J. Inst. Math. Jussieu (2026), 1–76
doi:10.1017/S1474748026102035

TOPOLOGICALLY 1-BASED T-MINIMAL STRUCTURES

BENJAMIN CASTLE[1], ASSAF HASSON[2] AND
WITH AN APPENDIX B BY WILL JOHNSON[3]

[1] *Mathematics, University of Illinois Urbana-Champaign, Urbana, United States*
(btcastl2@illinois.edu)
[2] *Mathematics, Ben-Gurion University of the Negev, Be'er Sheva, Israel*
(hassonas@math.bgu.ac.il)
[3] *School of Philosophy, Fudan University, Shanghai, China*
(willjohnson@fudan.edu.cn)

(Received 30 August 2025 UTC; revised 11 July 2026 UTC; accepted 26 August 2026 UTC)

Abstract We prove group-existence and structure theorems in a general setting of tame topological theories. More precisely, we identify a linear/non-linear dividing line – called *topological 1-basedness* – among the class of t-minimal theories with the independent neighbourhood property. This is a wide class including all visceral theories, as well as all dense weakly o-minimal and C-minimal theories (even those where exchange fails).

Now assume $\mathcal{M}$ is highly saturated and t-minimal with the independent neighbourhood property. We show that if $\mathcal{M}$ is non-trivial and topologically 1-based, it admits a type-definable abelian group $(G,+)$ with G an open subseteq of M. Moreover, we can ensure that G is a topological group with the subspace topology inherited from M; and in this case, we show that the induced structure on G satisfies an appropriate topological analogue of the Hrushovski-Pillay classification of 1-based stable groups.

Contents

1 Introduction **2**

2 T-minimality and germs **8**

3 Independent Neighbourhoods **14**

Keywords and phrases: t-minimality; 1-based; group configuration; type definable groups; topological model theory

Mathematics Subject Classification: Primary 03C65

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

4 Infinitesimal neighborhoods 22

5 Topological 1-basedness 27

6 The structure theorem for topologically 1-based types 36

7 Groupoid spines 42

8 Type-definability requirements 46

9 Construction of a regular groupoid spine 50

10 Topologizing the group 55

11 The structure of the group 61

Appendix A Proving the independent neighbourhood property 67

Appendix B Independent neighbourhoods in visceral theories, by Will Johnson 71

References 75

1. Introduction

In this paper, we introduce the notion of a *topologically 1-based t-minimal structure*, and prove that when such a structure is non-trivial, it admits an infinite type-definable abelian topological group, whose (local) structure we then proceed to study. These results can be viewed as a topological analogue of the literature on 1-based stable structures: On the one hand, the existence of a type-definable group resembles Hrushovski's characterization of locally modular regular types ([9], following earlier works of Zilber [21, §3] and Cherlin-Harrington-Lachlan [4]. On the other hand, our structural analysis of the group closely resembles the Hrushovski-Pillay analysis of 1-based stable groups [10].

Our main theorem reads as follows (see Theorem 11.15; the relevant terms are defined in Definitions 2.2, 3.1, 9.7 and 5.1):

Theorem 1. *Let $\mathcal{M}$ be a sufficiently saturated t-minimal structure with the independent neighbourhood property. Assume that $\mathcal{M}$ is non-trivial and topologically 1-based. Then there are a countable parameter set A, and an A-type-definable abelian group G, such that the following hold:*

(1) *G is open in $\mathcal{M}$.*

(2) *G is a topological group with the topology inherited from M.*

(3) *G is locally linear.*

We explain the terminology more below; for now, we note that the above theorem encompasses all dense weakly o-minimal and C-minimal structures; in particular, as far

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

as we know, it is the first group-existence theorem for (non-o-minimal) weakly o-minimal structures.

1.1. Background

To recall, 1-basedness in stable theories is an abstract notion viewed as capturing those stable theories of *linear* complexity. Classical results in stability show that in non-trivial 1-based theories, any instance of non-triviality can be explained in terms of type-definable abelian groups (e.g. [19, Corollary 5.4.13]. The key notion allowing the construction of such groups is the *germ* of a map on a complete (stationary) type – where the word *germ* here indicates that maps are viewed up to agreement on a non-forking extension of the domain.

Soon after the aforementioned works on 1-basedness in stable theories, it became clear that *o-minimality* provided a suitable tame (yet unstable) setting where many stability-like results could be developed. In particular, o-minimal structures were seen to exhibit a similar dimension theory to strongly minimal structures (and thus many notions could be translated from stability to o-minimality). To that end, early works of Peterzil and Loveys [16],[12] developed a version of linearity for o-minimal structures, and obtained similar group-existence and structure theorems (which ultimately became part of the o-minimal trichotomy theorem, [18]. The key notion here was again the germ of a function – but in the o-minimal setting, germs retain their original topological meaning (one treats functions up to agreement on an interval).

Many aspects of Peterzil's group-existence proof [16] are quite general, but the definition of linearity is rather specific to o-minimality. In a later work, Maalouf [13] adapted the definition of linearity to the related context of *C-minimality*, and proved a similar group-existence theorem – however, Maalouf needed to assume the *exchange lemma*, which can fail in a general C-minimal structure.

The above works on o-minimality and C-minimality naturally lead to the following question: is there a general notion of linearity for 'topologically tame' structures, and can one use it to prove abstract group-existence and structure theorems? For example, what can be said for *weakly o-minimal* structures, and for C-minimal structures without exchange?

1.2. T-minimality and work of Johnson

Answering the above question would require a general setting for topological tameness, and this general setting should admit a suitable dimension theory for definable sets and types. Recently, Johnson [11] has shown that *t-minimality* provides such a setting. Recall that a (sufficiently saturated) structure is t-minimal if (1) it admits a uniformly $\emptyset$-definable Hausdorff topology, and (2) a unary definable set is infinite if and only if it has non-empty interior with respect to the given topology. T-minimality encompasses all dense weakly o-minimal and C-minimal structures (regardless of exchange), in addition to many valued fields and expansions thereof (e.g. all 1-h-minimal fields, and all unstable dp-minimal expansions of fields).

Johnson's paper develops a rudimentary dimension theory assuming only t-minimality, then proceeds to develop stronger geometric properties (e.g. generic continuity of

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

functions) assuming a stronger condition called *viscerality* (saying roughly that the topology is 'uniform'). In our setting, we do not want to assume viscerality: the issue is that in many cases viscerality follows from the existence of a group structure, and we want to work a priori without a group structure. For example, while all o-minimal groups are visceral, there are also non-visceral linear o-minimal structures (and our results should apply to them).

Instead, we augment t-minimality with a weaker condition, which we call the *independent neighbourhood property*: roughly speaking, this property says that one can always 'zoom in' on a tuple without changing its dimension. The same property was already studied in the dp-minimal visceral context in [7], where it proved quite useful. The key features of the independent neighbourhood property in our case are:

- It encompasses all visceral, weakly o-minimal and C-minimal structures.
- It implies many of the same geometric properties (including generic continuity) that Johnson's paper proves for visceral structures.

In light of the above discussion, our main theorem shows that the known group-existence and structure theorems for o-minimal and C-minimal structures are special cases of a general result, valid for all t-minimal structures with the independent neighbourhood property. Our proof follows a similar (but not identical) outline to the original proof of [16], with various new challenges encountered throughout. In particular, note that we *do not* assume exchange (and this causes quite a few difficulties); as a trade-off, our version of non-triviality needs to be a bit stronger than the usual one (see Definition 9.7).

1.3. Topological 1-basedness

The main task in our proof is to identify and study a linearity notion – namely *topological 1-basedness* – in the absence of either exchange or an order (recall that Peterzil and Loveys' o-minimal version used both exchange and the order). Let us now discuss our version in more depth.

Just as *germs* play a key role in the stable and o-minimal group-existence theorems, they also play a key role in our setting. In fact, it follows from Johnson's paper that in a t-minimal structure, every complete type $\mathrm{tp}(a/A)$ has a well-defined germ at any realization a (see Lemma 2.9), and this germ can be coded as an imaginary element $\mathrm{germ}(a/A)$.

When one constructs groups in the stable setting, germs are manifested definably as *canonical bases*. Thus, it is natural to view the imaginary $\mathrm{germ}(a/A)$ as a sort of 'topological canonical base' of $\mathrm{tp}(a/A)$, localized to the point a. There is good evidence for this line of thinking – for example, assuming the independent neighbourhood property, if $B \supseteq A$, then one has $\dim(a/B) = \dim(a/A)$ if and only if $\mathrm{germ}(a/B) = \mathrm{germ}(a/A)$ (see Lemma 3.7).

In light of this analogy, our definition of topological 1-basedness essentially amounts to replacing canonical bases in the stability-theoretic definition with germs. There is one difference, however, which makes our definition a bit more general. Recall that in stability, 1-basedness requires the canonical base of a type to be algebraic over a single realization. Since the canonical base is imaginary, the word 'algebraic' here refers to the operator acl^{eq}. In geometric structures, Gagelman extends the dimension theory to imaginary elements, where the 0-dimensional closure may be a proper coarsening of acl^{eq}. In our

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

view, it is this 0-dimensional closure operator (rather than acl^{eq}) which correctly captures linearity in our setting. Of course, 0-dimensionality only makes sense assuming exchange (i.e. in geometric structures). However, one can also characterize it topologically. That is, in geometric t-minimal structures, 0-dimensional imaginaries are those associated with definable equivalence relations whose classes are *open* (or at least those which have interior in some ambient set). Thus, dropping the assumption of exchange and keeping only t-minimality (with independent neighbourhoods), our notion of linearity is the following: say that $\mathrm{tp}(a/Ab)$ is *topologically 1-based over* A if the following equivalent conditions hold (see Lemma 5.6):

- $\dim(b/A\mathrm{germ}(a/Ab)) = \dim(b/Aa)$.
- The map $y \mapsto \mathrm{germ}(a/Ay)$ is constant on a neighbourhood of b in $\mathrm{tp}(b/Aa)$.

Then say that $\mathcal{M}$ is topologically 1-based if $\mathrm{tp}(a/b)$ is topologically 1-based over $\emptyset$ for all real tuples a and b.

Note that this definition generalizes known versions of linearity in the literature. For example, our main structure theorem for topologically 1-based types, Theorem 6.4, shows that topological 1-basedness is equivalent to an appropriately localized version of Pillay's notion of weak normality. Outside of the stable context, in, e.g. [1], Berenstein and Vassiliev develop a version of linearity in the related context of *geometric structures*, which they call *weak 1-basedness*. When restricted to o-minimality, weak 1-basedness is shown to capture exactly the linear o-minimal structures in the sense of Peterzil and Loveys. Now suppose $\mathcal{M}$ is t-minimal with the independent neighbourhood property, and $\mathcal{M}$ moreover satisfies exchange. Then $\mathcal{M}$ is a geometric structure. In Corollary 6.5, we show that such an $\mathcal{M}$ is topologically 1-based if and only if it is weakly 1-based as a geometric structure – thus, topological 1-basedness also generalizes linearity in o-minimal structures (see Corollary 5.18).

1.4. Infinitesimals and regularity

Once we have the notion of topological 1-basedness, our task is to use it to build type-definable groups. There are two key notions involved here. The first, following earlier works such as Peterzil's o-minimal group configuration of [17], is the *infinitesimal neighbourhood* $\mu(a/A)$ of a tuple a over a parameter set A. This refers to the set of realizations of $\mathrm{tp}(a/A)$ from a highly saturated elementary extension which are 'infinitesimally close' to a in the sense of the original model (see Definition 4.3). Infinitesimal neighbourhoods allow us to replace finite-to-finite correspondences, locally, with homeomorphisms. For example, suppose a and b are interalgebraic over A; then using generic continuity, we show that $\mu(ab/A)$ is the graph of a homeomorphism $\mu(a/A) \to \mu(b/A)$ (see Lemma 4.9). By repeated applications of this fact, it follows that if a and b are interalgebraic over At, then $\mu(abt/A)$ is the graph of a $\mu(t/A)$-indexed family of homeomorphisms $\mu(a/A) \to \mu(b/A)$. Thus, we have a convenient tool for constructing families of homeomorphisms – and such objects will ultimately be the building blocks of a type-definable group action.

The other key notion is *regularity*. Using infinitesimal neighbourhoods, one can construct (for example) three points $a,b,c \in M$, a set A and three compatible type-

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

definable families of homeomorphisms $\mu(a/A) \to \mu(b/A)$, $\mu(b/A) \to \mu(c/A)$ and $\mu(a/A) \to \mu(c/A)$ (compatible meaning that any composition $\mu(a/A) \to \mu(b/A) \to \mu(c/A)$ agrees with one of the given maps $\mu(a/A) \to \mu(c/A)$). One then needs to find additional conditions assuring that (1) these families of maps arise from a group (or rather, at this level, a groupoid) and (2) the resulting groupoid is type-definable in an appropriate sense.

