivial. □

Example 3.19. The above is no longer true if we allow E to be refined by $\equiv$. For example, consider the theory $T = \mathrm{Th}(\mathbf{R}, +, \cdot, 0, 1, <)$ of real closed fields, and the total relation on the entire model (field). Clearly, it has a normal form $\{x = y\} \vee \bigvee_{n>0}(x = x)$, and the induced metric is just the discrete 0-1 metric, and in particular its only class (the entire model) has diameter 1. On the other hand, we might give it a normal form $\{x = y\} \vee \bigvee_{n>0}(\bigwedge_{m\geq n}(x = m \leftrightarrow y = m))$ (where m ranges over natural numbers). With respect to this normal form, any two distinct positive natural numbers k, l are at distance $\max(k, l)+1$. In particular, the diameter of the only class is infinite.

From Corollary 3.18, we deduce the following description of relatively F_σ equivalence relations with pseudo-closed classes.

Corollary 3.20. *Suppose E is a relatively F_σ equivalence relation refining $\equiv$. Then the following are equivalent:*

(1) *each class of E is pseudo-closed,*
(2) *each class of E has finite diameter with respect to any normal form of E,*
(3) *each class of E has finite diameter with respect to some normal form of E.*

In addition, they imply that E is refined by $\equiv_{KP}$ (restricted to the domain of E). (And we will see later in Example 5.5 that the converse does not hold.)

Proof. The first part follows immediately from the previous corollary.

"In addition" can be obtained thus: E refines $\equiv$, so it is enough to show that the restriction of E to any $p \in S(\emptyset)$ is refined by the restriction of $\equiv_{KP}$ to p. But any class in the restriction has finite diameter with respect to some normal form, and they all have the same diameter (by Fact 3.15), so in fact the restriction is type-definable and as such refined by $\equiv_{KP}$ (by Fact 2.4). □

Remark. If E is a type-definable equivalence relation, then its classes are trivially pseudo-closed, so by the above, if E is refined by $\equiv$, then for any normal form of E, all E-classes have finite diameter.

For technical reasons, later on we will rely on the action of an automorphism group, so we introduce the following definition.

Definition (Orbital equivalence relation, orbital on types equivalence relation)**.** Suppose E is an invariant equivalence relation on a set X.

- We say that E is orbital if there is a group $\Gamma \leq \mathrm{Aut}(\mathfrak{C})$ such that Γ preserves classes of E setwise and acts transitively on each class.

- We say that E is orbital on types if it refines type and the restriction of E to any complete $\emptyset$-type is orbital.

Remarks.

- An orbital equivalence relation always refines type. (So every orbital equivalence relation is orbital on types.)
- The relations $\equiv_L, \equiv_{KP}$ are orbital (as witnessed by $\operatorname{Aut} f_L(\mathfrak{C}), \operatorname{Aut} f_{KP}(\mathfrak{C})$).
- The group witnessing that a given relation is orbital can always be chosen as a normal subgroup of $\operatorname{Aut}(\mathfrak{C})$.

The following proposition shows that the definition of a orbital on types equivalence relation is, in a way, the weakest possible for the proof of Theorem 4.6.

Proposition 3.21. *An invariant equivalence relation E is orbital on types if and only if for any class C of E there is a group Γ such that Γ preserves E classes within the (complete $\emptyset$-)type p containing C and Γ acts transitively on C.*

Proof. First, notice that both conditions imply that E is refined by $\equiv$, so we can assume that.

The implication ($\Rightarrow$) is just a weakening; for ($\Leftarrow$), observe that $\operatorname{Aut}(\mathfrak{C})$ acts transitively on $X := p(\mathfrak{C})$, so for any class $C' \in X/E$ we have an automorphism σ which takes C to C'. It is easy to see that then $\sigma\Gamma\sigma^{-1}$ acts transitively on C' and preserves all E-classes in X setwise. From that we conclude that the normal closure of Γ in $\operatorname{Aut}(\mathfrak{C})$ witnesses that E restricted to X is orbital. □

The following simple corollary allows us to easily recognise some relations as orbital on types.

Corollary 3.22. *If E is an invariant equivalence relation on an invariant set X, refining $\equiv$, and the restriction of E to any complete type in X has at most two classes, then E is orbital on types.*

Proof. Without loss of generality we may assume that X is a single complete type, so $\operatorname{Aut}(\mathfrak{C})$ acts transitively on X. In particular, for any element $a \in X$, we have a set $S \subseteq \operatorname{Aut}(\mathfrak{C})$ such that $S \cdot a = [a]_E$. Since E is invariant, elements of S preserve $[a]_E$ and so does the group $\Gamma = \langle S \rangle$.

Of course Γ preserves X, so it also preserves the complement $X \setminus [a]_E$. But since E has at most two classes, this means that Γ preserves all classes, so by the previous proposition E is orbital on types. □

At a glance, it is not obvious whether the condition that E is orbital on types is any stronger than the condition that it refines type. The following examples show that it is indeed the case.

Example 3.23. Consider the permutation group

$$\begin{aligned} G &= \langle (1,2)(3,5)(4,6), (1,3,6)(2,4,5) \rangle \\ &= \{(), (1,2)(3,5)(4,6), (1,3,6)(2,4,5), \\ &\quad (1,4)(2,3)(5,6), (1,5)(2,6)(3,4), (1,6,3)(2,5,4)\} \end{aligned}$$

acting naturally on a 6-element set. Then the equivalence relation $\sim$ such that $1 \sim 2$, $3 \sim 4$, $5 \sim 6$ (and no other nontrivial relations) is preserved by G, but it is not the orbital equivalence relation of any subgroup (in fact, the subgroup of G preserving all $\sim$-classes setwise is trivial).

Let M_0 be a structure with base set $\{1, 2, 3, 4, 5, 6\}$, with a relation symbol E for $\sim$, and such that G is the automorphism group of M_0 (which we can obtain, for instance, by adding a predicate for the set of all orbits of G on M_0^6 to the language).

Then E is an invariant (even definable) equivalence relation which refines $\equiv$ and is not orbital on types.

We can extend Example 3.23 to an infinite model in a number of simple ways, for instance as follows.

Example 3.24. Let M_0 be as in the previous example, and let M_1 be a structure which is just a large set (larger than the desired saturation degree) with no non-logical symbols, and consider the structure $\mathfrak{C} = (M_0 \times M_1, M_i, \pi_i)_{i \in \{0,1\}}$ (with all the structure of M_0), where π_i is the projection $M_0 \times M_1 \to M_i$.

Then $\mathfrak{C}$ is saturated (it is clearly categorical in every cardinality), and the automorphism group is just the product of G from the previous example and the full permutation group of M_1. In particular, there is only one 1-type on the product sort of $\mathfrak{C}$, and the relation on it induced by E is not orbital, although it trivially refines type.

We finish with a little less artificial example.

Example 3.25. Consider a large algebraically closed field K of characteristic $p > 0$, and choose some $t \in K$, transcendental over the finite field $\mathbf{F}_p$, and consider $T = \mathrm{Th}(K, +, \cdot, t)$.

Let $n > 3$ be a natural number which is not divisible p, and X be the set of n-th roots of t in K (i.e., the roots of $x^n - t$). Notice that X generates a definable, finite additive group $\langle X \rangle$. Put

$$G = (\{a = (a_1, a_2) \in K^2 \mid a_1 + a_2 \in \langle X \rangle\}, +)$$

G is a definable group (definably isomorphic to $K \times \langle X \rangle$). Consider the equivalence relation on G defined by

$$a \mathrel{E} b \iff (a \equiv b \wedge a_1 + a_2 = b_1 + b_2)$$

We will show that E is not orbital on types, even though it is type-definable, bounded and refines $\equiv$. (Nb. this E is the conjunction of $\equiv$ and the relation of lying in the same coset of G^{000}, which in this case is equal to the classical model-theoretic connected component G^0.)

Let ξ be some primitive nth root of unity. Notice that for any $x_1, x_2 \in X$, the pairs (x_1, ξ) and (x_2, ξ^{-1}) have the same type:

- Consider the field extension $\mathbf{F}_p^{\mathrm{alg}} \subseteq \mathbf{F}_p^{\mathrm{alg}}(x_1)$.
- There is an automorphism of $\mathbf{F}_p^{\mathrm{alg}}$ which takes ξ to ξ^{-1}, and since the extension is purely transcendental, it extends to an automorphism of $\mathbf{F}_p^{\mathrm{alg}}(x_1)$ fixing x_1.
- x_1 and x_2 are transcendental over $\mathbf{F}_p^{\mathrm{alg}}$ and they generate $\mathbf{F}_p^{\mathrm{alg}}(x_1)$ (because their quotient is some root of unity), so they are conjugate by some automorphism of $\mathbf{F}_p^{\mathrm{alg}}(x_1)$ over $\mathbf{F}_p^{\mathrm{alg}}$.
- The composition of the two automorphisms is an automorphism of $\mathbf{F}_p^{\mathrm{alg}}(x_1)$ which takes (x_1, ξ) to (x_2, ξ^{-1}), and therefore fixes $x_1^n = x_2^n = t$, so it extends to an automorphism of K which fixes t.

From that it follows that all $a \in G$ of the form $(x, \xi^{\pm 1} x)$, where $x \in X$ have the same type, say $p_0 \in S_G(\emptyset)$.

But then in particular for any $x \in X$ we have $(x, \xi x) \mathrel{E} (\xi x, x)$. If E was orbital on types, there would be some automorphism $f \in \mathrm{Aut}(K/t)$ which takes x to ξx and ξx to x – therefore taking ξ to ξ^{-1} – which preserves setwise the E-classes within p_0. But then

$$b = f((\xi x, \xi^2 x)) = (x, \xi^{-1} x) \not\mathrel{E} (\xi x, \xi^2 x) = a \models p_0$$

because $a_1 + a_2 - b_1 - b_2 = x(\xi + \xi^2 - 1 - \xi^{-1}) = \xi^{-1}x(\xi^3 + \xi^2 - \xi^1 - 1)$ and ξ is algebraic of degree $n > 3$.

We have seen that the E-class of $(\xi x, \xi^2 x)$ is not preserved by any f which takes $(x, \xi x)$ to $(\xi x, x)$, and because $(\xi x, \xi^2 x) \models p_0$, it follows that E is not orbital on types.

3.4. **Invariant subgroups as invariant equivalence relations.** We have seen before that strong connected components of definable groups are, in a way, a special case of strong types. In this section, we will show that the correspondence is more general, and invariant subgroups are a special case of invariant equivalence relations.

Definition. Suppose G is a type-definable group and $H \leq G$ is invariant. We define E_H as the relation on G of lying in the same right coset of H:

$$\begin{aligned} g_1 \; E_H \; g_2 &\iff Hg_1 = Hg_2 \\ &\iff (\exists h_1, h_2 \in H)\, h_1 g_1 = h_2 g_2 \\ &\iff g_1 g_2^{-1} \in H \end{aligned}$$

Remark. Clearly, E_H is invariant, and it has $[G : H]$ classes, so H has bounded index if and only if E_H is a bounded equivalence relation.

It is not hard to see that invariant subgroups of type-definable groups correspond to invariant equivalence relations as shown in the following proposition.

Lemma 3.26. *Suppose G is a type-definable group and $H \leq G$ is Borel. Then E_H is Borel and its Borel class is the same as that of H.*

Proof. Consider the mapping $f\colon S_{G^2}(\emptyset) \to S_G(\emptyset)$ given by $\mathrm{tp}(a, b/\emptyset) \mapsto \mathrm{tp}(ab^{-1}/\emptyset)$. Since the operations in G are type-definable, this map is a well-defined and continuous surjection.