The key notion is regularity (inspired by regular group actions): say that a family of maps $f_t : X \to Y$ is *regular* if for any $(x,y) \in X \times Y$ there is t with $f_t(x) = y$, and if any two such maps either agree everywhere or agree nowhere. Regular families of maps can be seen as a manifestation of linearity (or 1-basedness), in the sense that a single point determines an entire map in the family. In Sections 7 and 8, we study type-definable systems of compatible regular families of maps in more detail, showing that in many cases such configurations give rise to type-definable group actions (this is in some sense a completely abstract group configuration theorem – see Theorem 8.8). Then in Section 9, we show that such collections of regular families arise naturally under the assumption of topological 1-basedness. The key input here is a structure theorem for topologically 1-based types (see Theorem 6.4), which roughly says that such types always look regular in an infinitesimal neighbourhood (as mentioned above, this is a topological analogue of the *weakly normal* decomposition in [10]):

Theorem 2. *Let $\mathcal{M}$ be a sufficiently saturated t-minimal structure with the independent neighbourhood property. Let a and b be real tuples, and A a parameter set. Then the following are equivalent:*

(1) $\mathrm{tp}(a/Ab)$ *is topologically 1-based over A.*

(2) *Any two fibres in the projection $\mu(ab/A) \to \mu(b/A)$ are either equal or disjoint.*

So, condition (2) of the above theorem is how we prove that the families of homeomorphisms we build are regular, and thus gives rise to a type-definable group.

1.5. Summary of the paper

The paper is roughly organized into four main parts:

- Sections 2–4 study t-minimal theories and the independent neighbourhood property in general, focusing on germs and generic continuity. These sections do not involve topological 1-basedness or groups. Section 2 constructs and studies germs in arbitrary t-minimal theories; Section 3 adds the independent neighbourhood property to the background assumptions and develops generic continuity; and Section 4 rewrites the key points of the previous sections in the language of infinitesimal neighbourhoods.
- Sections 5–6 introduce topological 1-basedness. Section 5 is concerned mainly with the definition and basic properties. For example, we prove that topological 1-basedness is preserved under adding and removing constants (Lemma 5.7) and taking reducts (Lemma 5.9), and that a topologically 1-based structure cannot admit an infinite definable field (Lemma 5.15). Then Section 6 is devoted to the main structure theorem for topologically 1-based types, Theorem 2 above.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

- Sections 7–8 leave the topological setting entirely and focus on a completely abstract version of group configurations, which we call *regular groupoid spines*. This is something like a regular groupoid, but where only certain ordered pairs of objects have morphisms between them. Section 7 shows that any regular groupoid spine with at least three objects arises from a genuine regular groupoid (this is really a group configuration theorem). Then Section 8 shows that if a regular groupoid spine is type-definable in an appropriate sense, so is the groupoid it arose from (and thus so is the automorphism group at any of its objects).
- Finally, Sections 9–11 combine the previous two parts to prove our main theorems. Suppose that $\mathcal{M}$ is non-trivial and topologically 1-based. In Section 9, we use infinitesimal neighbourhoods to construct a type-definable regular groupoid spine with three objects in $\mathcal{M}$ (precisely, every required property other than regularity comes from the general theory of infinitesimal neighbourhoods, and regularity comes from Theorem 2). Thus, by the results of Sections 7–8, we obtain a type-definable group in an elementary extension of $\mathcal{M}$. Then in Section 10, we show that the group can be transferred down to $\mathcal{M}$, and that – by virtue of generic continuity – it can be tweaked slightly to make it a topological group with the induced topology from $\mathcal{M}$. Finally, Section 11 shows that the group is locally abelian and locally linear. The main tool here is again Theorem 2. The idea is that any two infinitesimal translates of an infinitesimal neighbourhood (in an appropriate sense) are equal or disjoint – and this implies that the given infinitesimal neighbourhood is a coset of a subgroup. Thus, the local structure of definable sets can be analysed entirely in terms of subgroups and their cosets, in a similar manner to the Hrushovski-Pillay analysis of 1-based stable groups [10].

  We can also give some local analysis. Namely, using the notion of topologically 1-based types, we are able to define what it means for a one-dimensional type to be *topologically locally modular* (even if the ambient structure is not topologically 1-based as a whole). In this case (assuming non-triviality), we are able to build a type-definable topological group on an open subseteq of the realizations of the type (see Theorem 10.11). However, the structural results for the group (local linearity and local abelianity) need full topological 1-basedness, at least for all tuples from the group itself.

Finally, we give a brief appendix devoted to examples of t-minimal structures with the independent neighbourhood property, covering some non-visceral structures, most notably weakly o-minimal and C-minimal structures. We also include a complementary appendix by Will Johnson, proving that visceral theories have the independent neighbourhood property.

1.6. Notation and model-theoretic generalities

We use standard model-theoretic notation and terminology. Structures are denoted by calligraphic capital letters ($\mathcal{M}$,$\mathcal{N}$ etc.), and their universes are denoted by the corresponding Latin letters M, N. All tuples are assumed to be finite. As is customary, by *real* tuples, we refer to tuples in M^n (some n), as opposed to *imaginary* tuples which

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

may reside in $\mathcal{M}^{eq}$. **Unless explicitly stated otherwise, all parameter sets are allowed to contain imaginary elements.**

Our structures will typically be assumed to be *sufficiently saturated*. To clarify, when we say $\mathcal{M}$ is *sufficiently saturated*, we mean it is κ-saturated and κ-strongly homogeneous for some uncountable cardinal κ strictly larger than the cardinality of the language of $\mathcal{M}$. Here, κ is unspecified but fixed to go along with the structure $\mathcal{M}$. In this case, the word *small* always means 'of size less than κ' – and by a *parameter set*, we mean a small subseteq of $\mathcal{M}^{eq}$. Thus, smallness and parameter sets are sensitive to the particular choice of model. For example, when studying infinitesimal neighbourhoods, we will be working in an $|M|^+$-saturated model $\mathcal{N} \succ \mathcal{M}$. So M is small with respect to $\mathcal{N}$, but obviously not with respect to $\mathcal{M}$.

By a *type-definable set*, we mean the intersection of a small collection of definable sets. We use throughout that the projection of a type-definable set is again type-definable (over the same parameters). A complete type over a parameter set A is a maximally consistent type (in the language $\mathcal{L}$ expanded with constant symbols for all $a \in A$). For every tuple a and parameter set A, we denote by $\mathrm{tp}(a/A)$ the unique complete type over A realized by a.

As mentioned above, parameter sets are always allowed to contain imaginary elements. This is not always the case for tuples. In order not to overload notation and statements, we will not be pedantic about specifying whether a tuple is real or not, when this is clear from the context. E.g. all topological and geometric properties (such as dimension, genericity, having infinitesimal neighbourhoods etc.) will only refer to real tuples. Typically, unless stated otherwise, we only consider types of real tuples. So, for example, when denoting $\mathrm{tp}(a/b)$, we may allow b to be imaginary (e.g. a germ) but not a. Similarly, we distinguish between the operators $\mathrm{acl}(\cdot)$, $\mathrm{dcl}(\cdot)$ and their respective imaginary ones $\mathrm{acl}^{eq}(\cdot)$, $\mathrm{dcl}^{eq}(\cdot)$ only when confusion may arise. E.g. if b is explicitly allowed to be imaginary, we may use $b \in \mathrm{dcl}(a)$. But if b is not specified, and we want to stress that b could be imaginary, we write $b \in \mathrm{dcl}^{eq}(a)$.

2. T-minimality and germs

From now until the end of Section 6, we fix a sufficiently saturated first-order structure $\mathcal{M}$ and a Hausdorff topology τ on M with a $\emptyset$-definable basis. Note that there is also a $\emptyset$-definable basis for the product topology on M^n, given by products of n basic open sets in M.

We will frequently use the following easy consequence of compactness:

Fact 2.1. *Let $X \subseteq M^n$. Then the open subsets of X are closed under small intersections.*

Proof. Let $\{U_t : t \in T\}$ be a definable basis for the topology on M^n, and let $\{V_i : i \in I\}$ be a small collection of open subsets of X. Suppose we are given $a \in \bigcap V_i$. For each i, choose t_i with $X \cap U_{t_i} \subseteq V_i$. Consider the partial type in y saying that $y \in T$, $a \in U_y$ and $U_y \subseteq U_{t_i}$ for each i. This partial type is finitely satisfiable and small, so it has a realization in $\mathcal{M}$, say b. Then $a \in (X \cap U_b) \subseteq \bigcap V_i$, so a is in the interior of $\bigcap V_i$ in X. As a was arbitrary, this shows that $\bigcap V_i$ is open in X. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

2.1. T-minimality

Recall the following definition (originally appearing in [14]):

Definition 2.2. $\mathcal{M}$ is *t-minimal* if for all definable sets $X \subseteq M$, X is infinite if and only if it has non-empty interior.

Remark 2.3. An equivalent version of the definition is to require that τ has no isolated points and the boundary of any unary definable set is finite. For example, if $X \subseteq M$ is definable with infinite boundary, then the boundary $\partial(X)$ is an infinite definable set with no interior, contradicting the definition. We will use the finiteness of $\partial(X)$ (for unary definable X) freely throughout.

The t-minimal structures are a fairly wide class, including most known instances of topological tameness in model theory. In particular, t-minimality encompasses all weakly o-minimal and C-minimal structures, as well as Henselian valued fields of characteristic zero.

Johnson [11, Definition 2.26, 2.29] has developed a general dimension theory in t-minimal structures. In particular, for any tuple $a \in M^n$ and any set A, one has a dimension $\dim(a/A)$, defined as the length of any minimal tuple $b \subseteq a$ with $a \in \operatorname{acl}(Ab)$ (such a tuple is called a *basis* of a over A; one of the results of Johnson's paper – [11, Proposition 2.25] – is that this is well-defined, i.e. any two bases have the same length). As usual, one then defines the dimension of an A-definable set X (or an A-type-definable set) as the maximum value of $\dim(a/A)$ for $a \in X$ (and as usual, this does not depend on A, as long as it is small). We caution that the dimension theory nevertheless exhibits wild behaviour in general t-minimal structures, and refer the reader to Johnson's paper for some examples. For now, let us list some crucial properties of the dimension theory that we will use (all of these are in [11, Proposition 2.31 and Theorem 2.37]):

Fact 2.4. *Assume $\mathcal{M}$ is t-minimal, and consider real tuples a and b, parameter sets A and B, and type-definable sets X and Y.*

(1) $\dim(M^n) = n$ *for all* n.
(2) *If* $X \subseteq Y$, *then* $\dim(X) \leq \dim(Y)$.
(3) $\dim(X \times Y) = \dim(X) + \dim(Y)$.
(4) *If* $f : X \to Y$ *is relatively definable and surjective with finite fibres, then* $\dim(X) = \dim(Y)$.
(5) *If* $A \subseteq B$, *then* $\dim(a/B) \leq \dim(a/A)$.
(6) $\dim(ab/A) \leq \dim(b/A) + \dim(a/Ab)$.
(7) *If* $\operatorname{acl}(Aa) = \operatorname{acl}(Ab)$, *then* $\dim(a/A) = \dim(b/A)$.
(8) $\dim(a/A)$ *is the smallest dimension of an A-definable set containing A.*

Fact 2.4(6) is known as *sub-additivity*, and will be referenced as such throughout the paper.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Assumption 2.5. From now until the end of Section 6, assume $\mathcal{M}$ is t-minimal.

2.2. Basic witnesses

Suppose $a \in M^n$ is a real tuple, and A is a parameter set. Let $b \in M^d$ be a basis for a over A, and let $\pi : M^n \to M^d$ be a corresponding projection sending a to b. Since $a \in \mathrm{acl}(Ab)$, there is an A-definable set $X \subseteq M^n$ such that $a \in X$ and π is finite-to-one on X. We give this a separate name:

Definition 2.6. Let $a \in M^n$, A a parameter set and $b \in M^d$ a basis for a over A.

(1) By an a-witness for b over A, we mean an A-definable set $X \subseteq M^n$ with $a \in X$ and admitting a finite-to-one projection $\pi : X \to M^d$ with $\pi(a) = b$.

(2) Let X be an a-witness for b over A, witnessed by the projection $\pi : M^n \to M^d$. We say that X is a *minimal a-witness for b over A* if the fibre size $|\pi^{-1}(b) \cap X|$ is minimal among all such witnesses. That is, let Y be another a-witness for b over A, also witnessed by the projection π. Then $|\pi^{-1}(b) \cap X| \leq |\pi^{-1}(b) \cap Y|$.

(3) A set X is a *basic witness* (resp. *minimal basic witness*) of a over A if there is some basis b of a over A such that X is an a-witness (resp. minimal A-witness) for b over A.

So for any a and A, there is a minimal basic witness of a over A; and the basic witnesses (resp. minimal basic witnesses) form a cofinal family of sets in $\mathrm{tp}(a/A)$.

2.3. Germs

One of our main tools will be the existence of *germs* of types. This was developed in a slightly different setting in [3, §3.2]. In that paper, the existence of germs followed immediately from the frontier inequality (i.e. $\dim(\mathrm{Fr}(X)) < \dim(X)$ for definable X). In the present setting, the frontier inequality can fail (Johnson's paper gives precise examples, including even some with the independent neighbourhood property [Definition 3.1]). However, germs still exist for a different reason – and in many ways, they will serve as a replacement for the frontier inequality for the rest of the paper.

Definition 2.7. Let $a \in M^n$.

(1) If $X \subseteq M^n$ and $a \in X$, then the *germ of X at a*, denoted $\mathrm{germ}_a(X)$, is the equivalence class realized by X among all subsets of M^n modulo agreement in a neighbourhood of a. In this case, denoting $g = \mathrm{germ}_a(X)$, we say that X *realizes* g. In other words, $g = X/\sim_a$ where $Y_1 \sim_a Y_2$ if there exists an open $U \ni a$ such that $Y_1 \cap U = Y_2 \cap U$, and the realizations of g are just members of g.

(2) A *definable germ* at a is a germ at a realized by some definable subseteq of M^n.

(3) If g is a definable germ at a, then the *dimension* of g, denoted $\dim(g)$, is the smallest value of $\dim(X)$ where X is a definable set realizing g.

Remark 2.8. A restatement of Definition 2.7(3) is as follows: Let g be a definable germ at a, realized by the definable set X. Then $\dim(g)$ is the *local dimension of X at a* – that is, the smallest value of $\dim(X \cap U)$ where U is a definable neighbourhood of a.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Lemma 2.9. *Let $a \in M^n$, and A a parameter set. Then the formulas in* $\mathrm{tp}(a/A)$ *have a cofinal germ at a. Precisely, let X be any minimal basic witness of a over A. Then for every A-definable $Y \subseteq X$ containing a, we have* $\mathrm{germ}_a(X) = \mathrm{germ}_a(Y)$.

Proof. The case when a is independent (that is, $\dim(a/A) = n$) is, essentially, due to Johnson. More precisely, in this case, the lemma just says that if Y is A-definable and contains a, then Y contains a neighbourhood of a. But if not, then a belongs to $Y - \mathrm{Int}(Y)$, and thus $\dim(Y - \mathrm{Int}(Y)) = n$. By Theorem 1.7(3) of Johnson's paper [11], this implies that $Y - \mathrm{Int}(Y)$ has non-empty interior, which is absurd.

Now in the general case, let X be any minimal basic witness. So X is a minimal a-witness for some basis b, witnessed by the projection $\pi : M^n \to M^d$ with $\pi(a) = b$.

Let $Y \subseteq X$ be any other A-definable set containing a. Since $Y \subseteq X$, Y is also an a-witness for b over A (also witnessed by π). Since the fibre size above b in X is minimal among all such sets, this means X and Y agree at all points in the fibre above b. Now let $W \subseteq M^d$ be the set of points $u \in \pi(X) \cap \pi(Y)$ such that $\pi^{-1}(u) \cap X = \pi^{-1}(u) \cap Y$. Then W is A-definable and contains b. By the first paragraph (applied to W in M^d, using that $\dim(b/A) = d$), W contains a neighbourhood U of b. In particular, X and Y agree on the preimage of U in M^n, which is a neighbourhood of a. □

Definition 2.10. Let $a \in M^n$ and let A be a parameter set. We define the *germ of a over A*, denoted $\mathrm{germ}(a/A)$, to be $\mathrm{germ}_a(X)$ for any set X as in Lemma 2.9.

Thus, $\mathrm{germ}(a/A)$ is always a definable germ at a (admitting an A-definable realization). Moreover, by Fact 2.1, $\mathrm{germ}(a/A)$ is precisely the germ at a of the set of realizations of $\mathrm{tp}(a/A)$.

2.4. Operations on germs

Germs come with certain natural relations. We review these now.

Definition 2.11. Let g and h be germs at a point $a \in M^n$. We say that g is a *subgerm* of h, denoted $g \leq h$, if for some X realizing g and Y realizing h, we have $X \subseteq Y$. Equivalently, for any X realizing g and Y realizing h, we have $X \cap U \subseteq Y \cap U$ for some neighbourhood U of a.

The following definitions are well-defined since the topology on $X \times Y$ is the product topology.

Definition 2.12. Let g and h be germs at points $a \in M^m$ and $b \in M^n$. The *product of g and h*, denoted $g \times h$, is the germ at $(a,b) \in M^{m+n}$ realized by $X \times Y$ whenever X realizes g and Y realizes h.

Definition 2.13. Let $a \in M^m$ and $b \in M^n$, and let g be a germ at $(a,b) \in M^{m+n}$. The *fibre germ of g over b*, denoted g_b, is the germ at a realized by the fibre $X_b \subseteq X^n$ whenever X realizes g.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

One easily has:

Lemma 2.14. *Let $a \in M^m$ and $b \in M^n$.*

(1) *The relation $\leq$ gives a partial order on germs at a.*
(2) *If g and h are germs at a and b, respectively, then $(g \times h)_b = g$.*
(3) *If g and h are germs at (a,b), and $g \leq h$, then $g_b \leq h_b$.*
(4) *If g and h are definable germs at a and b, respectively, then*

$$\dim(g \times h) = \dim(g) + \dim(h).$$

The following facts are also apparent:

Lemma 2.15. *Let $a \in M^m$, $b \in M^n$ and $A \subseteq B$ parameter sets.*

(1) $\mathrm{germ}(a/B) \leq \mathrm{germ}(a/A)$.
(2) $\mathrm{germ}(ab/A) \leq \mathrm{germ}(a/A) \times \mathrm{germ}(b/A)$.
(3) *If X is A-definable and contains a, then* $\mathrm{germ}(a/A) \leq \mathrm{germ}_a(X)$.

2.5. Codes of germs

We remark here that definable germs can be viewed as elements of $\mathcal{M}^{eq}$. In particular, this will apply to $\mathrm{germ}(a/A)$ for any a and A.

Note that since the topology on $\mathcal{M}$ is $\emptyset$-definable, every automorphism σ of $\mathcal{M}$ is a homeomorphism. It follows that if g is a germ at $a \in M^n$, then $\sigma(g)$ is a well-defined germ at $\sigma(a)$. Now we define:

Definition 2.16. Let $a \in M^n$, and let g be a germ at a. We say that $c \in \mathcal{M}^{eq}$ is a *code* for g if the following holds: Let σ be any automorphism of $\mathcal{M}$. Then $\sigma(c) = c$ if and only if $\sigma(g) = g$ (so in particular $\sigma(a) = a$).

Remark 2.17. An immediate consequence of Definition 2.16 is that if g is a germ of a with code c, then $a \in \mathrm{dcl}(c)$.

Now we show:

Lemma 2.18. *Let $a \in M^n$ and let g be a definable germ at a. Then there is a code for g in $\mathcal{M}^{eq}$. In particular, for any parameter set A, there is a code for* $\mathrm{germ}(a/A)$ *in $\mathcal{M}^{eq}$.*

Proof. Let X be a definable set realizing g. Write $X = Y_b$ for some $\emptyset$-definable family $\{Y_t : t \in T\}$ of sets. Let $\sim$ be the equivalence relation on $M \times T$ given by $(x,s) \sim (y,t)$ if $x = y$ and Y_s and Y_t agree in a neighbourhood of x. Then $\sim$ is $\emptyset$-definable, so $T/\sim$ is a sort in $\mathcal{M}^{eq}$. Now let $c \in T/\sim$ be the equivalence class of (a,b). Then one checks easily that c is a code of g. □

Definition 2.16 should remind the reader of canonical bases in stability theory, particularly when considering germs of the form $\mathrm{germ}(a/A)$. In fact, germs will serve as a replacement for canonical bases when interpreting a group later on. The main difference

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

between the two notions is that germs at a contain the data of the point a, while canonical bases do not. This boils down to which equivalence relation is absorbed in the coding process: one should think of canonical bases (when they exist) as coding formulas in $\mathrm{tp}(a/A)$ up to agreeing *almost everywhere*, while germs code formulas in $\mathrm{tp}(a/A)$ up to agreement *in a neighbourhood of* a.

Note that there can be several codes of $\mathrm{germ}(a/A)$, but any two are interdefinable over a. For this reason – as is often the case with canonical bases – we will abuse notation by identifying all codes of a germ with the germ itself. Thus, if g is a code of $\mathrm{germ}(a/A)$, we will simply write $g = \mathrm{germ}(a/A)$.

2.6. D-approximations

Following [3, Definition 3.16], we also define:

Definition 2.19. Let $a \in M^n$ and let A be a parameter set. A *d-approximation* of $\mathrm{tp}(a/A)$ at a is a definable set X satisfying $a \in X$, $\dim(X) = \dim(a/A)$ and $\mathrm{germ}_a(X) = \mathrm{germ}(a/A)$.

Remark 2.20. Note that a d-approximation of $\mathrm{tp}(a/A)$ at a need not be A-definable.

For example, if $a \in M^n$ and $\dim(a/A) = n$, then a d-approximation of $\mathrm{tp}(a/A)$ at a is just a definable set containing a neighbourhood of a in M^n. In general, one should think of a d-approximation as analogous to a definable set of Morley degree 1 in an ω-stable theory: In that context, any complete type $\mathrm{tp}(a/A)$ is (up to parallelism and passing to $\mathrm{acl}(A)$) the unique generic type of some definable set X of degree 1. In our context, any complete type $\mathrm{tp}(a/A)$ is the unique 'generic at a' type of any d-approximation – where 'generic at a' means that no formula in the type decreases the germ at a.

We now give a couple of easy facts about d-approximations. First, a trivial but useful fact is the following:

Lemma 2.21. *Let $a \in M^n$, let A be a parameter set and let X be a d-approximation of $\mathrm{tp}(a/A)$ at X. Then any A-definable property which holds of a also holds on a neighbourhood of a in X.*

Proof. Let Y be an A-definable set containing a. By Lemma 2.14(3), we compute:

$$\mathrm{germ}_a(Y) \geq \mathrm{germ}(a/A) = \mathrm{germ}_a(X),$$

which is equivalent to the statement of the lemma. □

A restatement of Lemma 2.9 is the following:

Corollary 2.22. *Let $a \in M^n$ and A a parameter set. Then any minimal basic witness of a over A is a d-approximation of $\mathrm{tp}(a/A)$ at a.*

Proof. Let X be a minimal basic witness, corresponding to the projection $\pi : M^n \to M^d$ sending a to the basis $b = \pi(a)$. Lemma 2.9 gives that X realizes $\mathrm{germ}(a/A)$. Now π is finite-to-one on X, so by Fact 2.4, $\dim(X) \leq d$; while since $a \in X$, we also have $\dim(X) \geq d$. So $\dim(X) = d$, and thus X is a d-approximation. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

3. Independent Neighbourhoods

We will need a natural strengthening of t-minimality, which we call the *independent neighbourhood property*. This property has already been studied in the more restrictive context of SW-uniformities and related structures [7], [8, §3.1]. It is also the key geometric idea underlying the definition of *vicinic sets* in [8, §4]. In this section, we introduce the property in the abstract t-minimal setting and study its basic features. The theme will be that independent neighbourhoods allow for stronger geometric properties of definable sets and types.

3.1. Definition and examples

Definition 3.1. $\mathcal{M}$ has the *independent neighbourhood property* if for any tuple $a \in M^n$, any parameter set A and any neighbourhood U of a, there is a neighbourhood V of a such that $V \subseteq U$ and V is definable over a parameter t with $\dim(a/At) = \dim(a/A)$.

Definition 3.1 only requires the new parameters to preserve the dimension of the original tuple a. However, it follows that one can preserve the dimension of any tuple containing a:

Lemma 3.2. *Suppose $\mathcal{M}$ has the independent neighbourhood property. Let $a \in M^m$, $b \in M^n$, let A be a parameter set and let U be a neighbourhood of a. Then there is a neighbourhood $V \subseteq U$ of a defined over a tuple t with* $\dim(ab/At) = \dim(ab/A)$.

Proof. The set $U \times M^n$ is a neighbourhood of (a,b). By the independent neighbourhood property, there is a neighbourhood $W \subseteq U \times M^n$ of (a,b) defined over a parameter t with $\dim(ab/At) = \dim(ab/A)$. Let $\pi : M^m \times M^n \to M^m$ be the projection. Then $\pi(W) \subseteq U$ is a t-definable neighbourhood of a. □

Most natural examples of t-minimal theories have the independent neighbourhood property. For instance, it was shown in [7, Corollary 3.12] that all SW-uniformities have the independent neighbourhood property, covering all dp-minimal expansions of topological groups. More generally, Johnson shows in Appendix B that all visceral theories have the independent neighbourhood property, and thus so do *all* t-minimal topological groups. Finally, in Appendix A, we show that all C-minimal and (dense) weakly o-minimal theories have the independent neighbourhood property.