It is easy to see that $E_H = f^{-1}[H]$, and since $S_{G^2}(\emptyset), S_G(\emptyset)$ are compact and Polish, by Lemma 3.1, E_H has the same Borel class as H. □

Remarks.

- The previous proposition would remain true if we had taken for E_H the relation of lying in the same left coset, but right cosets will be technically more convenient in a short while.
- Equivalence relations E_H do not refine type, and in particular are not orbital on types, which is a desirable property. We will resolve this issue shortly by choosing a different equivalence relation to represent H, which will be Borel bireducible with E_H and orbital on types for normal H.

Until the end of this subsection, we fix a definable group G and recall from Theorem 2.6 the structure $(\mathfrak{C}, \mathfrak{X}, \cdot)$ where $\mathfrak{X}$ is a sort for a (left) principal homogeneous space for a group G definable in $\mathfrak{C}$, and $\cdot$ is the symbol for the left action.

Definition. Let H be an invariant subgroup of G. Then $E_{H,X}$ is the relation on $\mathfrak{X}$ of being in the same H-orbit.

Proposition 3.27. *The mapping $\Phi\colon H \mapsto E_{H,X}$ is a bijection between invariant subgroups of G and invariant equivalence relations on $\mathfrak{X}$.*

Proof. We fix some $x_0 \in \mathfrak{X}$, so as to apply the description of the automorphism group of $(\mathfrak{C}, \mathfrak{X}, \cdot)$ from Theorem 2.6.

First, choose some invariant $H \leq G$. We will show that $E_{H,X}$ is invariant.

Choose arbitrary $h \in H$, an original automorphism σ of $\mathfrak{C}$ and some $g, k \in G$, and denote by σ_g the automorphism of $(\mathfrak{C}, \mathfrak{X})$ induced by g. Since every pair of two $E_{H,X}$-related elements is of the form (kx_0, hkx_0), and the automorphism group of $(\mathfrak{C}, \mathfrak{X}, \cdot)$

is generated by automorphisms induced by some σ and automorphisms induced by some g (as in Theorem 2.6), it is enough to show that $\sigma(kx_0)\ E_{H,X}\ \sigma(hkx_0)$ and $kgx_0\ E_{H,X}\ hkgx_0$. The latter is immediate by the definition of $E_{H,X}$. For the former, just see that

$$\sigma(kx_0) = \sigma(k)x_0\ E_{H,X}\ \sigma(h)\sigma(k)x_0 = \sigma(hkx_0)$$

because $\sigma(h) \in H$ (by invariance of H).

To see that Φ is a bijection, choose an arbitrary invariant equivalence relation E on $\mathfrak{X}$, and let H be the setwise stabiliser of $[x_0]_E$. Choose arbitrary $h \in H$, $\sigma \in \mathrm{Aut}(\mathfrak{C})$. Then:

$$x_0\ E\ hx_0 \implies x_0 = \sigma(x_0)\ E\ \sigma(hx_0) = \sigma(h)x_0$$

therefore $\sigma(h) \in H$ and since h and σ were arbitrary, H is invariant. To see that $E = E_{H,X}$, notice that for any $x_1 = k_1x_0$ and $x_2 = k_2x_0$ we have

$$k_1x_0\ E\ k_2x_0 \iff x_0\ E\ k_2k_1^{-1}x_0 \iff k_2k_1^{-1} \in H \iff (\exists h \in H)\, hk_1x_0 = k_2x_0$$

(The first equivalence is obtained by applying the automorphism σ_{k_1} induced by k_1.)

□

Remark. An invariant subgroup $H \leq G$ has bounded index if and only if $E_{H,X}$ is a bounded equivalence relation.

Proposition 3.28. *Suppose $H \leq G$ is invariant and of bounded index, $H \leq K \leq G$ and K is pseudo-Borel.*

Then $E_{H,X}$ is Borel if and only if H is, and it has the same Borel class as H.

Moreover, if H is Borel, $E_H\restriction_K$ and $E_{H,X}\restriction_{K\cdot x_0}$ have the same Borel cardinality (in particular, E_H and $E_{H,X}$ have the same Borel cardinality).

Proof. Suppose H is Borel and bounded and fix an arbitrary K according to the assumptions (e.g. $K = G$) and a countable model $N \preceq (\mathfrak{C}, \mathfrak{X}, \cdot)$ containing x_0 – so that $N = (M, G(M) \cdot x_0)$.

Then K is Borel over M (by Corollary 3.5) and the map $f\colon S_G(N) \to S_{\mathfrak{X}}(N)$ defined by $f(\mathrm{tp}(g/N)) = \mathrm{tp}(g \cdot x_0/N)$ is a homeomorphism (because it is induced by an N-definable bijection), and f takes K_N to $(K \cdot x_0)_N$ and E_H^N to $E_{H,X}^N$.

In a similar fashion, the restriction map $g\colon S_G(N) \to S_G(M)$ (with the latter considered in the original structure $\mathfrak{C}$) is also a homeomorphism – it is clearly continuous and surjective, and it follows from Theorem 2.6 that it is injective (automorphisms of $(\mathfrak{C}, \mathfrak{X}, \cdot)$ fixing $(M, G(M) \cdot x_0)$ pointwise are the same as those which fix M and x_0 pointwise, and they have the same orbits in $\mathfrak{C}$ as the automorphisms of $\mathfrak{C}$ fixing M pointwise). It also takes E_H^N to E_H^M and K_N to K_M.

These two facts, along with Lemma 3.26 complete the proof:

(1) The Borel class of H is the same as that of E_H (by Lemma 3.26).
(2) By the (composition of) aforementioned homeomorphisms, E_H^M has the same Borel class as $E_{H,X}^N$, and the two have the same Borel classes as E_H and $E_{H,X}$, respectively (by Proposition 3.3).
(3) Similarly, $E_H^M\restriction_{K_M}$ is taken by a homeomorphism to $E_{H,X}^N\restriction_{(K\cdot x_0)_N}$, so they have the same Borel cardinalities, which are by definition the Borel cardinalities of $E_H\restriction_K$ and $E_{H,X}\restriction_{K\cdot x_0}$, respectively.

□

We finish with an observation that allows us to easily see that some $E_{H,X}$ are orbital.

Proposition 3.29. *Suppose H is a normal, invariant subgroup of G. Then $E_{H,X}$ is orbital.*

Proof. Consider the action $*$ of H on $(\mathfrak{C},\mathfrak{X},\cdot)$ by automorphisms (using Theorem 2.6, viewing H as a subgroup of $G\rtimes\operatorname{Aut}(\mathfrak{C})$). Then – because H is a normal subgroup of G – we have for any $x=g\cdot x_0\in\mathfrak{X}$ that

$$H*(g\cdot x_0)=(gH^{-1})\cdot x_0=(gH)\cdot x_0=(Hg)\cdot x_0=H\cdot(g\cdot x_0)=[x]_{E_{H,X}}$$

So $H\leq G\rtimes\operatorname{Aut}(\mathfrak{C})=\operatorname{Aut}((\mathfrak{C},\mathfrak{X},\cdot))$ witnesses that $E_{H,X}$ is orbital. $\square$

Remark. The converse of the previous proposition is not true: if we have $G=S_3$, $H=\langle(1,2)\rangle$ and $\operatorname{Aut}(\mathfrak{C})$ acting on G in such a way that any $\sigma\in\operatorname{Aut}(\mathfrak{C})$ acts on G either trivially or by conjugation by $(1,2)$, then although H is not normal, $E_{H,X}$ is orbital: for $\sigma\in\operatorname{Aut}(\mathfrak{C})$ acting nontrivially on G we have

$$((1,2),\sigma)(g\cdot x_0)=((1,2)\cdot g\cdot(1,2)^{-1})\cdot(1,2)^{-1}\cdot x_0=(1,2)\cdot(g\cdot x_0)$$

4. The technical theorem

4.1. **Descriptive-set-theoretic observations.** The following theorem is the core of the descriptive-set-theoretic part of the argument for Theorem 4.6. The next corollary will imply the latter, as soon as we show that its assumptions are satisfied.

Theorem 4.1 ([KMS13](Theorem 2.2))**.** *Suppose that P is a Polish space, R_n is a sequence of F_σ subsets of P^2, Γ is a group of homeomorphisms of P and $\mathcal{O}\subseteq P$ is an orbit of Γ such that for each n and open $U\subseteq P$ intersecting $\mathcal{O}$, there are distinct $x,y\in\mathcal{O}\cap U$ with $\mathcal{O}\cap(R_n)_x\cap(R_n)_y=\emptyset$. Then there is a continuous, injective homomorphism*

$$(2^{\mathbf{N}},E_0,\neg E_0)\to(\overline{\mathcal{O}},E^P_\Gamma,\neg\bigcup_n R_n)$$

(where E^P_Γ is the orbit equivalence relation of Γ on P).

As before, when E is an invariant, bounded equivalence relation, we denote by E^M the induced equivalence relation on $S(M)$. For the statement of the next corollary, we need to extend the notion of distance to the type spaces.

Definition. If E is a relatively F_σ equivalence relation induced by a metric d (coming from some normal form), then we also denote by d_M the induced distance on $S(M)$, i.e.

$$d_M(p_1,p_2)=\min_{a_1\models p_1,a_2\models p_2}d(a_1,a_2)$$

Remark. The classes of E^M are precisely the "metric components" of d_M, i.e. the maximal sets of types which are pairwise in finite distance of one another in the sense of d_M, though d_M might not satisfy the triangle inequality, so it is not in general a metric.

We link Theorem 4.1 to our context (bounded and invariant equivalence relations) by means of the following corollary.

Corollary 4.2 (based on [KMS13](Corollary 2.3))**.** *Suppose we have:*

- *a countable theory T with monster model $\mathfrak{C}$,*
- *a countable model $M\preceq\mathfrak{C}$,*
- *an invariant subset X of a countable product of sorts of $\mathfrak{C}$,*
- *a bounded, relatively F_σ equivalence relation E on X, with normal form $\bigvee_n\Phi_n$, inducing metric d,*
- *a pseudo-G_δ and E-saturated $Y\subseteq X$.*

Assume in addition that there is some $p\in Y_M\subseteq S_X(M)$ such that for every formula $\varphi\in p$ with parameters in M, and for all $N\in\mathbf{N}$, there is some $\sigma\in\operatorname{Aut}(\mathfrak{C})$ such that:

(1) *σ fixes M and all E-classes in Y setwise (and therefore Y itself as well),*

(2) *$\varphi \in \sigma(p)$ and $N < d_M(\sigma(p), p)$.*

Then there is a continuous, injective homomorphism

$$(2^{\mathbf{N}}, E_0, \neg E_0) \to (Y_M, E^M{\restriction}_{Y_M}, \neg(E^M{\restriction}_{Y_M}))$$

In particular, $E^M{\restriction}_{Y_M}$ is not smooth.

Proof. Let $\Gamma < \mathrm{Aut}(\mathfrak{C})$ be the group of all automorphisms fixing M and all classes of E in Y setwise, acting naturally on Y_M by homeomorphisms. For each n, put $R_n = \{(p, q) \in (Y_M)^2 \mid d_M(p, q) \leq n\}$, and $\mathcal{O} = \Gamma \cdot p$.

Denote by P the sort (or a product of sorts) containing X (the domain of E).

R_n are F_σ because they are closed – they are intersections of $(Y_M)^2$ with $R'_n = \{(p, q) \in S_P(M)^2 \mid d_M(p, q) \leq n\}$, which are compact, as continuous images of compact $[\Phi_n(x, y)] \subseteq S_{P^2}(M)$).

Choose an arbitrary natural n and a basic open set $[\psi]$ intersecting $\mathcal{O}$. Then for some $\gamma \in \Gamma$ we have $\psi \in \gamma(p)$. Let $\varphi = \gamma^{-1}(\psi)$. Then for some $\sigma \in \Gamma$ we have $\varphi \in \sigma(p)$ and $2n < d_M(\sigma(p), p)$. Let $x = \gamma(p)$ and $y = \gamma\sigma(p)$. So $d_M(x, y) = d_M(p, \sigma(p)) > 2n$ and so x, y are distinct elements of $[\psi] \cap \mathcal{O}$ and $\mathcal{O} \cap (R_n)_x \cap (R_n)_y = \emptyset$.

Now we can apply the previous theorem with $P = Y_M$ (which is Polish because it is G_δ in $S(M)$, by Corollary 3.2 and Corollary 3.4) and $\Gamma, \mathcal{O}, R_n$ as before, so as to obtain a continuous, injective homomorphism $(2^{\mathbf{N}}, E_0, \neg E_0) \to (\overline{\mathcal{O}}, E_\Gamma^{Y_M}, \neg \bigcup_n R_n)$ (with $\overline{\mathcal{O}}$ being the closure in Y_M), while inclusion gives a continuous, injective homomorphism from $(\overline{\mathcal{O}}, E_\Gamma^{Y_M}, \neg \bigcup_n R_n)$ to $(Y_M, E^M{\restriction}_{Y_M}, \neg(E^M{\restriction}_{Y_M}))$, since $\neg \bigcup_n R_n$ is just $\neg(E^M{\restriction}_{Y_M})$, and $E_\Gamma^{Y_M} \subseteq E^M$ (because Γ preserves E-classes intersecting Y), so by composing the two maps we complete the proof. □

4.2. **Generic formulas and proper types.** Throughout this subsection, $\Gamma \leq \mathrm{Aut}(\mathfrak{C})$ is an arbitrary group of automorphisms, while C is an arbitrary orbit of Γ (in some product of sorts of $\mathfrak{C}$). The facts will be cited from [KMS13], where Γ is $\mathrm{Aut\,f}_L(\mathfrak{C})$ (and so C is a Lascar strong type), but the proofs are exactly the same in the general case.

Definition.

- if φ, ψ are formulas, we say that $\varphi \vdash_C \psi$ if for all $c \in C$, if $\models \varphi(c)$, then $\models \psi(c)$ (i.e. $\varphi(\mathfrak{C}) \cap C \subseteq \psi(\mathfrak{C})$); likewise for one or both of φ, ψ replaced with types;
- we say that a formula (with parameters) φ is generic if there is a finite sequence $\tau_0, \ldots, \tau_{n-1}$ of elements of Γ such that $\bigvee_{i<n} \tau_i(\varphi)$ covers C; we say that a type is generic if every conjunction of its formulas is generic;
- for any (possibly incomplete, usually with parameters) type p, we say that p is proper if there is a non-generic formula ψ such that if φ is a conjunction of formulas in p, then the formula $\varphi \vee \psi$ is generic, or equivalently, $p \vee \psi$ is generic.

Remark. Considering C is Γ-invariant, whenever $\varphi \vdash_C \psi$ and φ is generic, so is ψ.

This fact allows us to extend a proper type to a proper complete type.

Fact 4.3 ([KMS13](Lemma 4.9))**.** *Let p be any proper type. Then for any formula φ (with parameters), one of $p \cup \{\varphi\}, p \cup \{\neg\varphi\}$ is proper.*

This fact will, in turn, help with extending a proper type to make it satisfy the assumptions of Corollary 4.2.

Fact 4.4 ([KMS13](Proposition 4.10))**.** *Suppose p is a proper type and $\varphi \in p$. Then there are $\tau_0, \ldots, \tau_{m-1} \in \Gamma$ such that for any $\sigma \in \Gamma$ there is $i < m$ such that $p \cup \{\sigma(\tau_i(\varphi))\}$ is proper.*

This lemma will be used for the base case for the recursion in Theorem 4.6.

Lemma 4.5 (Based on [KMS13](Proposition 4.11)). *Let $a \in C$. Then the type $q = (d_L(x,a) \leq 1)$ (where d_L is the Lascar distance; i.e. q is the type of an element of an infinite indiscernible sequence containing a – not to be confused with the type of any particular such element, as q is usually incomplete) is generic (and hence proper).*

4.3. **Proof of the technical theorem.** As mentioned in the introduction, a similar theorem has been proved, independently, in [KM13] using different methods. The proof we list here is a generalization of the main result of [KMS13], where the relation in question is the Lascar strong type.

Remark. The statement of Theorem 4.6 is somewhat technical, attempting to be almost as general as the proof allows. For variants which may be easier to digest, see Corollary 4.7 and Corollary 4.8.

Theorem 4.6 (based on [KMS13](Theorem 4.12)). *Suppose we have:*

- *an invariant set X (*in a countable product of sorts of $\mathfrak{C}$*),*
- *a bounded, relatively F_σ equivalence relation E on X, refining $\equiv$,*
- *a pseudo-G_δ and E-saturated set $Y \subseteq X$,*
- *an E-class $C \subseteq Y$ with infinite diameter with respect to some normal form of E,*
- *a group $\Gamma \leq \mathrm{Aut}(\mathfrak{C})$ preserving all E-classes setwise and acting transitively on C.*

Then $E{\restriction_Y}$ is not smooth.

Proof. By Proposition 3.16, we can choose a normal form $\bigvee_n \Phi_n$ for E such that $d_L(x,y) \leq 1 \vdash \Phi_1(x,y)$, where d_L is the Lascar distance. Denote the invariant metric associated to this normal form by d, so that for any $x, y \in X$ we have $d(x,y) \leq n \iff \models \Phi_n(x,y)$ (and $d \leq d_L$). Recall that d induces the distance on types $d_M(p_1,p_2) = \min(d(a_1,a_2))$ with min ranging over $a_1 \models p_1, a_2 \models p_2$.

By Corollary 3.18, C has infinite diameter with respect to d. We will apply the lemmas of subsection 4.2 with C, Γ as in the statement of the theorem.

Ultimately, the goal is to apply Corollary 4.2. To that end, we will construct a pair (M,p) such that M is a countable model of T and $p \in Y_M$ satisfies the hypotheses of Corollary 4.2.

Choose arbitrary $a \in C$ and let $q(x) = \Phi_1(x,a)$. We construct recursively three sequences σ_i, p_i, M_i with $i \in \mathbf{N}$ such that

1. each M_i is a finite set containing a, each p_i is a finite type over M_i, and each σ_i is in Γ,
2. M_i and p_i are both monotone with respect to $\subseteq$,
3. $p_i \cup q$ is proper for each i (in particular, it is consistent),
4. for every $i, N \in \mathbf{N}$ and $\varphi \in p_i$, there are some j, j' such that $d(\sigma_j(a), a) > N$ and $\sigma_j^{-1}(\varphi) \in p_{j'}$,
5. for each i and $\varphi \in L(M_i)$ consistent with the sorts of X, there is some j such that p_j contains one of $\varphi, \neg\varphi$,
6. for each M_i, if some formula $\varphi \in L_1(M_i)$ has a realisation, it has one in some M_j (Tarski-Vaught condition),
7. for every i, i', there is some $j \in \mathbf{N}$ such that $M_j \supseteq \sigma_{i'}[M_i] \cup \sigma_{i'}^{-1}[M_i]$.

Claim. If we take $M = \bigcup_i M_i$ and $p = \bigcup_i p_i$ as above, then they satisfy the assumptions Corollary 4.2, as witnessed by the σ_i.

Proof.

- $M = \bigcup_i M_i$ will be an elementary submodel of $\mathfrak{C}$ by (2), (6) along with Tarski-Vaught criterion, countable by (1);
- By (2) and (1), $p = \bigcup_i p_i$ is a type, by (5), $p \in S(M)$ (is a complete M-type),
- $a \in M$ and by (2) and (3), $p \cup q$ is consistent, so (considering p is complete) in fact p extends q, and in particular $p \in C_M \subseteq Y_M$,
- for any $\varphi \in p$ and any $N \in \mathbf{N}$, the automorphism σ_j witnessing (4) for φ and $(N+6)$ fixes M setwise (by (7)) as well as E-classes (because it comes from Γ), and Y (because it is an union of E-classes); it also takes p to a point $\sigma_j(p) \ni \varphi$ (directly by (4)). Furthermore, if we take any $c \models p$, then

  $$d_M(p, \sigma_j(p)) \geq d(c, \sigma_j(c)) - 4$$

  (By triangle inequality, the fact that two elements realizing the same complete M-type are at Lascar distance at most two (Fact 2.2), and the inequality $d \leq d_L$.) On the other hand (by the previous bullet), $c \models q$, therefore $d(c, a) \leq 1$, and

  $$d(c, \sigma_j(c)) \geq d(a, \sigma_j(a)) - d(\sigma_j(a), \sigma_j(c)) - d(a, c) \geq d(a, \sigma_j(a)) - 2$$

  so $d_M(p, \sigma_j(p)) > (N+6) - 2 - 4 = N$ and the claim is proved.

$\square$(claim)

It remains to show

Claim. The recursive construction can be accomplished as announced.

Proof. For the base case, we take $M_0 = \{a\}$, $p_0 = \emptyset$, $\sigma_0 = \mathrm{id}$. The relevant assumptions are satisfied: the only non-trivial one is (3), but it follows from Lemma 4.5 because $(d_L(x, a) \leq 1) \vdash q(x)$ and we assume that Γ acts transitively on C. Now, suppose we have p_i, M_i, σ_i up to $i = n$. We need to do the following:

- for each natural number N and formula $\varphi \in p_n$, we need to find $\sigma \in \Gamma$ as in (4), such that p_n can be extended by $\sigma^{-1}(\varphi)$ while remaining proper. For specific N, φ we can do it by taking the $\tau_0, \ldots, \tau_{m-1}$ as in Fact 4.4 for $p_n \cup q$ and φ, and a $\tau \in \Gamma$ such that $d(\tau(a), a) > N + \max_i d(\tau_i(a), a)$ (which exists because $C = \Gamma \cdot a$ has infinite diameter). Then by triangle inequality and invariance of d, for each i we have

  $$\begin{aligned} &d((\tau \circ \tau_i)^{-1}(a), a) = d(a, \tau(\tau_i(a)) \geq \\ &\geq d(a, \tau(a)) - d(\tau(a), \tau(\tau_i(a))) = d(a, \tau(a)) - d(a, \tau_i(a)) > N. \end{aligned}$$

  Then we choose i such that $p_{n+1} := p_n \cup q \cup \{\tau(\tau_i(\varphi))\}$ is proper and put $\sigma_{n+1} := (\tau \circ \tau_i)^{-1}$. For all N, φ we use a book-keeping argument (scheduling for all pairs (N, φ) in advance);
- for each formula $\varphi \in L_\alpha(M_n)$, we need to extend p_n to include one of $\varphi, \neg\varphi$ while keeping it proper; for a single formula, we can do it by Fact 4.3, whereas for all formulas, it can be done by a simple book-keeping argument (in a manner consistent with the previous point);
- extend M_n to satisfy the Tarski-Vaught condition (6); we cannot do it in one step for all φ, but it can also be easily done using the standard book-keeping argument;
- extend M_n to include its images and preimages by $\sigma_0, \ldots, \sigma_n$ – we can do it in a single step.

$\square$(claim)

$\square$

Remarks.