The simplest example of a t-minimal theory without the independent neighbourhood property is likely one pointed out by Johnson. Namely, consider a real closed field with the Sorgenfrey topology (generated by the half-open intervals of the form $[a,b)$). This gives a t-minimal theory. However, consider a generic singleton a, and any $b > a$. If $V \subset [a,b)$ is any neighbourhood of a, then a is definable over the canonical parameter of V (as the minimum element of V). It follows that no such V can be defined independently of a.

Assumption 3.3. From now until the end of Section 6, assume $\mathcal{M}$ has the independent neighbourhood property.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

3.2. Geometric consequences of independent neighbourhoods

Independent neighbourhoods imply stronger geometric properties of definable sets. We now give various examples. The first lemma below shows that the local dimension of a d-approximation of $\mathrm{tp}(a/A)$ at a is always $\dim(a/A)$. In contrast, without the independent neighbourhood property, one can have infinite discrete definable subsets of M^2 (consider $\{(x,-x) : x \in \mathbb{R}\} \subset \mathbb{R}^2$ with the Sorgenfrey topology; in that case, if $a \in \mathbb{R}$ is generic, then $\mathrm{germ}(a,-a/\emptyset)$ is a single point).

Lemma 3.4. *Let $a \in M^n$, let A be a parameter set and let X be a d-approximation of $\mathrm{tp}(a/A)$ at a.*

(1) *Let $Y \subseteq X$ be any definable set containing a neighbourhood of a in X. Then Y is also a d-approximation of $\mathrm{tp}(a/A)$ at a. Thus, the local dimension of X at a is $\dim(a/A)$.*

(2) *In particular, $\dim(\mathrm{germ}(a/A)) = \dim(a/A)$.*

Proof. We show (1); then (2) follows by definition.

To show (1), first note that Y trivially realizes $\mathrm{germ}(a/A)$, because X does. The point is that we can also show $\dim(Y) = \dim(a/A)$. For this, by assumption, there is a neighbourhood U of a such that X and Y agree on U. By the independent neighbourhood property, we may assume U is definable over At where $\dim(a/At) = \dim(a/A)$. Now $Y \cap U = X \cap U$ is At-definable and contains a, so $\dim(Y \cap U) \geq \dim(a/A)$, and thus $\dim(Y) \geq \dim(a/A)$ as well. Equality then follows since $Y \subseteq X$ and $\dim(X) = \dim(a/A)$. □

Lemma 3.4 implies a convenient reformulation of the independent neighbourhood property in terms of d-approximations:

Lemma 3.5. *Let $a \in M^n$, and let A be a parameter set.*

(1) *Let X be any d-approximation of $\mathrm{tp}(a/A)$ at a. Then there is a d-approximation $Y \subseteq X$ of $\mathrm{tp}(a/A)$ at a, such that Y is definable over At for some tuple t with $\dim(a/At) = \dim(a/A)$.*

(2) *In particular, if U is any neighbourhood of a in M^n, then there is a d-approximation $Y \subseteq U$ of $\mathrm{tp}(a/A)$ at a which is definable over a tuple t with $\dim(a/At) = \dim(a/A)$.*

Proof.

(1) Let Z be an A-definable d-approximation of $\mathrm{tp}(a/A)$ at a (for example, any minimal basic witness of a over A). Then X and Z agree on some neighbourhood of a, say U. By the independent neighbourhood property, we may assume U is At-definable, where $\dim(a/At) = \dim(a/A)$. Now take $Y = Z \cap U$. Then Y is At-definable and realizes $\mathrm{germ}(a/A)$; while we have $Y = Z \cap U = X \cap U \subseteq X$, so Y is a d-approximation by Lemma 3.4.

(2) Let X be any d-approximation of $\mathrm{tp}(a/A)$. Then $X \cap U$ is also a d-approximation of $\mathrm{tp}(a/A)$, by Lemma 3.4. Now apply (1) to $X \cap U$. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Finally, we show that the independent neighbourhood property implies a more natural characterization for d-approximations. Recall that under only t-minimality, a *minimal* basic witness is always a d-approximation. Assuming independent neighbourhoods, the same holds of *all* basic witnesses. This will be very useful when characterizing independence in the next subsection.

Lemma 3.6. *Let $a \in M^n$, A a parameter set and X a basic witness of a over A. Then X is a d-approximation of* $\mathrm{tp}(a/A)$ *at a.*

Proof. As in the proof of Corollary 2.22, we automatically have $\dim(X) = d$. To prove the lemma, we show that $\mathrm{germ}_a(X)$ is cofinal in $\mathrm{tp}(a/A)$. So let $Y \subseteq X$ be A-definable with $a \in Y$. We will show that X and Y agree in a neighbourhood of a.

Let $\pi : M^n \to M^d$ demonstrate that X is a basic witness; so $a \in X$, π is finite-to-one on X and $b = \pi(a)$ is a basis of a over A.

Since π is finite-to-one on X, clearly there is a neighbourhood U of a such that $X \cap U$ and $Y \cap U$ agree on the fibre above b (one chooses U small enough that $X \cap U$ has only the single point a lying above b). By the independent neighbourhood property, we may assume U is At-definable where $\dim(a/At) = \dim(a/A)$. So

$$\begin{aligned}&\dim(a/A) = \dim(a/At) = \dim(ab/At) \leq\\ &\dim(b/At) + \dim(a/Abt) = \dim(b/At) \leq\\ &\dim(a/At) = \dim(a/A),\end{aligned}$$

and it follows that b is still a basis for a over At. Thus, $\dim(b/At) = d$, and so any At-definable property of b holds on a neighbourhood of b. In particular, this implies that $X \cap U$ and $Y \cap U$ agree on all fibres above points in some neighbourhood of b; and as in Lemma 2.9, this implies the desired statement. □

3.3. Independence

Due to the lack of exchange, there are two natural notions of independence in our setting. Namely, given tuples a and b, one could only ask that $\dim(a/b) = \dim(a)$, or one could ask that the same also holds with a and b reversed. In this subsection, we show that the independent neighbourhood property implies a geometric characterization of each of these notions of independence.

Lemma 3.7. *Let $a \in M^n$, and $A \subseteq B$ parameter sets. The following are equivalent:*

(1) $\dim(a/B) = \dim(a/A)$.
(2) $\mathrm{tp}(a/A)$ *and* $\mathrm{tp}(a/B)$ *have a common d-approximation at a.*
(3) $\mathrm{tp}(a/A)$ *and* $\mathrm{tp}(a/B)$ *have exactly the same d-approximations at a.*
(4) $\mathrm{germ}(a/B) = \mathrm{germ}(a/A)$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Proof. (1) ⇒(2): Let X be a basic witness of a over A. If $\dim(a/B) = \dim(a/A)$, then by definition, X is also a basic witness of a over B (This is a subtle use of independent neighbourhoods: if we only knew that minimal basic witnesses were d-approximations, the same argument would fail). By Lemma 3.6, X is a d-approximation of both $\mathrm{tp}(a/A)$ and $\mathrm{tp}(a/B)$ at a.

(2)⇒(3): By definition of d-approximations, (2) gives that $\mathrm{tp}(a/A)$ and $\mathrm{tp}(a/B)$ have the same dimension and the same germ at a, and thus they have the same d-approximations at a.

(3)⇒(4): Clear, since a d-approximation contains the data of the relevant germ.

(4)⇒(1): Assume $\mathrm{germ}(a/B) = \mathrm{germ}(a/A)$. Then by Lemma 3.4, we have

$$\dim(a/B) = \dim(\mathrm{germ}(a/B)) = \dim(\mathrm{germ}(a/A)) = \dim(a/A). \qquad \square$$

Remark 3.8. Lemma 3.7 fails without the independent neighbourhood property, again witnessed by $\mathbb{R}$ with the Sorgenfrey topology and the inverted diagonal. In that case, for generic $a \in \mathbb{R}$, we have $\mathrm{germ}(a, -a/a) = \mathrm{germ}(a, -a/\emptyset)$ (each is a single point), while $0 = \dim(a, -a/a) < \dim(a, -a/\emptyset) = 1$.

Definition 3.9. Let $a \in M^n$ and $A \subseteq B$ parameter sets. We say that a is *independent from B over A* if the equivalent conditions of Lemma 3.7 hold.

We now turn to the obvious symmetric notion of independence.

Lemma 3.10. *Let a and b be tuples, and A a parameter set. The following are equivalent:*

(1) $\dim(a/Ab) = \dim(a/A)$ *and* $\dim(b/Aa) = \dim(b/A)$.

(2) $\dim(ab/A) = \dim(a/A) + \dim(b/A)$.

(3) *There are sets X and Y such that X is a d-approximation of* $\mathrm{tp}(a/A)$ *at a, Y is a d-approximation of* $\mathrm{tp}(b/A)$ *at b and $X \times Y$ is a d-approximation of* $\mathrm{tp}(ab/A)$ *at (a,b).*

(4) *Whenever X is a d-approximation of* $\mathrm{tp}(a/A)$ *at a, and Y is a a-approximation of* $\mathrm{tp}(b/A)$ *at b, we have that $X \times Y$ is a d-approximation of* $\mathrm{tp}(ab/A)$ *at (a,b).*

(5) $\mathrm{germ}(ab/A) = \mathrm{germ}(\mathrm{a}/\mathrm{A}) \times \mathrm{germ}(b/A)$.

Proof. (1)⇒(2): Assume (1). Let c be a basis of a over A, and let d be a basis of b over A. Then $ab \in \mathrm{acl}(Acd)$. If we can show that the concatenated tuple cd is independent over A, it will follow that cd is a basis of ab over A, which implies (2). To show this, assume cd is not independent over A. Then there is some coordinate from cd that is algebraic over A with all other coordinates of cd. Call this coordinate $e \in M$. If e occurs in c, then c is not independent over Ad, which implies that $\dim(a/Ab) < \dim(a/A)$. Similarly, if e occurs in d, then $\dim(b/Aa) < \dim(b/A)$. In either case, we contradict (1).

(2)⇒(3): Let X and Y be basic witnesses for a and b over A, respectively. Assuming (2), we have

$$\dim(X \times Y) = \dim(X) + \dim(Y) = \dim(a/A) + \dim(b/A) = \dim(ab/A),$$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

which implies easily that $X \times Y$ is a basic witness for ab over A (As in Lemma 3.7, this argument fails with minimal basic witnesses, so we are really using independent neighbourhoods). This implies (3) by Lemma 3.6.

(3)⇒(4): Assuming (3), we conclude that $\dim(ab/A) = \dim(a/A) + \dim(b/A)$ and $\mathrm{germ}(ab/A) = \mathrm{germ}(a/A) \times \mathrm{germ}(b/A)$. This implies (4) by definition.

(4)⇒(5): By definition of d-approximations.

(5)⇒(1): Assume $\mathrm{germ}(ab/A) = \mathrm{germ}(a/A) \times \mathrm{germ}(b/A)$. Then by Lemmas 3.4 and 2.14(4), we have

$$\dim(ab/A) = \dim(\mathrm{germ}(ab/A)) = \dim(\mathrm{germ}(a/A) \times \mathrm{germ}(b/A))$$

$$= \dim(\mathrm{germ}(a/A)) + \dim(\mathrm{germ}(b/A)) = \dim(a/A) + \dim(b/A).$$

That is, we have shown that (5) implies (2). But (2) easily implies (1) by sub-additivity. □

Definition 3.11. Let a and b be tuples, and A a parameter set. We say that a and b are *symmetrically independent over* A if the equivalent conditions of Lemma 3.10 are satisfied.

3.4. Generic continuity

Johnson [11, Remark 1.16] points out that generic continuity can fail in a general t-minimal structure. His example is yet again a real closed field with the Sorgenfrey topology and the inverted diagonal (namely $x \mapsto -x$ is nowhere continuous). In this subsection, we show that this is no coincidence: in fact, independent neighbourhoods imply strong generic continuity results. In particular, a corollary will be that generic continuity holds for all definable functions with domain a power of M.

Most of our generic continuity results need some local instance of additivity. Thus, we define:

Definition 3.12. Let a and b be tuples, and A a parameter set. We say that (a,b) is *additive over* A if $\dim(ab/A) = \dim(b/A) + \dim(a/Ab)$.

So if $\mathcal{M}$ satisfies exchange, then all pairs are additive.

Our main use of additivity, which we will exploit frequently, is the following:

Lemma 3.13. *Let (a,b) be additive over A, and let $B \supset A$ with $\dim(ab/B) = \dim(ab/A)$. Then:*

(1) *b is independent from B over A, i.e. $\dim(b/B) = \dim(b/A)$.*

(2) *a is independent from Bb over Ab, i.e. $\dim(a/Bb) = \dim(a/Ab)$.*

(3) *(a,b) is additive over B.*

Proof. Combining additivity over A with sub-additivity over B, we have:

$$\dim(ab/B) \leq \dim(b/B) + \dim(a/Bb) \leq \dim(b/A) + \dim(a/Ab) = \dim(ab/A).$$

Since $\dim(ab/B) = \dim(ab/A)$, all of the above inequalities must be equalities. This implies (1), (2) and (3). □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Now the main technical input in generic continuity is the following:

Lemma 3.14. *Let a and b be tuples so that (a,b) is additive over A. Let π be the projection sending (a,b) to b. If X is any d-approximation of $\mathrm{tp}(ab/A)$ at (a,b), then $\pi(X)$ contains a d-approximation of $\mathrm{tp}(b/A)$ at b.*

Proof. By Lemma 2.21, we may assume X is definable over At where $\dim(ab/At) = \dim(ab/A)$. By Lemma 3.13, we have $\dim(b/At) = \dim(b/A)$. It follows by Lemma 3.7 that $\mathrm{germ}(b/At) = \mathrm{germ}(b/A)$.

On the other hand, note that $\pi(X)$ is an At-definable set containing b. Applying Lemma 2.15(3), we obtain

$$\mathrm{germ}(b/A) = \mathrm{germ}(b/At) \leq \mathrm{germ}_b(\pi(X)).$$

Now let Y be any d-approximation of $\mathrm{tp}(b/A)$ at b. Then $\mathrm{germ}_b(Y) \leq \mathrm{germ}_b(\pi(X))$. It follows that $Y \cap \pi(X)$ contains a neighbourhood of b in Y, so by Lemma 3.4, it is a d-approximation of $\mathrm{tp}(b/A)$ at b. Since $Y \cap \pi(X)$ is contained in $\pi(X)$, the lemma is now proved. □

We conclude that projections are locally open at additive pairs:

Lemma 3.15. *Let (a,b) be additive over A, and let π be the projection sending (a,b) to b. Then there are d-approximations X of $\mathrm{tp}(ab/A)$ at (a,b), and Y of $\mathrm{tp}(b/A)$ at b, so that $\pi(X) = Y$ and $\pi : X \to Y$ is open.*

Proof. Fix A-definable d-approximations X and Y of $\mathrm{tp}(ab/A)$ at (a,b) and $\mathrm{tp}(b/A)$ at b, respectively. Shrinking X if necessary, we may assume $\pi(X) \subseteq Y$ (here we use that $X \cap \pi^{-1}(Y)$ is also A-definable, so is also a d-approximation). Recall that:

(1) Every neighbourhood of (a,b) in X is also a d-approximation of $\mathrm{tp}(ab/A)$ at (a,b), by Lemma 3.4; and

(2) Every d-approximation of $\mathrm{tp}(b/A)$ at b contains a neighbourhood of b in Y, by definition.

In particular, combining (1) and (2) with Lemma 3.14 gives that for every neighbourhood N of (a,b) in X, $\pi(N)$ contains a neighbourhood of b in Y. This is an A-definable property of (a,b) (since it is enough to verify it for the definable neighbourhood basis), so by Lemma 2.21, it holds on a neighbourhood of (a,b) in X; that is, there is a neighbourhood X' of (a,b) in X such that $\pi : X' \to Y$ is open. Let $Y' = \pi(X')$, so that Y' is open in Y and $\pi : X' \to Y'$ is open and surjective. Then by Lemma 3.4, X' and Y' are d-approximations of $\mathrm{tp}(ab/A)$ at (a,b) and $\mathrm{tp}(b/A)$ at b, respectively, and this proves the desired statement. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

In the case of algebraic pairs, we can improve 'open' to 'homeomorphism':

Lemma 3.16. *Let a and b be tuples, and A a parameter set, so that $a \in \operatorname{acl}(Ab)$. Let π be the projection sending (a,b) to b. Then there are d-approximations X of $\operatorname{tp}(ab/A)$ at (a,b), and Y of $\operatorname{tp}(b/A)$ at b, so that π restricts to a homeomorphism $X \to Y$.*

Proof. Since $a \in \operatorname{acl}(Ab)$, we have $\dim(ab/A) = \dim(b/A)$ and $\dim(a/Ab) = 0$; thus (a,b) is additive over A. By Lemma 3.15, we can choose d-approximations X and Y so that $\pi : X \to Y$ is surjective and open. As a projection, π is automatically continuous. The only missing ingredient is ensuring that π is injective on X.

In addition to the above, since $a \in \operatorname{acl}(Ab)$, there is an A-definable set Z containing (a,b) so that π is finite-to-one on Z. By Lemma 2.15(3),

$$\operatorname{germ}_{(a,b)}(Z) \geq \operatorname{germ}(ab/A) = \operatorname{germ}_{(a,b)}(X).$$

Thus, $X \cap Z$ contains a neighbourhood of (a,b) in X, so is also a d-approximation of $\operatorname{tp}(ab/A)$ at (a,b) by Lemma 3.4.

Next, since π is finite-to-one on Z, there is a neighbourhood U of (a,b) so that

$$\pi^{-1}(b) \cap Z \cap U = \{(a,b)\}.$$

By Lemma 3.4(1), $X \cap Z \cap U$ is still a d-approximation of $\operatorname{tp}(ab/A)$ at (a,b). Then by Lemma 3.5(1), there is moreover a d-approximation $X' \subseteq (X \cap Z \cap U)$ which is definable over At for some tuple t with $\dim(ab/At) = \dim(ab/A)$. Let $Y' = \pi(X')$. Then Y' is open in Y, so is still a d-approximation of $\operatorname{tp}(b/A)$ at b by Lemma 3.4, and $\pi : X' \to Y'$ is still open and continuous.

Since $a \in \operatorname{acl}(Ab)$, it follows as in the proof of Lemma 3.6 that also $\dim(b/At) = \dim(b/A)$, so that Y' is also a d-approximation of $\operatorname{tp}(b/At)$ at b. Now let $Y'' \subseteq Y'$ be the set of $y \in Y'$ so that $\pi^{-1}(y) \cap X'$ has only one element. Then Y'' is At-definable and contains b, and thus Y'' contains a neighbourhood of b in Y'. So by Lemma 3.4(1), Y'' is still a d-approximation of $\operatorname{tp}(b/A)$ at b. Finally, let $X'' = X' \cap \pi^{-1}(Y'')$. Then X'' contains a neighbourhood of (a,b) in X', so is still a d-approximation of $\operatorname{tp}(ab/A)$ at (a,b); and now the map $\pi : X'' \to Y''$ is simultaneously open, continuous, surjective and injective, and is thus a homeomorphism. □

We conclude with generic continuity of functions on powers of M. In fact, we will be a bit more general. In his paper, Johnson [11, Definition 2.47] defines a *very weak k-cell* to be a definable set X admitting a finite-to-one projection $X \to M^k$ whose image has non-empty interior. One checks easily that any open subseteq of M^k is a very weak k-cell, and that any very weak k-cell has dimension k. Now we show:

Corollary 3.17. *Let $f : X \to Y$ be an A-definable function, where X is a very weak k-cell. Let $a \in M^n$ with $\dim(a/A) = k$. Then f is continuous in a neighbourhood of a.*

Proof. Let $\pi : X \to M^k$ be a projection whose image has non-empty interior, as provided by the assumption that X is a very weak k-cell. Let $e = \pi(a)$, and let $b = f(a)$. The assumptions imply that $\dim(e/A) = \dim(a/A) = \dim(ab/A) = k$, and thus e is a basis for both a and ab over A. Letting $\Gamma \subseteq X \times Y$ be the graph of f, it now follows that Γ and

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

X are basic witnesses of ab and a over A, respectively, and thus are d-approximations of $\mathrm{tp}(ab/A)$ and $\mathrm{tp}(a/A)$ at (a,b) and a, respectively.

Now by Lemma 3.16, there are d-approximations X' of $\mathrm{tp}(a/A)$ and Γ' of $\mathrm{tp}(ab/A)$ such that the projection $(a,b) \mapsto a$ induces a homeomorphism $\Gamma' \to X'$. Then $\mathrm{germ}_a(X') = \mathrm{germ}_a(X)$, and $\mathrm{germ}_{(a,b)}(\Gamma') = \mathrm{germ}_{(a,b)}(\Gamma)$. So, letting $X'' = X \cap X'$ and $\Gamma'' = \Gamma \cap \Gamma'$, it follows that X'' is a neighbourhood of a in X and Γ'' is a neighbourhood of (a,b) in Γ. Note that if $x \in X''$, then $(x, f(x)) \in \Gamma''$ (and vice versa) – thus Γ'' projects homeomorphically to X''. Finally, on X'', f factors as $X'' \to \Gamma'' \to Y$, a composition of two continuous maps. □

3.5. Uniqueness of the topology

We end this section with an interesting fact that we nevertheless will not need. Recall that a structure can be t-minimal with respect to more than one topology (for example, consider a real closed field with either the order topology or the Sorgenfrey topology). On the other hand, it turns out that there can only be *one* such topology with the independent neighbourhood property (at least up to perturbing a finite set) – and if this topology exists, it is the coarsest t-minimal topology on $\mathcal{M}$ (again up to a finite set). Thus, if one t-minimal topology on a structure has independent neighbourhoods, one might think of that topology as the 'canonical' one.

Recall that we have a fixed topology τ on $\mathcal{M}$ with the independent neighbourhood property. Now we show:

Lemma 3.18. *Let τ' be another $\emptyset$-definable topology on $\mathcal{M}$ with respect to which $\mathcal{M}$ is t-minimal.*

(1) *There is a cofinite set $M' \subseteq M$ so that τ' refines τ on M'.*

(2) *If $\mathcal{M}$ also has the independent neighbourhood property with respect to τ', then there is a cofinite set M' such that τ and τ' agree on M'.*

Proof. We show (1). Then (2) follows by applying (1) twice (the second time with τ and τ' reversed).

Now to prove (1), we will show that for cofinitely many $a \in M$, every definable τ-neighbourhood of a contains a definable τ'-neighbourhood of a. Since this is a definable property of a, by compactness, it will suffice to prove it for any $a \in M - \mathrm{acl}(\emptyset)$. So fix such an a.

Since $a \in M - \mathrm{acl}(\emptyset)$, we have $\dim(a) = 1$. Now let U be a definable τ-neighbourhood of a. It is harmless to shrink U, so by the independent neighbourhood property, we may assume U is definable over a tuple t with $\dim(a/t) = 1$. By t-minimality, the τ'-boundary of U is finite (Remark 2.3), and is thus contained in $\mathrm{acl}(t)$. So a does not belong to the τ'-boundary of U, and thus a belongs to the τ'-interior of U. This shows that U' contains a definable τ'-neighbourhood of a, as desired. □

Remark 3.19. If $\mathcal{M}$ has the independent neighbourhood property with respect to τ, then one can completely describe τ on $M - \mathrm{acl}(\emptyset)$ using only the model theory of $\mathcal{M}$ (with no topology). Let $a \in M - \mathrm{acl}(\emptyset)$. Then by the independent neighbourhood property, a has

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

a neighbourhood basis consisting of all t-definable sets containing a, where t ranges over all tuples with $a \notin \mathrm{acl}(t)$.

4. Infinitesimal neighborhoods

We now introduce the two-model setup that has become customary in treating infinitesimals in topological settings.

Assumption 4.1. From now until the end of Section 6, we fix $\mathcal{N}$, an elementary extension of $\mathcal{M}$ which is sufficiently saturated with respect to $\mathcal{M}$.

By default (unless stated otherwise), we still perform our model-theoretic analysis in $\mathcal{M}$. In particular, parameter sets will continue to be small subsets of $\mathcal{M}^{eq}$.

Following previous works (see, e.g. [17]), we use $\mathcal{N}$ to discuss infinitesimal elements, and in particular infinitesimal neighbourhoods of elements of $\mathcal{M}$. The idea is to prove statements about $\mathcal{M}$ by translating them to cleaner statements about infinitesimals in $\mathcal{N}$.

The present section is divided into four parts. The first subsection makes the framework of infinitesimals precise. There should be no surprises here, as similar frameworks have become well known in topological settings. The next subsection then rewrites some of our main topological lemmas on d-approximations in terms of infinitesimal neighbourhoods. Again, this is a straightforward translation procedure, and we will be brief.

Finally, the next two subsections give two isolated but interesting results. First, we use infinitesimals to prove a very useful remnant of additivity that holds even if $\mathcal{M}$ does not satisfy exchange. Then we show that if $\mathcal{M}$ is saturated in its own cardinality, every infinitesimal neighbourhood is the locus of a complete type.

4.1. Basics

We begin with basic definitions:

Definition 4.2. Let $a \in M^n$ and $b \in N^n$. We say that b is *$\mathcal{M}$-infinitesimally close to a* if b belongs to every $\mathcal{M}$-definable neighbourhood of a. The set of all elements $\mathcal{M}$-infinitesimally close to a is called the *infinitesimal neighbourhood of a*, and denoted $\mu(a)$.

Clearly, in the setup of Definition 4.2, b is infinitesimally close to a if and only if each coordinate of b is infinitesimally close to the corresponding coordinate of a. That is, we have $\mu(a,b) = \mu(a) \times \mu(b)$ for any tuples a and b from $\mathcal{M}$.