- We can always take for Γ the group of all automorphisms preserving E-classes setwise. (In which case $\Gamma \trianglelefteq \mathrm{Aut}(\mathfrak{C})$.)
- The methods of [KMS13] – with adjustments akin to those made here – could likely also be employed to prove a weak analogue of Theorem 4.6 in case of uncountable language and/or relations E on X in uncountable products of sorts, but this is beyond the scope of this paper.

4.4. **Variants of Theorem 4.6.** In this subsection, we consider some alternate forms of Theorem 4.6.

We obtain the following immediate corollary (a partial contrapositive) of Theorem 4.6. We will use it in the next section as one of the main tools for characterisation of smooth equivalence relations (Theorem 5.8).

Corollary 4.7. *Suppose E is a bounded, relatively F_σ and orbital on types equivalence relation on an invariant set X. Then if $Y \subseteq X$ is E-saturated, pseudo G_δ and $E{\restriction}_Y$ is smooth, then all E-classes contained in Y have finite diameter.*

Proof. The proof is by contraposition. Choose any E-class $C \subseteq Y$ of infinite diameter and some $a \in C$.

Then $C \subseteq [a]_\equiv$, so also $C \subseteq [a]_\equiv \cap Y$. The restriction of E to $[a]_\equiv$ is orbital, which is witnessed by some $\Gamma \leq \mathrm{Aut}(\mathfrak{C})$. Since C has infinite diameter, we may apply Theorem 4.6 to the relation $E{\restriction}_{[a]_\equiv}$, the set $Y \cap [a]_\equiv$ and the group Γ, deducing that $E{\restriction}_{Y\cap[a]_\equiv}$ is not smooth. But that trivially implies that $E{\restriction}_Y$ is not smooth (by inclusion mapping). □

We may simplify a little more if Y is contained in a single complete type. The next corollary can be seen as a strengthening of Theorem 3.17 in case of E which are orbital on types (because a relation with countably many classes is smooth).

Corollary 4.8. *Suppose E is a bounded, F_σ and orbital on types equivalence relation on an invariant set X. Let $a \in X$ be arbitrary, and assume that $Y \subseteq [a]_\equiv$ is E-saturated, pseudo-G_δ with $a \in Y$. Fix any normal form $\bigvee_n \Phi_n$ for E. Then the following are equivalent:*

1. *$E{\restriction}_Y$ is smooth,*
2. *$E{\restriction}_{[a]_\equiv}$ is type-definable,*
3. *all E-classes in $[a]_\equiv$ have finite diameter with respect to $\bigvee_n \Phi_n$,*
4. *all E-classes in $[a]_\equiv$ are pseudo-closed,*
5. *$[a]_E$ has finite diameter with respect to $\bigvee_n \Phi_n$,*
6. *$[a]_E$ is pseudo-closed.*

Proof. We may assume without loss of generality that $X = [a]_\equiv$. Then E is orbital.

If $E{\restriction}_Y$ is smooth, we can apply Corollary 4.7 to deduce that all E-subclasses of Y have finite diameter. But then by Fact 3.15, all E-classes in $X = [a]_\equiv$ have the same, finite diameter $N \in \mathbf{N}$, so E is type-definable by Φ_N, and by Corollary 3.18, all other conditions immediately follow.

On the other hand, if $[a]_E$ is pseudo-closed, then again by Corollary 3.18, it means that it has finite diameter $N \in \mathbf{N}$. But then by Fact 3.15, all classes have diameter N, so E is type-definable by Φ_N, and by Fact 3.7, $E{\restriction}_Y$ is smooth.

It is clear that all conditions imply that $[a]_E$ is pseudo-closed, and since this implies that $E{\restriction}_Y$ is smooth – which in turn implies all the other conditions – they are all equivalent. □

Remark. Note that Corollary 4.8 is, in a way, a strongest possible result. This is to say, there are examples of bounded, F_σ and orbital equivalence relations whose Borel cardinality is exactly that of $\mathbf{E}_0$ (cf. [KPS13], Example 3.3), so we cannot replace the condition that $E{\restriction_Y}$ is smooth with some weaker upper bound on Borel cardinality.

For relations refining $\equiv_{KP}$, we may be even more specific.

Corollary 4.9. *Suppose E is bounded, relatively F_σ and orbital on types. Suppose in addition that it refines $\equiv_{KP}$. Then for any a in the domain of E, we have that $E{\restriction_{[a]_{\equiv_{KP}}}}$ is trivial (i.e. total on $[a]_{\equiv_{KP}}$) if and only if it is smooth. (In particular, if E is defined on a Borel set X and it is smooth, then it is equal to $\equiv_{KP}{\restriction}X$.)*

Proof. The implication from left to right is trivial. To prove the converse, choose any a in domain of E.

The set $[a]_{\equiv_{KP}}$ is E-saturated (because E refines $\equiv_{KP}$), type-definable over a and contained in $[a]_{\equiv}$, so we can assume without loss of generality that E is defined on $[a]_{\equiv}$. Then we can apply Corollary 4.8, which tells us that if $E{\restriction_{[a]_{\equiv_{KP}}}}$ is smooth, then E is type-definable. But in this case E is refined by $\equiv_{KP}$ (by Fact 2.4), and therefore equal to $\equiv_{KP}$ restricted to $[a]_{\equiv}$ and so $E{\restriction_{[a]_{\equiv_{KP}}}}$ is trivial. $\square$

Remark. In Theorem 4.6 and the above corollaries, we do not really need to assume that E is relatively F_σ, only that it is F_σ on types (i.e. that its restrictions to complete types are F_σ, or even just types intersecting Y). This is because the proofs all of those only really work with the restrictions of E to a specific complete type intersecting Y. (And in fact whenever we show non-smoothness in those, we show that the restriction to some complete type (or its subset) is non-smooth.)

We infer an analogous result for invariant subgroups of bounded index of definable groups, which we will employ later in the context of definable group extensions, in the final section (specifically in Theorem 6.2).

Corollary 4.10. *Suppose that G is a definable group and $H \trianglelefteq G$ is an invariant, normal subgroup of bounded index, which is F_σ (equivalently, generated by a countable family of type-definable sets). Suppose in addition that $K \geq H$ is a pseudo-G_δ subgroup of G. Then $E_H{\restriction_K}$ is smooth if and only if H is type-definable.*

Proof. If H is type-definable, then by Lemma 3.26 E_H is a type-definable equivalence relation (on a type-definable set), and as such it is immediately smooth by Fact 3.7, and so is its restriction to K.

The proof in the other direction will proceed by contraposition: assume that H is not type-definable. Recall Proposition 3.27: consider, once again, the sorted structure $(\mathfrak{C}, \mathfrak{X}, \cdot)$.

By Proposition 3.28, H corresponds to a bounded F_σ equivalence relation $E_{H,X}$ on $\mathfrak{X}$ (which is not type-definable, since H is not), which is only defined on a single type, and – owing to the assumption that H is normal and Proposition 3.29 – orbital.

Evidently $K \cdot x_0$ is $E_{H,X}$-saturated and pseudo-G_δ, so we can apply Corollary 4.8 to $E = E_{H,X}$ and $Y = K \cdot x_0$, deducing that $E_{H,X}{\restriction_{K\cdot x_0}}$ is not smooth, and therefore (by Proposition 3.28) neither is $E_H{\restriction_K}$. $\square$

5. Characterisation of smooth equivalence relations and Borel cardinalities

In this section, we will attempt to characterise the bounded, orbital on types and (relatively) F_σ equivalence relations which are smooth, and in particular compare smoothness and type-definability. Firstly, we analyse several examples showing us some of the limitations of this attempt.

5.1. **Counterexamples.**

Proposition 5.1. *Suppose E is a type-definable equivalence relation on a type-definable set X, and that there are countably many complete $\emptyset$-types on X, and infinitely many of them are not covered by singleton E-classes. Then E has a normal form such that the classes of E have unbounded diameter (that is, there is no uniform bound on the diameter).*

Proof. Let p_n with $n > 0$ be some enumeration of complete $\emptyset$-types on X. Then put (for $n > 0$)

$$\Phi_n(x,y) = (x = y) \vee \left(E \wedge \bigvee_{m_1,m_2 \leq n} p_{m_1}(x) \wedge p_{m_2}(x) \right)$$

It is easy to see that for each n, $\Phi_n(x,y)$ is a type-definable equivalence relation and Φ_n is increasing, so $\bigvee_n \Phi_n(x,y)$ is trivially a normal form. In addition, any non-singleton E-class intersecting p_n has diameter at least $n+1$.

There are infinitely many p_n which intersect an E class which is not a singleton, so in particular, the non-singleton classes have no (finite) uniform bound on diameter. □

Example 5.2. Let $T = ACF_0$ be the theory of algebraically closed fields of characteristic 0. Consider $E = \equiv_{KP}$ as a relation on $\mathfrak{C}^2$. The space $S_2(\mathbf{Q}^{\mathrm{alg}})$ is countable, because T is ω-stable, and since $\mathbf{Q}^{\mathrm{alg}}$ is a model, we know (due to Proposition 3.3) that $\equiv_{KP}$ has only countably many classes (on the set of pairs). It is also, of course, smooth, orbital and even type-definable.

Despite being rather well-behaved, E still has a normal form with respect to which the classes have arbitrarily large diameter, by the preceding proposition:

(1) (The set of realisations of) each type of the form $\mathrm{tp}(q,t/\emptyset)$ with $q \in \mathbf{Q}$ and t transcendental is a single, infinite $\equiv_{KP}$-class (because it is the set of realisations of a single type over $\mathbf{Q}^{\mathrm{alg}}$), and in particular, it is not covered by singleton classes.
(2) Furthermore, $S_2(\emptyset)$ is countable (because T is ω-stable), so by Proposition 5.1, E has a normal form with respect to which its classes have arbitrarily large diameter.

Proposition 5.3. *Suppose there is a non-isolated complete $\emptyset$-type p_0 such that $p_0(\mathfrak{C})$ is not contained in a single class of some definable, bounded (equivalently, with finitely many classes) equivalence relation E. Then the relation*

$$E'(x,y) = (E(x,y) \vee \neg p_0(x)) \wedge (x \equiv y)$$

is F_σ and smooth, but not type-definable.

Furthermore, if $E \cap \equiv$ is orbital on types, then so is E'.

Proof. A definable and bounded equivalence relation has only finitely many classes, so E' differs from $\equiv$ only in that one class is divided into finitely many pieces. Fix a countable model M and a Borel reduction $f\colon S(M) \to X$ of $\equiv^M$ as an equivalence relation on $S(M)$ to $\Delta(X)$, equality on a Polish space X (which exists because $\equiv$ is smooth, being type-definable).

Let $[p_0]_\equiv / E = \{A_1, \ldots, A_n\}$. Then define $\tilde{f}\colon S(M) \to X \sqcup \{1,\ldots,n\}$ (where $\sqcup$ is the disjoint union and $\{1,\ldots,n\}$ has discrete topology) by

$$\tilde{f}(\mathrm{tp}(x/M)) = \begin{cases} f(\mathrm{tp}(x/M)) & x \not\models p_0 \\ j & x \in A_j \end{cases}$$

Then clearly $\tilde{f}$ is Borel and witnesses that E' is smooth.

E' is easily seen to be F_σ, as it is the intersection of the open (and therefore F_σ, as the language is countable) set $(E(x,y) \vee \neg p_0(x))$ and the closed set $(x \equiv y)$.

It remains to show that E' is not type-definable. For that, we need the following

Claim. For any formula (without parameters) $\psi \in p_0$, there is some $x \models p_0$ and $x' \not\models p_0$, such that $x' \models \psi$ and $E(x,x')$. In fact, we can find such x' for any $x \models p_0$.