We also work with infinitesimal neighbourhoods over parameter sets. In this case, instead of intersecting open sets, we intersect d-approximations:

Definition 4.3. Let $a \in M^n$ and A a parameter set. The *infinitesimal neighbourhood of a over A*, denoted $\mu(a/A)$, is the intersection of all $\mathcal{M}$-definable d-approximations of $\mathrm{tp}(a/A)$ at a, viewed as a type-definable subseteq of N^n.

Remark 4.4. Note that $\mu(a)$ is NOT the same as $\mu(a/\emptyset)$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Thus, in the above setting, $\mu(a/A)$ is type-definable over M. Recall that the dimension theory from $\mathcal{M}$ extends to type-definable sets, and the dimension of a type-definable set is the smallest dimension of any definable set containing it. It follows easily that $\dim(\mu(a/A)) = \dim(a/A)$ for any a and A as above.

We now give two equivalent definitions of $\mu(a/A)$, matching the corresponding notion in other settings (e.g. [17]):

Lemma 4.5. *Let $a \in M^n$, $b \in N^n$ and A a parameter set. The following are equivalent:*

(1) $b \in \mu(a/A)$.

(2) $b \in \mu(a)$ *and b belongs to some $\mathcal{M}$-definable d-approximation of* $\mathrm{tp}(a/A)$ *at a.*

(3) $b \in \mu(a)$ *and* $b \models \mathrm{tp}(a/A)$.

Proof. (1)⇒(2): Assume $b \in \mu(a/A)$. Let X be any $\mathcal{M}$-definable d-approximation of $\mathrm{tp}(a/A)$ at a. By definition, $b \in X$. So to prove (2), it suffices to show that $b \in \mu(a)$. For this, let U be any $\mathcal{M}$-definable neighbourhood of a. By Lemma 3.4, $X \cap U$ is a d-approximation of $\mathrm{tp}(a/A)$ at a, and is also $\mathcal{M}$-definable. Thus, $b \in X \cap U \subseteq U$, as desired.

(2)⇒(3): Assume (2). We need to show that $b \models \mathrm{tp}(a/A)$. For this, let X be A-definable with $a \in X$. By (2), there is an $\mathcal{M}$-definable d-approximation Y of $\mathrm{tp}(a/A)$ at a with $b \in Y$. Then

$$\mathrm{germ}_a(X) \geq \mathrm{germ}(a/A) = \mathrm{germ}_a(Y),$$

so there is a neighbourhood U of a with $Y \cap U \subseteq X \cap U$. As this is an elementary property, and X and Y are $\mathcal{M}$-definable, we can choose U to also be $\mathcal{M}$-definable. By (2), $b \in \mu(a)$, and thus $b \in U$. So $b \in Y \cap U \subseteq X \cap U \subseteq X$, as desired.

(3)⇒(1): Assume (3). Fix an A-definable d-approximation Y of $\mathrm{tp}(a/A)$ at a. To show (1), let X be any $\mathcal{M}$-definable d-approximation of $\mathrm{tp}(a/A)$ at a. Then $\mathrm{germ}_a(X) = \mathrm{germ}_a(Y)$, so there is a neighbourhood U of a so that $X \cap U = Y \cap U$. Since X and Y are $\mathcal{M}$-definable, we can choose U to also be $\mathcal{M}$-definable. By (3), $b \in Y \cap U$, and thus $b \in X \cap U \subseteq X$, as desired. □

We also point out that the term 'neighbourhood' is justified: that is, $\mu(a/A)$ is actually a neighbourhood of a in an appropriate sense:

Lemma 4.6. *Let $a \in M^n$ and A a parameter set.*

(1) *If $X \subseteq M^n$ is an $\mathcal{M}$-definable d-approximation of* $\mathrm{tp}(a/A)$ *at a, then $\mu(a/A)$ is open in X.*

(2) *In particular,* $\mathrm{germ}_a(\mu(a/A)) = \mathrm{germ}(a/A)$.

Proof. It is clear that (1) implies (2), because any d-approximation of $\mathrm{tp}(a/A)$ at a realizes $\mathrm{germ}(a/A)$. We show only (1). So let X be an $\mathcal{M}$-definable d-approximation of $\mathrm{tp}(a/A)$ at a. Then Lemma 4.5(2) gives that $\mu(a/A)$ is the intersection of a small number of open subsets of X – namely those of the form $X \cap U$ where U is an $\mathcal{M}$-definable open neighbourhood of a (here 'small' is in the sense of $\mathcal{N}$, not $\mathcal{M}$). (1) now follows by applying Fact 2.1 in the structure $\mathcal{N}$. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

4.2. Independence and generic continuity revisited

We now rewrite key lemmas from Section 4 using infinitesimal neighbourhoods. We will be terse, as most of the statements we give are straightforward exercises.

We begin with independence:

Lemma 4.7. *Let $a \in M^n$, and $A \subseteq B \subset \mathcal{M}^{eq}$.*

(1) $\mu(a/B) \subset \mu(a/A)$.

(2) $\mu(a/B) = \mu(a/A)$ *if and only if a is independent from B over A.*

Proof.

(1) This follows by Lemma 4.5(3), because every realization of $\mathrm{tp}(a/B)$ also realizes $\mathrm{tp}(a/A)$.

(2) First, suppose a is independent from B over A. Then by Lemma 3.7(3), $\mathrm{tp}(a/A)$ and $\mathrm{tp}(a/B)$ have the same d-approximations at a, which by definition gives that $\mu(a/A) = \mu(a/B)$. Conversely, if $\mu(a/A) = \mu(a/B)$, then

$$\dim(a/A) = \dim(\mu(a/A)) = \dim(\mu(a/B)) = \dim(a/B),$$

so a is independent from B over A. □

Lemma 4.8. *Let $a \in M^m$, $b \in M^n$ and $A \subset \mathcal{M}^{eq}$.*

(1) $\mu(ab/A) \subset \mu(a/A) \times \mu(b/A)$*; thus the natural projections map $\mu(ab/A)$ into $\mu(a/A)$ and $\mu(b/A)$, respectively.*

(2) $\mu(ab/A) = \mu(a/A) \times \mu(b/A)$ *if and only if a and b are symmetrically independent over A.*

Proof.

(1) By Lemma 4.5(3). Namely, suppose $(x,y) \in \mu(ab/A)$. Then $(x,y) \in \mu(a,b)$, so $x \in \mu(a)$ and $y \in \mu(b)$. Moreover, we have $(x,y) \models \mathrm{tp}(ab/A)$, and thus $x \models \mathrm{tp}(a/A)$ and $y \models \mathrm{tp}(b/A)$. Thus, by Lemma 4.5(3), $x \in \mu(a/A)$ and $y \in \mu(b/A)$.

(2) First, suppose a and b are symmetrically independent over A. By (1), we only need to show that if $x \in \mu(a/A)$ and $y \in \mu(b/A)$, then $(x,y) \in \mu(ab/A)$. It is automatic that $(x,y) \in \mu(a,b)$, since $x \in \mu(a)$ and $y \in \mu(b)$. So by Lemma 4.5(2), it suffices to show that (x,y) belongs to some $\mathcal{M}$-definable d-approximation of $\mathrm{tp}(ab/A)$ at (a,b). Now by Lemma 3.10, there are $\mathcal{M}$-definable d-approximations X of $\mathrm{tp}(a/A)$ at a and Y of $\mathrm{tp}(b/A)$ at b, so that $X \times Y$ is a d-approximation of $\mathrm{tp}(ab/A)$ at (a,b). Since $x \in \mu(a/A)$ and $y \in \mu(b/A)$, we have $x \in X$ and $y \in Y$, and thus $(x,y) \in X \times Y$, as desired.

Now assume $\mu(ab/A) = \mu(a/A) \times \mu(b/A)$. Then

$$\dim(\mu(ab/A)) = \dim(\mu(a/A)) + \dim(\mu(b/A)),$$

so $\dim(ab/A) = \dim(a/A) + \dim(b/A)$, and thus a and b are symmetrically independent over A. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Now we turn to generic continuity:

Lemma 4.9. *Let a,b be real tuples, and A a parameter set so that (a,b) is additive over A. Then:*

(1) *The projection $\mu(ab/A) \to \mu(b/A)$ is open and surjective.*
(2) *If $a \in \operatorname{acl}(Ab)$, then the projection $\mu(ab/A) \to \mu(b/A)$ is a homeomorphism.*
(3) *In particular, if a and b are interalgebraic over A, then $\mu(ab/A)$ is the graph of a homeomorphism $\mu(a/A) \to \mu(b/A)$.*

Proof.

(1) Let π be the projection sending (a,b) to b. By Lemma 3.15, there are $\mathcal{M}$-definable d-approximations X of $\operatorname{tp}(ab/A)$ at (a,b), and Y of $\operatorname{tp}(b/A)$ at b, so that π induces an open surjective map $X \to Y$. By Lemma 4.6, $\mu(ab/A)$ and $\mu(b/A)$ are open in X and Y, respectively, which implies that the projection $\mu(ab/A) \to \mu(b/A)$ is also open.

To show $\mu(ab/A) \to \mu(b/A)$ is surjective, let $y \in \mu(b/A)$. By compactness, it suffices to show that for any $\mathcal{M}$-definable d-approximation Z of $\operatorname{tp}(ab/A)$ at (a,b), we have $y \in \pi(Z)$. But by Lemma 3.14, $\pi(Z)$ contains an $\mathcal{M}$-definable d-approximation of $\operatorname{tp}(b/A)$ at b; thus $\pi(Z)$ contains $\mu(b/A)$, and thus $y \in \pi(Z)$.

(2) Repeat the proof of (1), but starting with d-approximations X and Y so that $X \to Y$ is moreover injective (as provided by Lemma 3.16).

(3) Apply (2) to both $\mu(ab/A) \to \mu(a/A)$ and $\mu(ab/A) \to \mu(b/A)$. □

4.3. The fibre germ equation

We now give a rather surprising consequence of independent neighbourhoods. Namely, we prove a fairly strong additivity-like statement that holds even if additivity fails. This will be an extremely useful fact throughout the rest of the paper.

Recall that our main uses of additive pairs will ultimately stem from Lemma 3.13. Roughly, that lemma said that independence is transitive for additive pairs: if (a,b) is additive over A and independent from B over A, then all dimension computations in the additive decomposition of $\dim(ab/A)$ remain unchanged over B.

The result of this subsection, then, is that one of the clauses of Lemma 3.13 (namely Lemma 3.13(2)) does not require additivity at all. This seems to be a unique and pleasant feature of the topological setting (compared to other sub-additive dimension theories).

Lemma 4.10. *Let a and b be real tuples, and A a parameter set.*

(1) *$\mu(ab/A)_b = \mu(a/Ab)$. That is, the fibre above b of the projection $\mu(ab/A) \to \mu(b/A)$ is precisely $\mu(a/Ab)$.*
(2) *$\operatorname{germ}(ab/A)_b = \operatorname{germ}(a/Ab)$ – that is, the fibre germ of $\operatorname{germ}(ab/A)$ above b is $\operatorname{germ}(a/Ab)$.*
(3) *If $B \supset A$ and $\dim(ab/B) = \dim(ab/A)$, then $\dim(a/Bb) = \dim(a/Ab)$.*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Proof. (1) follows from Lemma 4.5(3): we have $x \in \mu(a/Ab)$ if and only if $x \in \mu(a)$ and $x \models \mathrm{tp}(a/Ab)$; and this is equivalent to the assertion that $(x,b) \in \mu(a,b)$ and $(x,b) \models \mathrm{tp}(ab/A)$.

(2) now follows from (1) by applying Lemma 4.6(2) to both sides of the equation. That is, by Lemma 4.6(2), we can rewrite the desired statement as $\mathrm{germ}_a(\mu(ab/A)_b) = \mathrm{germ}_a(\mu(a/Ab))$. But (1) gives that $\mu(ab/A)_b = \mu(a/Ab)$, so of course the two sets have the same germ at a.

Now for (3), assume $B \supset A$ and $\dim(ab/B) = \dim(ab/A)$. By Lemma 3.7, $\mathrm{germ}(ab/B) = \mathrm{germ}(ab/A)$. Now applying (2) gives

$$\mathrm{germ}(a/Bb) = (\mathrm{germ}(ab/B))_b = (\mathrm{germ}(ab/A))_b = \mathrm{germ}(a/Ab).$$

Then by Lemma 3.7 again, $\dim(a/Bb) = \dim(a/Ab)$. □

In fact, Lemma 4.10 implies:

Corollary 4.11. *Let a and b be real tuples, and A a parameter set. If $(x,y) \in \mu(ab/A)$, then the fibre $\mu(ab/A)_y$ realizes* $\mathrm{germ}(x/Ay)$.

Proof. Let X be an A-definable d-approximation of $\mathrm{tp}(ab/A)$ at (a,b). By Lemma 4.10(2), the fibre X_b realizes $\mathrm{germ}(a/Ab)$. Now if $(x,y) \in \mu(ab/A)$, then $\mathrm{tp}(xy/A) = \mathrm{tp}(ab/A)$. So X_y must realize $\mathrm{germ}(x/Ay)$. Since $\mu(ab/A)$ is open in X, the same germ is also realized by $\mu(ab/A)_y$. □

4.4. Infinitesimal neighbourhoods as complete types

We end this section with another interesting fact about infinitesimal neighbourhoods. Namely, in the case that $\mathcal{M}$ is saturated in its own cardinality, Lemma 4.5 (in combination with Lemma 3.5) implies that every infinitesimal neighbourhood is, in fact, the locus of a complete type. We will not explicitly use this fact, as we might not have access to saturated models. Nevertheless, this seems interesting in its own right, so we include a proof sketch. Note that if we had a saturated model, some of our arguments could be expressed a bit more elegantly using this fact.

Lemma 4.12. *Assume $\mathcal{M}$ is saturated. Let $a \in M^n$ and A a parameter set. Then there is a set $B \subseteq M^{eq}$ such that $A \subseteq B$, $\dim(a/B) = \dim(a/A)$ and $\mu(a/A)$ is precisely the set of realizations of* $\mathrm{tp}(a/B)$ *(note that B will be small in the sense of $\mathcal{N}$, but not in the sense of $\mathcal{M}$).*

Proof. Let $\{X_\alpha : \alpha < \kappa\}$ enumerate all $\mathcal{M}$-definable d-approximations of $\mathrm{tp}(a/A)$ at a. We may assume that κ is a cardinal and $\kappa \leq |M|$. We will inductively build b_α and Y_α, for $\alpha < \kappa$, such that:

(1) b_α is a finite tuple from $\mathcal{M}$.

(2) $Y_\alpha \subseteq X_\alpha$ is an Ab_α-definable d-approximation of $\mathrm{tp}(a/A)$ at a.

(3) $\dim(a/A_\alpha b_\alpha) = \dim(a/A_\alpha)$, where $A_\alpha = A \cup \{b_\beta : \beta < \alpha\}$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Indeed, to build b_α and Y_α, one just needs to apply Lemma 3.5 to X_α with parameter set A_α. The key point making this work is that $|A_\alpha| < |M|$ for all α, and hence (as $\mathcal{M}$ is saturated) such a b_α can be found inside $\mathcal{M}$.

Now suppose we have built b_α and Y_α for all α. Let $B = A_\kappa$ be A together with all b_α for $\alpha < \kappa$ – equivalently, the union of all A_α for $\alpha < \kappa$. It follows by induction that $\dim(a/A_\alpha) = \dim(a/A)$ for all $\alpha \leq \kappa$ (for the limit step, use that any drop in dimension is witnessed by a formula with finitely many parameters). In particular, $\dim(a/B) = \dim(a/A)$.

Now we show that $\mu(a/A)$ is the set of realizations of $\mathrm{tp}(a/B)$. First, assume $a' \models \mathrm{tp}(a/B)$. Then for each $\alpha < \kappa$, we have $\mathrm{tp}(a'/Ab_\alpha) = \mathrm{tp}(a/Ab_\alpha)$, and thus $a' \in Y_\alpha \subseteq X_\alpha$. So a' belongs to all $\mathcal{M}$-definable d-approximations of $\mathrm{tp}(a/A)$ at a, and thus $a' \in \mu(a/A)$.

Conversely, suppose $a' \in \mu(a/A)$. If $\mathrm{tp}(a'/B) \neq \mathrm{tp}(a/B)$, then there is $\alpha < \kappa$ so that $\mathrm{tp}(a'/A_\alpha) \neq \mathrm{tp}(a/A_\alpha)$. But A_α is a small set in $\mathcal{M}^{eq}$, and $\dim(a/A_\alpha) = \dim(a/A)$; so we have $\mu(a/A_\alpha) = \mu(a/A)$, by Lemma 4.5 and the definition of Y_α. Thus, $a' \in \mu(a/A_\alpha)$, and thus $\mathrm{tp}(a'/A_\alpha) = \mathrm{tp}(a/A_\alpha)$, a contradiction. □

5. Topological 1-basedness

Throughout this section, we continue to fix a sufficiently saturated t-minimal structure $\mathcal{M}$ with the independent neighbourhood property, as well as a highly saturated elementary extension $\mathcal{N}$ where we treat infinitesimal neighbourhoods of elements of $\mathcal{M}$.

The purpose of this section is to introduce our notion of topological 1-basedness and study its basic properties. The first two subsections give some motivation and equivalent definitions. The next three subsections show that topological 1-basedness is preserved under topological reducts and expansions by constants, and that (similarly to strongly minimal theories) if exchange holds, then topological 1-basedness can be checked at the level of 2-types. Finally, the last three subsections give examples of topologically 1-based and non-1-based structures.

5.1. The definition

We begin by defining topological 1-basedness. In this section, we will regularly use germs as parameters (e.g. if $g = \mathrm{germ}(a/A)$, we may use expressions like $\dim(a/g)$). For this, we remind that by convention, we are identifying germs with their codes in $\mathcal{M}^{eq}$ (Lemma 2.18). In particular, recall that if $g = \mathrm{germ}(a/A)$, then $a \in \mathrm{dcl}(g)$ and $g \in \mathrm{dcl}(Aa)$.

Definition 5.1. Let $a \in M^m$, $b \in M^n$ and A a parameter set. We say that $\mathrm{tp}(a/Ab)$ is *topologically 1-based over* A if $\dim(b/Aa) = \dim(b/A\mathrm{germ}(a/Ab))$. If $A = \emptyset$, we simply say that $\mathrm{tp}(a/b)$ is topologically 1-based.

Definition 5.2. We say that $\mathcal{M}$ is *topologically 1-based* if $\mathrm{tp}(a/b)$ is 1-based (over $\emptyset$) for all real tuples a and b.

Remark 5.3. Letting $g := \mathrm{germ}(a/Ab)$, an equivalent statement in Definition 5.1 would be that $\dim(b/Aag) = \dim(b/Aa)$, i.e. b is independent from g over Aa. This is equivalent because $a \in \mathrm{dcl}(g)$, so $\dim(b/Aag) = \dim(b/Ag)$ is automatic.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Remark 5.4. If $\mathcal{M}$ satisfies exchange, Definition 5.1 makes sense even if b is an imaginary tuple (that is, in that case, the values $\dim(b/Ag)$ and $\dim(b/Aa)$ are well-defined). Without exchange, dimension theory only works for real tuples, and thus we must assume b is a real tuple. Note that if $\mathcal{M}$ satisfies exchange, one could also allow imaginary b in Definition 5.2, and this would result in the same notion of topological 1-basedness (this is essentially Lemma 5.5).

Let us motivate the above definitions. As we stated in the introduction, our intention is to treat germs as the topological analogue of canonical bases. In stability, one would say $\mathrm{tp}(a/b)$ is 1-based if $\mathrm{Cb}(\mathrm{tp}(a/\mathrm{acl}(b))) \in \mathrm{acl}(a)$. In our case, there is no need to pass from b to $\mathrm{acl}(b)$, as Lemma 3.7 gives that $\mathrm{tp}(a/b)$ and $\mathrm{tp}(a/\mathrm{acl}(b))$ have the same germ at a. Thus, the most obvious translation of the definition from stability would simply be that $\mathrm{germ}(a/b) \in \mathrm{acl}(a)$.

Our chosen formulation is a bit more general. Indeed, let $g = \mathrm{germ}(a/b)$. Notice, then, that if $g \in \mathrm{acl}(a)$, then $\dim(b/g) = \dim(b/a)$, which is exactly our definition. On the other hand, the latter property is all we will need, so we have chosen to stick with it. In fact, we do not know whether the two potential definitions are equivalent.

5.2. Equivalent formulations

Recall that if $\mathcal{M}$ satisfies exchange, one can extend the dimension theory to imaginary sorts (see [6, §3]) – with the caveat that there can be infinite 0-dimensional interpretable sets. For example, if $\mathcal{M} \models \mathrm{ACVF}$, then the set of all balls of a given radius is 0-dimensional (the intuition is that the data of a ball is 'locally constant', so the set of balls is topologically 'small'). In this light, the following should help clarify our definition:

Lemma 5.5. *Assume $\mathcal{M}$ satisfies exchange.*

(1) *Let a and b be real tuples, A a parameter set and $g = \mathrm{germ}(a/Ab)$. Then $\mathrm{tp}(a/Ab)$ is topologically 1-based over A if and only if $\dim(g/Aa) = 0$.*

(2) *$\mathcal{M}$ is topologically 1-based if and only if for all real tuples a and parameter sets A, we have $\dim(\mathrm{germ}(a/A)/a) = 0$.*

Proof.

(1) For ease of notation, assume $A = \emptyset$. Since $a \in \mathrm{dcl}(g)$, we have $\dim(b/g) = \dim(b/ga)$. Moreover, since g is a code for the germ of a over b, we get $g \in \mathrm{dcl}(ab)$. Therefore, $\dim(ab) = \dim(abg)$. Now combining with additivity, we compute:

$$\begin{aligned}\dim(a) + \dim(b/a) &= \dim(ab) = \dim(abg)\\ &= \dim(a) + \dim(g/a) + \dim(b/ga)\\ &= \dim(a) + \dim(g/a) + \dim(b/g).\end{aligned}$$

Rearranging now gives

$$\dim(b/a) = \dim(g/a) + \dim(b/g).$$

In particular, $\dim(g/a) = 0$ holds if and only if $\dim(b/a) = \dim(b/g)$, as desired.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

(2) The right-to-left direction is trivial by (1); we show left to right. So assume $\mathcal{M}$ is topologically 1-based. Let a be a real tuple, A be a parameter set and $g = \mathrm{germ}(a/A)$. Let c be a finite tuple from A so that $\dim(a/c) = \dim(a/A)$, and thus $\mathrm{germ}(a/c) = g$. Then let b be a real tuple so that $c \in \mathrm{dcl}(b)$. Arguing by compactness, we may choose b to be independent from a over c; thus $\mathrm{germ}(a/b) = g$ as well. Then since $\mathcal{M}$ is topologically 1-based, and by (1), we get that

$$\dim(\mathrm{germ}(a/b)/a) = \dim(g/a) = 0,$$

which proves (2). □

Returning to the general case, one should think of topologically 1-based types as 'locally linear' – namely, $\mathrm{tp}(a/b)$ is topologically 1-based if in some neighbourhood of b, any conjugate $b' \equiv_a b$ satisfies $\mathrm{germ}(a/b) = \mathrm{germ}(a/b')$. In some sense, this is the most natural way to adapt the stability-theoretic definition using canonical bases to the topological setting. Namely, in a stable theory, the assertion that $\mathrm{Cb}(\mathrm{stp}(a/b)) \in \mathrm{acl}(a)$ really says that $\mathrm{Cb}(\mathrm{stp}(a/b'))$ is constant on a 'big' set of realizations $b' \models \mathrm{tp}(b/a)$ (here 'big' means non-forking over a). In the topological setting, we make exactly the same statement but interpret 'big' as 'on a neighbourhood of b in $\mathrm{tp}(b/a)$' (see Lemma 5.6 below). In general, passing from 'acl^{eq}' to 'constant on a neighbourhood' seems to be the correct philosophy in finding topological analogues of stability-theoretic phenomena.

Lemma 5.6 makes the above discussion precise:

Lemma 5.6. *Let a and b be real tuples, and A a parameter set. The following are equivalent:*

(1) $\mathrm{tp}(a/Ab)$ *is topologically 1-based over A.*

(2) *The map $y \mapsto \mathrm{germ}(a/Ay)$ is constant for $y \in \mu(b/Aa)$.*

Proof. We may assume $A = \emptyset$. Let $g = \mathrm{germ}(a/b)$. First, assume $\mathrm{tp}(a/b)$ is topologically 1-based, and let $b' \in \mu(b/a)$. We show that $g = \mathrm{germ}(a/b')$. Indeed, by 1-basedness, $\dim(b/a) = \dim(b/ga)$, and thus $\mu(b/a) = \mu(b/ga)$ by Lemma 4.7(2). So $b' \in \mu(b/ga)$, and thus $b' \models \mathrm{tp}(b/ga)$ (Lemma 4.5(3)). So $\mathrm{tp}(abg) = \mathrm{tp}(ab'g)$, and since $g = \mathrm{germ}(a/b)$, we conclude that also $g = \mathrm{germ}(a/b')$.

Conversely, suppose $y \mapsto \mathrm{germ}(a/y)$ is constant on $\mu(b/a)$. Then for any $b' \in \mu(b/a)$, we have $\mathrm{tp}(abg) = \mathrm{tp}(ab'g)$: indeed, $\mathrm{tp}(ab) = \mathrm{tp}(ab')$ follows by Lemma 4.5(3), and then we can add g to each side because g is defined over both ab and ab' by the same formula.