Proof. The proof is by contraposition: we assume that there are no such x, x' for ψ, and we will show that p_0 is isolated. Let

$$E''(x,y) = (E(x,y) \wedge \psi(x) \wedge \psi(y)) \vee (\neg\psi(x) \wedge \neg\psi(y))$$

Then E'' is a definable equivalence relation which has finitely many classes (at most 1 more than E) and (by the assumption), $p_0(\mathfrak{C})$ is a union of E''-classes, of which there are only finitely many, so $p_0(\mathfrak{C})$ is definable with some parameters. But since it is invariant, it implies that it is definable without parameters, and therefore p_0 is isolated.

Once we have some x' for an $x \models p_0$, we may obtain one for each of them simply by applying automorphisms. □(claim)

Now we choose a sequence φ_n of formulas such that $\bigwedge_n \varphi_n \vdash p_0$ and $\varphi_{n+1} \vdash \varphi_n$. Let $x_0, y_0 \models p_0$ be such that $\neg E(x_0, y_0)$ (which we can find because p_0 is not contained in a single E-class), and let x_n be a sequence of elements satisfying φ_n but not p_0, and simultaneously satisfying $E(x_n, x_0)$ (this sequence exists by the claim), and let y_n be a sequence such that each (x_0, x_n) is conjugate to (y_0, y_n) (so that $x_n \equiv y_n$ and $y_n \models \varphi_n$ and $E(y_0, y_n)$).

Then any limit point of the sequence $\mathrm{tp}(x_n, y_n/\emptyset)$ in $S_2(\emptyset)$ is not in E', even though each $\mathrm{tp}(x_n, y_n/\emptyset)$ is in E', so E' is not type-definable.

The "furthermore" part is obvious, since E' agrees with $E \cap \equiv$ on p_0 and is total when restricted to any other type. □

Example 5.4. Consider $T = \mathrm{Th}(\mathbf{Z}, +)$ (the theory of additive group of integers) and the type $p_0 = \mathrm{tp}(1/\emptyset)$ (the type of an element not divisible by any natural number).

The type p_0 is not isolated, and it is not contained in a single class of the definable relation E of equivalence modulo 3, while $E \cap \equiv$ has at most two classes in each complete type, so it is orbital on types due to Corollary 3.22.

In particular – by the preceding proposition – the relation $E'(x,y)$ which says that $x \equiv y$ and they either have the same residue modulo 3 or else they are both divisible by some natural number (i.e. they are not of the same type as 1), is F_σ, orbital on types and smooth, but not type-definable.

Example 5.5. Suppose there is some $a \in \mathfrak{C}$ such that $\equiv_{KP}$ has two classes on $[a]_\equiv$ (like $a = \sqrt{2}$ for $\mathfrak{C} \models ACF_0$), so that $\equiv_{KP}\restriction_{[a]_\equiv} \sim_B \Delta(2)$. Consider the infinite disjoint union of copies of $\mathfrak{C}$, i.e. the multi-sorted structure $(\mathfrak{C}_n)_{n\in\mathbf{N}}$ where each $\mathfrak{C}_n$ is a distinct sort isomorphic to $\mathfrak{C}$ (without any relations between elements of $\mathfrak{C}_n$ and $\mathfrak{C}_m$ for $n \neq m$). Then consider $\overline{a} = (a_n)_{n\in\mathbf{N}}$ where a_n is the element of $\mathfrak{C}_n$ corresponding to a. Then $[\overline{a}]_\equiv = \prod_n [a_n]_\equiv$ and similarly

$$(b_n)_n \equiv_{KP} (c_n)_n \iff \bigwedge_n b_n \equiv_{KP} c_n$$

(by [Cas+01], Lemma 3.7(iii)). Now consider the following relation E on $[\overline{a}]_\equiv$:

$$(b_n)_n \; E \; (c_n)_n \iff \{n \mid b_n \not\equiv_{KP} c_n\} \text{ is finite}$$

Then E is refined by $\equiv_{KP}$, but on the other hand

$$E \sim_B (\equiv_{KP}\restriction_{[a]_\equiv})^{\mathbf{N}}/\mathrm{Fin} \sim_B \Delta(2)^{\mathbf{N}}/\mathrm{Fin} = \mathbf{E}_0$$

(This can be seen either directly or by considering E as an equivalence relation on $[\overline{a}]_{\equiv}/{\equiv_{KP}}$, as in Corollary 3.12.)

In particular, E is not smooth, it is easy to see that E is F_σ (because $\equiv_{KP}$ is type-definable and there are countably many finite subsets of $\mathbf{N}$), and it is also orbital, as its classes are just the orbits of the group

$$\left\{ (\sigma_n)_n \in \prod_{n\in\mathbf{N}} \mathrm{Aut}(\mathfrak{C}_n) \;\middle|\; \text{for all but finitely many } n,\, \sigma_n \in \mathrm{Aut\,f}_{KP}(\mathfrak{C}_n) \right\}$$

Additionally, E is only defined on a single type and is not type-definable, so by Corollary 4.8, all its classes have infinite diameter.

Example 5.6. Consider a saturated model K of the theory $T = \mathrm{Th}(\mathbf{R}, +, \cdot, 0, 1, <)$ of real closed fields. For each $n \in \mathbf{N}^+$ we have a $\emptyset$-type-definable equivalence relation $\Phi_n(x, y) = \bigwedge_{k\geq n}(x < k \leftrightarrow y < k)$. Consider the relation $E = \bigvee_n \Phi_n$ (with $\Phi_0(x, y) = \{x = y\}$, as before):

- E is an F_σ equivalence relation (and since Φ_n is an increasing sequence of equivalence relations, it is easy to see that $\bigvee_n \Phi_n$ is its normal form).
- E has two classes: the class C_{fin} of elements bounded from above by some natural number, and its complement C_∞. Therefore, it is bounded and smooth.
- C_{fin} is a class which is not pseudo-closed (otherwise, by compactness, it would intersect $\bigwedge_n x > n = C_\infty$).

This combination of features is possible because E does not refine $\equiv$ (and therefore it is not orbital on types), so we cannot apply Corollary 4.7 to it.

Example 5.7 (Example 3.39 in [KM13])**.** Let T be the theory of an infinite dimensional vector space over $\mathbf{F}_2$ in the language $(+, 0, U_n)_{n\in\mathbf{N}}$ (i.e. an infinite abelian group of exponent 2), where U_n are predicates for independent subspaces of codimension 1 (i.e. subgroups of index 2).

Consider $G = \mathfrak{C} \models T$ as a definable (additive) group, and let $H \leq G$ be the intersection of all U_n. Then $[G : H] = \mathfrak{c}$, and cosets of H are exactly the types $X_\eta = \bigcap_n U_n^{\eta_n}$, where $\eta\colon \mathbf{N} \to \{0, 1\}$, while $U_n^0 = U_n$ and $U_n^1 = \mathfrak{C} \setminus U_n$.

Consider the subspaces $W_\theta \leq G$ defined as $W_\theta = \pi^{-1}[\ker(\theta)]$, where $\pi\colon G \to G/H$ is the quotient map, and θ is a nonzero functional $G/H \to \mathbf{F}_2$. Each θ is uniquely determined by W_θ (since its value is 0 on $\pi[W_\theta]$ and 1 elsewhere and $\ker\pi$ is contained in all W_θ), so there are $|(G/H)^*| = \mathfrak{c} > \aleph_0$ distinct W_θ, in particular some $W = W_\theta$ is not definable.

On the other hand, W is invariant, as it is the union of some X_η which are type-definable, and $[G : W] = 2$ (because W has codimension 1), so W is not type-definable (if it was, its complement would also be type-definable, as it is invariant and a coset of W).

Now, recall Proposition 3.27. Let us extend $\mathfrak{C}$ to $(\mathfrak{C}, \mathfrak{X}, \cdot)$, where $\mathfrak{X}$ is a principal homogeneous space for G. Then W induces an invariant equivalence relation $E_{W,\mathfrak{X}}$ on $\mathfrak{X}$ which has two classes, is orbital by Proposition 3.29 (or Corollary 3.22) and not type-definable by Proposition 3.28.

Remark. The equivalence relation in the previous example is not type-definable, and it is unlikely to even be Borel, as the subspace W is the kernel of an almost arbitrary linear functional, which can be very “wild”. It does show, however, that we need some “definability” hypotheses beyond invariance for the likes of Theorem 3.17.

5.2. **Main characterisation theorem.**

Theorem 5.8. *Let E be an F_σ, bounded, orbital on types equivalence relation on a type-definable set X, and d be the invariant metric induced by a normal form of E. Then consider the four conditions:*

(1) *E classes have uniformly bounded diameter with respect to d.*
(2) *E is type-definable.*
(3) *E is smooth.*
(4) *E classes have finite diameter with respect to d (but possibly unbounded).*

These conditions are related as follows:

- (1) *implies* (2),(3),(4),
- (2) *implies* (3) *and* (4)*, but not* (1)
- (3) *implies* (4)*, but not* (2) *or* (1).
- (4) *does not imply* (2) *or* (1).

If we assume, in addition, that E refines $\equiv_{KP}$ (on X), then conditions (2),(3),(4) *are equivalent (and equivalent to simply $E = \equiv_{KP}\restriction_X$) and are implied by, but do not imply* (1).

If we assume instead that E is only defined on a single complete $\emptyset$-type, then all conditions are equivalent.

Proof. For the first part:

- (1) trivially implies (2)
- That (2) implies (3) follows from Fact 3.7.
- That (3) implies (4) follows immediately from Corollary 4.7 with $Y = X$.
- That (2) does not imply (1) follows from Example 5.2.
- That (3) does not imply (2) is demonstrated by Example 5.4.
- Other listed implications (or lack thereof) are logical consequences of the ones above.

To show that the last three conditions are equivalent if E refines KP-type, it is enough to show that (4) implies (2). But it follows easily from Corollary 3.20. That it does not imply (1) can be seen in Example 5.2.

To show that all four are equivalent if E is defined on a single type, it is enough to notice that (4) implies (1). But this is immediate from the fact that all classes have the same diameter (by Fact 3.15).

□

Remarks.

- The most significant in the previous corollary are the conditions (2) and (3), as they are inherent to E, whereas the two others depend on the choice of the metric inducing E. In particular, for bounded E which are F_σ, orbital on types and either defined on a single complete type, or refining $\equiv_{KP}$, we have that E is type-definable if and only if it is smooth.
- The property that E has only countably many classes implies (2),(3),(4), but not (1) (and is not implied by any of the conditions). (2) follows immediately from Theorem 3.17, while the others follow as a consequence of Theorem 5.8. That having countably many classes does not imply (1), we have seen in Example 5.2. That none of the conditions imply that there is only a countable number of classes can be seen by examining $\equiv$ in a non-small theory.
- Example 5.5 along with Corollary 3.20 show that the condition that E is refined by $\equiv_{KP}$ is strictly weaker than all the conditions in Theorem 5.8, even with the added assumption that E is only defined on a single type. (Of course, it is not *strictly* weaker if we assume that E refines $\equiv_{KP}$, although trivially so.)

5.3. **Possible extensions of Theorem 5.8 and related questions.** The picture of logical relations between the four conditions in Theorem 5.8 is almost complete, except for a single implication, raising the question.

Question 1. *In Theorem 5.8, is (3) equivalent to (4)?*

A counterexample, if it exists, would be an equivalence relation that is smooth when restricted to any single complete type is smooth, but is not on the entirety of its domain.

If, in the first part of Theorem 5.8, we drop the assumption that E is orbital on types (so we allow E to not refine type), then (2) does not imply (4) – as witnessed by Example 3.19 (though (1) certainly still implies the other conditions and (2) implies (3)).

We can, however, replace instead the assumption that E is orbital on types by the assumption that it refines type. In this case, it is unknown whether (2) (or (3)) implies (4), but otherwise the implications hold (even if we assume that E refines $\equiv_{KP}$ or is defined on a single type).

Question 2. *In Theorem 5.8, can we replace "orbital on types" by "refines type"? (Recall that by Example 3.23 and Example 3.25, this is not a trivial replacement.)*

This is closely related to a more specific question.

Question 3. *In Corollary 4.8, can we replace "orbital on types" by "refines type"? (Note that in this case Theorem 3.17 holds, as it has no orbitalness assumptions, so a counterexample would have Borel cardinality of exactly $\Delta(2^{\mathbf{N}})$)*

It is conceivable that this question (or the previous one) could be answered in the positive by showing that, given an invariant equivalence relation, we can find some different first order language such that the relation in question remains invariant, but becomes orbital on types (with some care taken so as to ensure that all the relevant assumptions and conclusions are preserved when we change the language).

If we drop the requirement that E is F_σ, points (1) and (4) do not have a clear interpretation (though we could imagine that a relation E could be induced by some invariant metric for which $d(x, y) \leq n$ is not type-definable, but satisfies some weaker conditions). However, the other two do make sense, so another question that arises naturally is the following.

Question 4. *Suppose E is a Borel equivalence relation on a type-definable set. Under what conditions does smoothness of E imply that it is type-definable?*

Notice that while we do know that an F_σ equivalence relation defined on a single type and with countably many classes is automatically type-definable (by Theorem 3.17). Which brings us to another, more concrete (and perhaps easy) question.

Question 5. *Suppose E is a bounded, invariant equivalence relation on a single complete type with two classes (and therefore orbital by Corollary 3.22). If E is Borel, then is it necessarily type-definable?*

Note that this is no longer true without at least some weak "definability" assumptions, as illustrated by Example 5.7. On the other hand, it is conceivable that some weaker assumptions than Borelness would suffice.

5.4. **Possible Borel cardinalities.** In this subsection, we will explore the problem of possible Borel cardinalities. We will not prove much, as the specifics are beyond the scope of this paper, but we will state some related questions and conjectures and see how they relate to the main results we have seen before.

In [KPS13], the authors suggest the following conjecture.

Conjecture 6 (Conjecture 2 in [KPS13])**.** *Any non-smooth K_σ equivalence relation can be represented [in the sense of Borel cardinality] by some $\equiv_L\restriction_{[a]_{\equiv_{KP}}}$ (in some theory T).*

Because $\equiv_L$ is F_σ (and therefore for a countable model M, its realisation $\equiv_L^M$ is K_σ when restricted to a compact set such as $([a]_{\equiv_{KP}})_M$), we know that $\equiv_L\restriction_{[a]_{\equiv_{KP}}}$ has the Borel cardinality of a K_σ equivalence relation, and it follows from Corollary 4.9 that if the relation is non-trivial, it is non-smooth (because $\equiv_L$ is orbital and refines $\equiv_{KP}$), so this is the strongest result of this kind we can hope for.

In this context, Theorem 3.17, Corollary 4.8 and Theorem 5.8 can be interpreted as providing a lower bound on the Borel cardinality in some cases. The above conjecture is a part of a bigger question.

Question 7. *Suppose E is a Borel, bounded equivalence relation on a type-definable set X, while $Y \subseteq X$ is E-saturated and pseudo-closed. Consider the following conditions:*

(1) *E is F_σ,*
(2) *E is orbital on types,*
(3) *E refines $\equiv$,*
(4) *E refines $\equiv_{KP}$,*
(5) *etcetera.*

For a given conjunction of the above conditions, what are the possible Borel cardinalities of $E, E\restriction_Y$? What if Y is a single $\equiv_{KP}$-class?

The author has found some other partial results in this vein, but they are beyond the scope of this paper.

6. Applications to definable group extensions

6.1. **Introduction to extensions by abelian groups.** This section will show an application of the main result (more precisely, of Corollary 4.10) to definable extensions by abelian groups. More specifically, we deal with short exact sequences of groups of the form

$$(\dagger) \qquad 0 \to A \to \widetilde{G} \to G \to 0$$

where A is an abelian group. In this case, there is a full algebraic description of $\widetilde{G}$ in terms of an action of G on A by automorphisms (induced by conjugation in $\widetilde{G}$) and a (2-)cocycle $h\colon G^2 \to A$ (to be defined shortly). We define multiplication on $A \times G$ by the formula

$$(\dagger\dagger) \qquad (a_1, g_1) \cdot (a_2, g_2) = (a_1 + g_1 \cdot a_2 + h(g_1, g_2), g_1 g_2)$$

And a cocycle is defined thus.

Definition. Let G be a group acting on an abelian group A. A function $h\colon G^2 \to A$ is a 2-cocycle if it satisfies, for all $g, g_1, g_2, g_3 \in G$, the following equations:

$$h(g_1, g_2) + h(g_1 g_2, g_3) = h(g_1, g_2 g_3) + g_1 \cdot h(g_2, g_3)$$
$$h(g, e) = h(e, g) = e$$

The equation ($\dagger\dagger$) endows $A \times G$ with group structure – with inverse $(a, g)^{-1} = (-g^{-1} \cdot a - h(g^{-1}, g), g^{-1})$ – which is compatible with the exact sequence ($\dagger$), and any $\widetilde{G}$ in such a short exact sequence has this form. In this language, the properties of the action and of h reflect the properties of the extension, e.g. central extensions correspond to trivial actions of G on A. More information about this subject (in abstract algebraic terms) can be found in e.g. [Rot02] (section 10.3., and in particular the part up to and including Theorem 10.14).

Definition. A definable extension of a definable group G by a definable abelian group A is a tuple $(G, A, *, h)$ where $*$ is a definable action of G on A by automorphisms and $h\colon G^2 \to H$ is a definable cocycle. We will also call that the group $\widetilde{G} = A \times G$ with multiplication defined as in (††).

Remark. The group $\widetilde{G}$ defined as above is definable.

In [GK13], the authors have shown that such extensions can, under some additional assumptions, give new examples of definable groups with $G^{00} \neq G^{000}$, building upon and extending the intuitions from the original example of model-theoretic universal cover of $SL_2(\mathbf{R})$, published in [CP12]. They also pose some questions and conjectures, one of which will be proved at the end of this section. To state their main result, we need the following definition.

Definition. A 2-cocycle $h\colon G^2 \to A$ is split via $f\colon G \to A$ if for all $g_1, g_2 \in G$ we have

$$\begin{aligned} h(g_1, g_2) = df(g_1, g_2) &:= f(g_1) + g_1 \cdot f(g_2) - f(g_1 g_2) \\ f(e) &= 0 \end{aligned}$$

Below is Theorem 2.2 of [GK13], one of its main results, which (along with the conjectures we will recall later this section) draws attention to definable group extensions.

Theorem 6.1. *Let G be a group acting by automorphisms on an abelian group A, where G, A and the action of G on A are $\emptyset$-definable in a structure $\mathcal{G}$, and let $h\colon G \times G \to A$ be a 2-cocycle which is B-definable in $\mathcal{G}$ and with finite range* $\mathrm{rng}(h)$ *contained in* $\mathrm{dcl}(B)$ *(the definable closure of B) for some finite parameter set $B \subset \mathcal{G}$. By A_0 we denote the subgroup of A generated by* $\mathrm{rng}(h)$. *Additionally, let A_1 be a bounded index subgroup of A which is type-definable over B and which is invariant under the action of G. Finally, suppose that:*

(i) *the induced 2-cocycle $\overline{h}_{|G_B^{00} \times G_B^{00}}\colon G_B^{00} \times G_B^{00} \to A_0/(A_1 \cap A_0)$ is non-split via B-invariant functions (i.e. there is no B-invariant function $f\colon G_B^{00} \to A_0/(A_1 \cap A_0)$ such that $\overline{h}\restriction_{G_B^{00}}$ is split via f),*

(ii) *$A_0/(A_1 \cap A_0)$ is torsion free (and so isomorphic with $\mathbf{Z}^n$ for some natural n).*

Then $\widetilde{G}_B^{000} \neq \widetilde{G}_B^{00}$ (where $\widetilde{G}$ is $G \times A$ with the group structure defined according to (††)).

6.2. **Main theorem for definable group extensions.** In this subsection, we will fix some (arbitrary) definable extension of a definable group G by an abelian definable group A, the group $\widetilde{G} = A \times G$ with cocycle h and the corresponding short exact sequence

$$0 \to A \to \widetilde{G} \xrightarrow{\pi} G \to 0$$

We denote the quotient map $\widetilde{G} \to G$ by π.

We intend to prove this theorem, and to that end, we need Corollary 4.10 and two lemmas.

Theorem 6.2. *Suppose $\widetilde{H} \trianglelefteq \widetilde{G}$ is invariant and of bounded index, generated by a countable family of type-definable sets, contained in a type-definable group $\overline{H} \leq \widetilde{G}$. Put $A_H := \overline{H} \cap A$ and $H = \pi[\overline{H}]$ (both of these are naturally type-definable groups).*

Assume that $H = \pi[\widetilde{H}]$ and that it acts trivially on $A_H/\widetilde{H} \cap A$, and $\widetilde{H} \cap A$ is type-definable. Then $\widetilde{H}$ is type definable.

In the following, until the end of the proof of Theorem 6.2, we will assume that we have everything as in its assumptions. We also fix a countable model M.

Lemma 6.3 (generalization of a part of the proof of [GK13](Proposition 2.14))**.** *Given an assignment $g \mapsto a_g$ such that for $g \in H$ we have $(a_g, g) \in \widetilde{H}$, the formula $\Phi((a,g)\cdot \widetilde{H}) = a - a_g + (\widetilde{H} \cap A)$ yields a well-defined, injective function $\overline{H}/\widetilde{H} \to A_H/(\widetilde{H}\cap A)$.*

Proof. Consider any $(a_1, g_1), (a_2, g_2) \in \overline{H}$.

Since $(a_{g_1}, g_1), (a_{g_2}, g_2) \in \widetilde{H}$, we have

$$
\begin{aligned}
(1)\quad (a_{g_1} - g_1 g_2^{-1}\cdot a_{g_2} - g_1\cdot h(g_2^{-1}, g_2) + h(g_1, g_2^{-1}), g_1 g_2^{-1}) = \\
= (a_{g_1}, g_1)(-g_2^{-1} a_{g_2} - h(g_2^{-1}, g_2), g_2^{-1}) = (a_{g_1}, g_1)(a_{g_2}, g_2)^{-1} \in \widetilde{H}
\end{aligned}
$$

we also have, by simple calculation – for arbitrary $g \in G$ and $a, a' \in A$ –

$$
(2)\qquad (a,g)^{-1}(a',g) = (0,g)^{-1}(-a,e)(a',e)(0,g) = (g^{-1}\cdot(a'-a), e)
$$

We want to show that $(a_1, g_1)(a_2, g_2)^{-1} \in \widetilde{H}$ if and only if $(a_1 - a_{g_1}) - (a_2 - a_{g_2}) \in \widetilde{H} \cap A$. The first condition says that

$$
(a_1 - g_1 g_2^{-1}\cdot a_2 - g_1 \cdot h(g_2^{-1}, g_2) + h(g_1, g_2^{-1}), g_1 g_2^{-1}) = (a_1, g_1)(a_2, g_2)^{-1} \in \widetilde{H}
$$

Multiplying it on the left by the inverse of the LHS of (1), and applying (2), we infer that it is equivalent to

$$
(g_1 g_2^{-1})^{-1}\left((a_1 - a_{g_1}) - g_1 g_2^{-1}\cdot(a_2 - a_{g_2})\right) \in \widetilde{H} \cap A
$$

On the other hand, because the action of H on the cosets of $\widetilde{H}\cap A$ is trivial, and $g_1, g_2 \in H$, so we can cancel both $g_1 g_2^{-1}$, and in conclusion the first condition is equivalent to

$$
(a_1 - a_{g_1}) - (a_2 - a_{g_2}) \in \widetilde{H} \cap A
$$

which is just the second condition. □

For any topological space X, by $F(X)$ we denote the Effros Borel space of closed subsets of X, with the σ-algebra generated by sets of the form $B_U := \{F \in F(X) \mid F \cap U \neq \emptyset\}$ for open (or equivalently just basic open, if X is second-countable) $U \subseteq X$. We have the following theorem ([Kec95], Theorem 12.13).

Theorem 6.4 (Kuratowski–Ryll-Nardzewski selection theorem)**.** *For any Polish space X, there is a Borel function $d\colon F(X) \to X$, such that for any nonempty closed $F \subseteq X$ we have $d(F) \in F$. (In fact, there is a sequence of such functions, picking out a dense subset of each nonempty closed set.)*

We apply this to prove the second lemma.

Lemma 6.5 (generalization of [KPS13](Proposition 4.4))**.** *Let $A_1 \subseteq \widetilde{H}$ be a type-definable subgroup of A_H of bounded index. Let F be the relation on the compact, Polish (by Corollary 3.11) group A_H/A_1 of lying in the same coset of $(\widetilde{H}\cap A)/A_1$. Then $E_{\widetilde{H}}\restriction_{\overline{H}} \leq_B F$, where $E_{\widetilde{H}}$ is the equivalence relation of lying in the same coset of $\widetilde{H}$ on $\widetilde{G}$ (note that right cosets are the same as left cosets, as $\widetilde{H}$ is normal).*

Proof. As the first step, we need to show this.

Claim. There is a Borel function $\Psi\colon S_H(M) \to S_{\widetilde{H}}(M)$ such that $\Psi(\mathrm{tp}(g/M)) = \mathrm{tp}(a_g, g/M)$.

Proof. Each $p \in S_H(M)$ defines naturally a closed $[p] \subseteq S_{\overline{H}}(M)$. Since $\widetilde{H}$ is generated by a countable family of type-definable sets, we can write $S_{\widetilde{H}}(M) = \bigcup_{i\in\mathbf{N}} D_i$ for closed $D_i \subseteq S_{\overline{H}}(M)$. Then, we put

$$
D_i' = \{p \in S_H(M) \mid [p] \cap D_i \neq \emptyset \wedge (\forall j < i)[p] \cap D_j = \emptyset\}
$$

(Putting it differently, each D'_i is the projection of D_i onto $S_H(M)$ minus the projections of all D_j for $j < i$.)

Each D'_i is a Borel set in $S_H(M)$ (as a difference of two compact sets), and $\pi[\widetilde{H}] = H$, so in fact $S_H(M) = \coprod_i D'_i$ (where $\coprod$ is the disjoint union).

Consider $F(S_{\overline{H}}(M))$. We can now put $\Phi\colon S_H(M) \to F(S_{\overline{H}}(M))$ defined by $\Phi(p) = [p] \cap D_i$ if $p \in D'_i$. We need Φ to be Borel. To show that, it is enough to show that its restriction to any D'_i is Borel (because they form a countable, Borel cover of $S_H(M)$). But for each restriction $\Phi_i = \Phi\restriction_{D'_i}$ and open U

$$\Phi_i^{-1}[B_U] = D'_i \cap \{p \in S_H \mid [p] \cap D_i \cap U \neq \emptyset\}$$

which is Borel because $D_i \cap U$ is F_σ, so Φ_i is Borel and so is Φ.

Then the function $\Psi = d \circ \Phi$, where $d\colon F(S_{\overline{H}}(M)) \to S_{\overline{H}}(M)$ is given by the Kuratowski–Ryll-Nardzewski theorem (which we can apply because $\overline{H}$ is type-definable, so $S_{\overline{H}}(M)$ is a compact Polish space) is Borel as a composition of two Borel maps.

It is easy to see that Ψ is into $S_{\widetilde{H}}(M)$, as each D_i is a subset of $S_{\widetilde{H}}(M)$, and since $\Psi(p) \in [p]$, we immediately get the claim. □(claim)

Next, we define a function $f\colon S_{\overline{H}}(M) \to A_H/A_1$ by

$$f(\mathrm{tp}((a,g)/M)) = (a - a_g) + A_1$$

where a_g is such that $\Psi(\mathrm{tp}((a,g)/M)) = \mathrm{tp}((a_g,g)/M)$. We need to show that:

(1) f is a well-defined, Borel function

(2) $\mathrm{tp}((a_1,g_1)/M) E_{\widetilde{H}} \mathrm{tp}((a_2,g_2)/M) \iff (a_1 - a_{g_1}) - (a_2 - a_{g_2}) \in \widetilde{H} \cap A$

For the first point, we can see that $f = f_3 \circ f_2 \circ f_1$, where

$$\begin{aligned} f_1\colon & S_{\overline{H}}(M) \to S_{A_H}(M) \times S_H(M) \\ & \mathrm{tp}(a,g/M) \mapsto (\mathrm{tp}(a/M), \mathrm{tp}(g/M)) \\ f_2\colon & S_{A_H}(M) \times S_H(M) \to S_{A_H}(M) \times S_{\widetilde{H}}(M) \\ & (\mathrm{tp}(a/M), \mathrm{tp}(g/M)) \mapsto (\mathrm{tp}(a/M), \Psi(\mathrm{tp}(g/M))) \\ f_3\colon & S_{A_H}(M) \times S_{\widetilde{H}}(M) \to A_H/A_1 \\ & (\mathrm{tp}(a/M), \mathrm{tp}(b,g)/M) \mapsto (a-b)/A_1 \end{aligned}$$

f_1, f_2 are clearly Borel and well-defined, so it remains to show that so is f_3.

Suppose $a \equiv_M a'$ and $(b,g) \equiv_M (b',g')$. Then also $b \equiv_M b'$, so $a - a' \in (A_H)^{000}$ and $b - b' \in (A_H)^{000}$ (by Proposition 3.3, because lying in the same coset of $(A_H)^{000}$ is clearly a bounded and invariant equivalence relation), and hence

$$(a - a') - (b - b') = (a - b) - (a' - b') \in (A_H)^{000}$$

But A_1 is invariant of bounded index, so $(A_H)^{000} \subseteq A_1$ and so f_3 is well defined. It is not hard to see that it is continuous: the preimage of a closed set in A_H/A_1 is the set of pairs $\mathrm{tp}(a/M), \mathrm{tp}(b,g/M)$ with $a - b$ belonging to some type-definable (over M, e.g. by Corollary 3.10) set. This is clearly a type-definable (over M) condition about a, b, so the preimage is closed.

The second point follows immediately from the previous lemma. □

Having proved all the necessary lemmas, we have all but finished the proof of the main theorem.

Proof of Theorem 6.2. Put $A_1 = \widetilde{H} \cap A$. Then A_1 satisfies the hypotheses of the previous lemma, and in this case F is simply equality on A/A_1 (a Polish space by Corollary 3.11) and hence smooth, so $E_{\widetilde{H}}\restriction_{\overline{H}}$ is smooth. On the other hand, by virtue of Corollary 4.10, this implies that $\widetilde{H}$ is type-definable. □

6.3. **Corollaries of Theorem 6.2.** In [GK13], the authors state the following two (equivalent, as we will see shortly) conjectures.

Conjecture 8 (Conjecture 2.11. in [GK13])**.** *Suppose we have a definable extension of a definable group G by a definable abelian group A, corresponding to the short exact sequence*

$$0 \to A \to \widetilde{G} \to G \to 0$$

with the 2-cocycle $h\colon G^2 \to A$. Assume that h has finite range and that $G^{00} = G^{000}$.

Then the conjecture is that for any invariant $\widetilde{H} \leq \widetilde{G}$ of bounded index, and such that $\widetilde{H} \cap A$ is type-definable, we have

$$\widetilde{G}^{00} \cap A \subseteq \widetilde{H} \cap A$$

Conjecture 9 (Conjecture 2.10 in [GK13])**.** *Assume we have G, A, h as in the first paragraph of the previous conjecture. Additionally let A_1 be a type-definable subgroup of A of bounded index and invariant under the action of G.*

Then the conjecture is that

$$\widetilde{G}^{00} \cap A \subseteq A_1 \iff \widetilde{G}^{000} \cap A \subseteq A_1$$

Remark. The preceding conjectures are important mostly for two reasons:

(1) If proven in general (without assumptions of countability), it would imply that Corollary 2.8 in [GK13] holds in general, i.e. that in Theorem 6.1, if $G^{00}_B = G^{000}_B$, then the assumption (i) (about non-splitting of the modified 2-cocycle, with other assumptions notwithstanding) not only implies, but is equivalent to $\widetilde{G}^{00} \neq \widetilde{G}^{000}$ (this explained in more detail in [GK13], but omitted here, as we only prove the countable case of the conjecture).
(2) It implies that, in a rather general context, the quotient $\widetilde{G}^{00}/\widetilde{G}^{000}$ is isomorphic (algebraically) to the quotient of a compact group by a finitely generated dense subgroup (this will be revisited at the end of this section).

The next fact tells us that the two conjectures are equivalent, and more.

Fact 6.6. *Suppose $\widetilde{G}$ is a definable extension of a definable group G by a definable abelian group A (so the cocycle might have infinite range, but it must be definable). Then the following are equivalent:*

(1) *For any invariant $\widetilde{H} \leq \widetilde{G}$ of bounded index, such that $\widetilde{H} \cap A$ is type-definable, we have $\widetilde{G}^{00} \cap A \subseteq \widetilde{H}$.*
(2) *For any $A_1 \leq A$ type-definable, G-invariant, such that $\widetilde{G}^{000} \cap A \subseteq A_1$ we also have $\widetilde{G}^{00} \cap A \subseteq A_1$.*

In particular, Conjecture 8 and Conjecture 9 are equivalent.

Proof. To see that the (1) implies (2), take an A_1 as in the assumptions of (2). $\widetilde{G}^{000}$ is normal in $\widetilde{G}$, so $H = A_1\widetilde{G}^{000}$ is a subgroup of $\widetilde{G}$. Furthermore, it is invariant (as a product of two invariant sets) and of bounded index (because it contains $\widetilde{G}^{000}$), and its intersection with A is just A_1 (because A_1 contains $\widetilde{G}^{000} \cap A$), which is type-definable, so by (1) it contains $\widetilde{G}^{00} \cap A$.

To see the opposite implication, notice that if we have $\widetilde{H}$ as in (1), then $A_1 = \mathrm{Core}_{\widetilde{G}}(\widetilde{H} \cap A)$ (where $\mathrm{Core}_K(L)$ is the largest normal subgroup of K contained in L) satisfies the assumptions of (2):

- it is type-definable because $\widetilde{H} \cap A$ is type-definable, and A_1 is an invariant intersection of its conjugates (which are pseudo-closed),
- it contains $\widetilde{G}^{000} \cap A$ (and in particular has bounded index in A), because $\widetilde{H}$ contains $\widetilde{G}^{000}$ and $\widetilde{G}^{000} \cap A$ is normal in $\widetilde{G}$, and

- it is G-invariant because it is normal in $\widetilde{G}$ (by the definition).

Therefore, A_1 contains $\widetilde{G}^{00} \cap A$. But A_1 is contained in $\widetilde{H}$, so $\widetilde{H}$ contains $\widetilde{G}^{00} \cap A$. $\square$

Remark. The authors of [GK13] actually allow a finite parameter set B over which the cocycle h is definable, they calculate the connected components over this set, and they assume that H is B-invariant in Conjecture 8 and A_1 is type-definable over B in Conjecture 9. But since B is finite, we may add constants for its elements to the language, and it remains countable, while none of the properties relevant to the previous conjectures change, so we assume without loss of generality that $B = \emptyset$.

We will also drop the requirement that h has finite range, as it is not needed for the subsequent discussion, which leaves us in the general context of Theorem 6.2, only with the additional assumption that $G^{00} = G^{000}$.

The assumption that $G^{00} = G^{000}$ allows us to make the following observation.

Fact 6.7 ([GK13](Remark 2.1(iii)))**.** *If A_1 is a type-definable, bounded index subgroup of A, invariant under the action of G (i.e. normal in $\widetilde{G}$), then G^{00} acts trivially on A/A_1.*

Proof. A_1 is G-invariant, so G acts naturally on A/A_1, yielding an abstract homomorphism $f\colon G \to S(A/A_1)$ (where $S(A/A_1)$ is the abstract permutation group). Notice that

$$\ker(f) = \{g \in G \mid (\forall a)\, g \cdot a - a \in A_1\}$$

is a type-definable subgroup of G, and it has bounded index, because $S(A/A_1)$ is small. Therefore, it contains G^{00}. $\square$

This leads us to the next corollary.

Corollary 6.8. *Suppose $\widetilde{H} \trianglelefteq \widetilde{G}$ is an invariant subgroup of bounded index, contained in $\widetilde{G}^{00}$ and F_σ (i.e. generated by a countable family of type-definable sets).*

Then $\widetilde{H}$ is type definable (and therefore equal to $\widetilde{G}^{00}$) if and only if $\widetilde{H} \cap A$ is type-definable and $\pi[\widetilde{H}] \leq G$ is type-definable.

Proof. $\Rightarrow$ is clear. For $\Leftarrow$ notice that $\pi[\widetilde{H}]$ being type-definable implies that it is, in fact, equal to G^{00}, and by the previous fact it acts trivially on $A/(\widetilde{H} \cap A)$, so the result follows immediately from Theorem 6.2 with $\overline{H} = \widetilde{G}^{00}$. $\square$

And finally we obtain a proof of Conjecture 8.

Corollary 6.9. *Suppose $\widetilde{H} \leq \widetilde{G}$ is an invariant subgroup of bounded index, and that $G^{00} = G^{000}$. Suppose in addition that $\widetilde{H} \cap A \cap \widetilde{G}^{00}$ is type-definable.*

Then $\widetilde{H} \cap A \supseteq \widetilde{G}^{00} \cap A$. (In particular, the countable case of Conjecture 8 holds, without any assumptions on h beyond definability.)

Proof. We intend to apply the preceding corollary, so we need to modify H to satisfy its assumptions.

(1) We can assume without loss of generality that $\widetilde{H}$ is normal, because we can replace it with $\mathrm{Core}(\widetilde{H}) = \bigcap_{g \in \widetilde{G}} g\widetilde{H}g^{-1}$: the latter is obviously normal and invariant, it has bounded index, because it contains G^{000}, and because $A, \widetilde{G}^{00}$ are normal, we have the identity

$$\mathrm{Core}(\widetilde{H}) \cap A \cap \widetilde{G}^{00} = \mathrm{Core}(\widetilde{H} \cap A \cap \widetilde{G}^{00})$$

so in particular, the left hand side is type-definable, and hence $\mathrm{Core}(\widetilde{H})$ satisfies the assumptions of this corollary and it is contained in $\widetilde{H}$.

(2) Secondly, can assume without loss of generality that $\widetilde{H} \leq \widetilde{G}^{00}$, by replacing $\widetilde{H}$ with $\widetilde{H} \cap \widetilde{G}^{00}$.
(3) Thirdly, we can also assume that $\widetilde{H}$ is generated by a countable family of type-definable sets by replacing it with $(\widetilde{H} \cap A) \cdot \widetilde{G}^{000}$ This is a subgroup because $\widetilde{G}^{000}$ is normal and it is generated by $\widetilde{H} \cap A$ and the countable family of type-definable sets generating $\widetilde{G}^{000}$, which exists by Fact 2.5. Moreover, $\widetilde{G}^{000} \leq \widetilde{H}$, so $(\widetilde{H} \cap A) \cdot \widetilde{G}^{000} \leq \widetilde{H}$.
(4) Then we can apply the previous corollary to deduce that $H = \widetilde{G}^{00}$, so trivially $H \cap A \supseteq \widetilde{G}^{00} \cap A$.

□

Corollary 6.10 (As suggested by Remark 2.16 in [GK13])**.** *Suppose that $\widetilde{G}$ is a definable extension of a definable group G by an abelian group A, and that $G^{00} = G^{000}$.*

Let $A_1 \leq \widetilde{G}^{000} \cap A$ be a G-invariant, type-definable subgroup of A of bounded index contained in $\widetilde{G}^{000}$. Then $\left(\widetilde{G}^{000} \cap A\right)/A_1$ is dense in $\left(\widetilde{G}^{00} \cap A\right)/A_1$ (with logic topology).

Proof. Let A_2 be the preimage of the closure of $(\widetilde{G}^{000} \cap A)/A_1$ by the quotient map $\pi\colon A \to A/A_1$. Then A_2 is a type-definable, G-invariant subgroup of A, which contains $\widetilde{G}^{000} \cap A$.

Since A_2 is G-invariant, $\widetilde{H} = A_2 \cdot \widetilde{G}^{000}$ is an invariant subgroup of $\widetilde{G}$, and $\widetilde{H} \cap A = A_2$ is type-definable, as is $\widetilde{H} \cap A \cap \widetilde{G}^{00}$, so by Corollary 6.9, $A_2 = \widetilde{H} \cap A \supseteq \widetilde{G}^{00} \cap A$, and therefore A_2/A_1 – the closure of $(\widetilde{G}^{000} \cap A)/A_1$ in A/A_1 – contains $(\widetilde{G}^{00} \cap A)/A_1$, which was to be shown. □

This corollary, along with the comments following [GK13](Proposition 2.14), implies that in a somewhat general setting, the group $\widetilde{G}^{00}/\widetilde{G}^{000}$ is (abstractly) isomorphic to the quotient of a compact abelian group by a dense subgroup (which is finitely generated if the range of h is finite) – which is analogous of the fact that for G definable in o-minimal expansions of a real closed field (e.g. those considered in [CP12]), G^{00}/G^{000} is abstractly isomorphic to the quotient of a compact, abelian Lie group by a dense, finitely generated subgroup. This furthers the analogy between [GK13] and [CP12].

Remark. As stated in the introduction, the facts we have shown apply in the countable case. It is likely that similar (though more elaborate) methods could yield a proof of Conjecture 9 (or equivalently Conjecture 8) in general, but this is beyond the scope of this paper.

References

[GK13] Jakub Gismatullin and Krzysztof Krupiński. *On model-theoretic connected components in some group extensions.* Version 2. (Submitted.) 2013. arXiv: 1201.5221v2. (Visited on 06/02/2014) (cit. on pp. 4, 6, 7, 36, 37, 39–41).

[KM13] Itay Kaplan and Benjamin D. Miller. *An embedding theorem of $\mathbb{E}_0$ with model theoretic applications.* Version 1. (Submitted.) Aug. 26, 2013. arXiv: 1308.5515. (Visited on 06/02/2014) (cit. on pp. 4, 6, 26, 32).

[KMS13] Itay Kaplan, Benjamin Miller, and Pierre Simon. *The Borel cardinality of Lascar strong types.* Version 1. (Accepted pending revision.) 2013. arXiv: 1301.1197. (Visited on 06/02/2014) (cit. on pp. 4, 6, 18, 24–26, 28).

[KPS13] Krzysztof Krupiński, Anand Pillay, and Sławomir Solecki. "Borel equivalence relations and lascar strong types". In: *J. Math. Log.* 13.2 (Dec. 2013). DOI: 10.1142/S0219061313500086 (cit. on pp. 4–6, 12, 14, 15, 28, 34, 35, 37).

[CP12] Annalisa Conversano and Anand Pillay. "Connected components of definable groups and -minimality I". In: *Adv. Math.* 231.2 (Oct. 1, 2012), pp. 605–623. DOI: 10.1016/j.aim.2012.05.022 (cit. on pp. 36, 41).

[Gis11] Jakub Gismatullin. "Model theoretic connected components of groups". In: *Isr. J. Math.* 184.1 (Aug. 1, 2011), pp. 251–274. DOI: 10.1007/s11856-011-0067-8 (cit. on p. 11).

[GN08] Jakub Gismatullin and Ludomir Newelski. "G-Compactness and Groups". In: *Arch. Math. Logic* 47.5 (2008), pp. 479–501 (cit. on pp. 11, 12).

[Kan08] Vladimir Kanovei. *Borel Equivalence Relations*. American Mathematical Soc., 2008. 260 pp. (cit. on p. 9).

[Pet07] Y. Peterzil. "Pillay's conjecture and its solution – a survey". In: *Logic Colloquium 2007*. 2007 (cit. on p. 12).

[Pil04] Anand Pillay. "Type-Definability, Compact Lie Groups, and o-Minimality". In: *J. Math. Log.* 4.2 (2004), pp. 147–162 (cit. on p. 12).

[New03] Ludomir Newelski. "The diameter of a Lascar strong type". In: *Fund. Math.* 176.2 (2003), pp. 157–170. DOI: 10.4064/fm176-2-4 (cit. on pp. 6, 18).

[KN02] Krzysztof Krupiński and Ludomir Newelski. "On Bounded Type-Definable Equivalence Relations". In: *Notre Dame J. Formal Logic* 43.4 (2002), pp. 231–242. DOI: 10.1305/ndjfl/1074396308 (cit. on p. 16).

[Rot02] Joseph J. Rotman. *Advanced Modern Algebra*. American Mathematical Soc., 2002. 1026 pp. (cit. on p. 35).

[Cas+01] E. Casanovas, D. Lascar, A. Pillay, and M. Ziegler. "Galois Groups of First Order Theories". In: *J. Math. Log.* 1.2 (2001), pp. 305–319 (cit. on pp. 5, 9–11, 31).

[Kec99] Alexander S. Kechris. "New Directions in Descriptive Set Theory". In: *Bull. Symbolic Logic* 5.2 (June 1999), pp. 161–174 (cit. on p. 8).

[Kim98] Byunghan Kim. "A Note on Lascar Strong Types in Simple Theories". In: *J. Symbolic Logic* 63.3 (Sept. 1998), pp. 926–936 (cit. on p. 10).

[BK96] Howard Becker and A. S. Kechris. *The Descriptive Set Theory of Polish Group Actions*. Cambridge University Press, Dec. 5, 1996. 152 pp. (cit. on p. 9).

[Kec95] Alexander S. Kechris. *Classical Descriptive Set Theory*. Springer, 1995. 427 pp. (cit. on pp. 8, 13, 37).

[Las82] Daniel Lascar. "On the Category of Models of a Complete Theory". In: *J. Symbolic Logic* 47.2 (June 1982), pp. 249–266 (cit. on p. 9).