Now since $\dim(\mu(b/a)) = \dim(b/a)$, by compactness, there is an element $b' \in \mu(b/a)$ with $\dim(b'/abg) = \dim(b/a) = \dim(b'/a)$. Because $a \in \mathrm{dcl}(g)$, it follows that $\dim(b'/ag) = \dim(b'/g) = \dim(b'/a)$ as well, and since $\mathrm{tp}(abg) = \mathrm{tp}(ab'g)$, this implies $\dim(b/g) = \dim(b/a)$. □

5.3. Expansion by constants

Recall that the class of locally modular strongly minimal theories is preserved under naming constants and taking reducts (as is the analogous class of 'linear' o-minimal theories). In the next two subsections, we prove the same for topological 1-basedness.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Lemma 5.7. *Let C be a parameter set (from $\mathcal{M}^{eq}$), and let $\mathcal{M}_C$ be $\mathcal{M}$ expanded by naming the elements of C. Then $\mathcal{M}_C$ is topologically 1-based if and only if $\mathcal{M}$ is topologically 1-based.*

Proof. Before starting the proof, we should note that $\mathcal{M}_C$ is also t-minimal with the same topology. Moreover, $\mathcal{M}_C$ also has independent neighbourhoods. Indeed, given any a and A, the independent neighbourhood property for $\mathcal{M}$ gives us arbitrarily small neighbourhoods U of a defined over parameters t with $\dim(a/CAt) = \dim(a/CA)$; and this is equivalent to the independent neighbourhood property for $\mathcal{M}_C$. Thus, for the rest of the proof, we can use all of the results collected to this point on independent neighbourhoods.

We now proceed with the proof. First, suppose $\mathcal{M}_C$ is topologically 1-based. To show $\mathcal{M}$ is topologically 1-based, let a and b be real tuples. We wish to show that $\mathrm{tp}(a/b)$ is topologically 1-based. As the definition of topologically 1-based only depends on $\mathrm{tp}(ab)$, we may replace (a,b) with any other realization of the same type. In particular, we may assume that (a,b) is independent from C, i.e. $\dim(ab/C) = \dim(ab)$. By Lemma 4.10, this implies that $\dim(a/Cb) = \dim(a/b)$ and $\dim(b/Ca) = \dim(b/a)$.

Now let $g = \mathrm{germ}(a/b)$. Since $\dim(a/Cb) = \dim(a/b)$, we also have $g = \mathrm{germ}(a/Cb)$. Then the fact that $\mathcal{M}_C$ is topologically 1-based gives $\dim(b/Cg) = \dim(b/Ca)$. But now we compute:

$$\dim(b/g) \geq \dim(b/Cg) = \dim(b/Ca) = \dim(b/a),$$

and this implies that $\mathrm{tp}(a/b)$ is topologically 1-based.

Now assume $\mathcal{M}$ is topologically 1-based; we show that $\mathcal{M}_C$ is topologically 1-based. To do this, fix real tuples a and b, and set $g = \mathrm{germ}(a/Cb)$. We want to show that $\mathrm{tp}(a/Cb)$ is topologically 1-based over C – i.e. that $\dim(b/Cg) = \dim(b/Ca)$.

Let c be a finite tuple from C such that $\dim(a/Cb) = \dim(a/cb)$, and $\dim(b/Cg) = \dim(b/cg)$. Then let d be a real tuple such that $c \in \mathrm{dcl}(d)$. Replacing d with an independent realization of $\mathrm{tp}(d/c)$, we may assume that ab and d are symmetrically independent over c (for example, this is, essentially, [11, Corollary 2.16]).

Let $g = \mathrm{germ}(a/Cb)$. By the choice of c, we also have $g = \mathrm{germ}(a/cb)$. Now by the choice of d, we have $\dim(ab/d) = \dim(ab/c)$, and thus Lemma 4.10(2) gives $\mathrm{germ}(a/bd) = \mathrm{germ}(a/bc) = g$. Since $\mathcal{M}$ is topologically 1-based, it follows that $\dim(bd/g) = \dim(bd/a)$. By Lemma 4.10(3), this implies $\dim(b/dg) = \dim(b/da)$.

Next, recall that $\dim(ab/d) = \dim(ab/c)$. By Lemma 4.10(3) again, this implies $\dim(b/da) = \dim(b/ca)$. Then, combining with the previous paragraph, we have:

$$\begin{aligned}\dim(b/Cg) &= \dim(b/cg) \geq \dim(b/dg)\\ &= \dim(b/da) = \dim(b/ca) \geq \dim(b/Ca),\end{aligned}$$

which implies that $\mathrm{tp}(a/Cb)$ is topologically 1-based over C. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

5.4. Reducts

Now we show that topological 1-basedness is preserved under reducts in an appropriate sense. There is a subtlety here in that one needs to preserve the topology and independent neighbourhoods for this to even make sense. Thus, we define:

Definition 5.8. An *admissible reduct of* $\mathcal{M}$ is a structure $\mathcal{M}^-$ such that:

(1) The universe of $\mathcal{M}^-$ is $\mathcal{M}$.

(2) Every (parametrically) definable set in $\mathcal{M}^-$ is also parametrically definable in $\mathcal{M}$.

(3) With respect to the same topology τ, $\mathcal{M}^-$ is also t-minimal with the independent neighbourhood property (in particular, some basis for τ is $\emptyset$-definable in $\mathcal{M}^-$).

In all known examples of t-minimal structures with independent neighbourhoods, the independent neighbourhoods are witnessed by a fixed definable basis. This means there is a $\emptyset$-definable basis $\{U_t\}$ such that for all $a \in M^n$ and A, there are arbitrarily small $U_{t_1} \times \cdots \times U_{t_n}$ containing a and satisfying $\dim(a/At_1 \dots t_n) = \dim(a/A)$. In this case, any reduct of $\mathcal{M}$ containing the same family $\{U_t\}$ is admissible. In contrast, we do not know whether any reduct defining the same topology is admissible.

Now we show:

Lemma 5.9. *Let $\mathcal{M}^-$ be an admissible reduct of $\mathcal{M}$. If $\mathcal{M}$ is topologically 1-based, then $\mathcal{M}^-$ is topologically 1-based.*

Proof. By Lemma 5.7, we may assume all basic relations of $\mathcal{M}^-$ are definable without parameters in $\mathcal{M}$. Now we begin with:

Claim 5.10. *If $X \subseteq M^n$ is $\mathcal{M}^-$-definable, then* $\dim(X)$ *is the same in $\mathcal{M}^-$ and $\mathcal{M}$.*

Proof. Working first in $\mathcal{M}^-$, let $\{X_i\}$ be a weak cell-decomposition of X (in the sense of [11, Proposition 2.50]). So each X_i is a weak k_i-cell for some k_i. Then X_i is still a weak k_i-cell in the sense of $\mathcal{M}$, and it follows that in both structures $\dim(X)$ is equal to the largest k_i. □

The claim also implies the analogous statement for germs: the dimension of a definable germ in $\mathcal{M}^-$ remains the same when viewed as a definable germ in $\mathcal{M}$ (on the other hand, the data of which germ is $\mathrm{germ}(a/A)$ for fixed a and A may change).

Now towards a proof of the lemma, fix real tuples a and b. We want to show that $\mathrm{tp}_{\mathcal{M}^-}(a/b)$ is topologically 1-based. We can freely replace (a,b) by any other realization of $\mathrm{tp}_{\mathcal{M}^-}(ab)$. In particular, by a compactness argument (using the above claim), one finds a realization such that $\dim_{\mathcal{M}}(ab) = \dim_{\mathcal{M}^-}(ab)$. Now let c be a basis of (a,b) in the sense of $\mathcal{M}^-$; then one still has $(a,b) \in \mathrm{acl}_{\mathcal{M}}(c)$, and it follows that c is still a basis in the sense of $\mathcal{M}$. In particular, if X is an (a,b)-witness for c in the sense of $\mathcal{M}^-$, then X is also an (a,b)-witness for c in the sense of $\mathcal{M}$. It follows that $\mathrm{germ}_{\mathcal{M}^-}(a/b)$ and $\mathrm{germ}_{\mathcal{M}}(a/b)$ have a common realization, and thus are equal. Call this common germ g. By Lemma 4.10(1), $\mathrm{germ}_{\mathcal{M}^-}(b/a)$ and $\mathrm{germ}_{\mathcal{M}}(b/a)$ are both the fibre germ g_a, and thus

$$\dim_{\mathcal{M}^-}(b/a) = \dim(\mathrm{germ}_{\mathcal{M}^-}(b/a)) = \dim(\mathrm{germ}_{\mathcal{M}}(b/a)) = \dim_{\mathcal{M}}(b/a).$$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Now since $\mathcal{M}$ is topologically 1-based, we have $\dim_{\mathcal{M}}(b/g) = \dim_{\mathcal{M}}(b/a)$. If $\mathrm{tp}_{\mathcal{M}^-}(a/b)$ were not topologically 1-based, then we would have

$$\dim_{\mathcal{M}^-}(b/g) < \dim_{\mathcal{M}^-}(b/a) = \dim_{\mathcal{M}}(b/a) = \dim_{\mathcal{M}}(b/g).$$

So there is some g-definable set in $\mathcal{M}^-$ which contains b and has dimension less than $\dim_{\mathcal{M}}(b/g)$. But this assertion is still valid in $\mathcal{M}$, so one obtains $\dim_{\mathcal{M}}(b/g) < \dim_{\mathcal{M}}(b/g)$, a contradiction. □

5.5. Reduction to plane curves

A fundamental result in the theory of strongly minimal sets says that if such a set is locally modular, it is also 1-based; equivalently, non-1-basedness can always be witnessed in the plane. A similar result was proven for weakly 1-based geometric theories in [1, Proposition 2.15]. We do not know whether the analogous statement holds of topological 1-basedness in complete generality – however, it does hold assuming exchange (essentially because of the result of [1]). Here we give a short topological proof.

Lemma 5.11. *Assume $\mathcal{M}$ has exchange and is not topologically 1-based. Then there are $a,b \in M$ and a real tuple t so that $\dim(a/t) = \dim(b/t) = \dim(ab/t) = 1$ and $\mathrm{tp}(ab/t)$ is not topologically 1-based over $\emptyset$.*

Proof. First, note that if $a \in M^n$ and $t \in M^m$ and $\dim(a/t)$ is either 0 or n, then $\dim(a/t)$ is automatically topologically 1-based. Indeed, in this case, $g := \mathrm{germ}(a/t)$ is either a single point or the full germ in M^n at a. In either case, this implies $g \in \mathrm{dcl}(a)$, and thus g and a are interdefinable, and thus automatically $\dim(t/g) = \dim(t/a)$. In particular, this shows that topological 1-basedness is automatic if $n = 1$ (i.e. all 1-types are topologically 1-based).

Now let $a \in M^n$ and $t \in M^m$ so that $\mathrm{tp}(a/t)$ is not topologically 1-based over $\emptyset$, and so that the value of n is minimized among all such pairs. Moreover, write $d := \dim(a/t)$. By the above remarks, we have $n \geq 2$ and $0 < d < n$. Write $a = (e_1, \ldots, e_d, f_1, \ldots, f_{n-d})$ where $e := (e_1, \ldots, e_d)$ is a basis for a over t.

Let us first verify that our choice of n assures that $d = n - 1$. By Lemma 4.9(3), $g = \mathrm{germ}(a/t)$ is realized by the graph of a function $\mu(e) \to \mu(f_1) \times \cdots \times \mu(f_{n-d})$ (i.e. g equals the germ at a of the function's graph). This function is determined by its coordinate components $\mu(e) \to \mu(f_i)$, and each of these components realizes $g_i := \mathrm{germ}(ef_i/t)$. In other words, $g \in \mathrm{dcl}(g_1, \ldots, g_{n-d})$. Now by assumption (and Lemma 5.5), we have $\dim(g/a) > 0$, and thus some $\dim(g_i/a) > 0$. Then $\mathrm{tp}(e_1, \ldots, e_d, f_i/t)$ is not topologically 1-based over $\emptyset$. By the minimality of n, we must have $a = (e_1, \ldots, e_d, f_i)$, so that $d = n - 1$.

We will show that $n = 2$ and thus $d = 1$. So for now, assume towards a contradiction that $n \geq 3$. Let $t' \in \mu(t/a)$ be independent from t over a. A short computation (using that $d = n - 1$ and $n \geq 3$) yields that $a \notin \mathrm{acl}(tt')$. So we can write $a = (b,c)$ where $b \in M$, $c \in M^{n-1}$ and $\dim(b/tt') = 1$.

Let X be a $\emptyset$-definable d-approximation of $\mathrm{tp}(cbt)$ at cbt. By Lemma 4.10, the fibre X_{bt} realizes $\mathrm{germ}(c/bt)$. Since $\mathrm{tp}(cbt) = \mathrm{tp}(cbt')$, $X_{bt'}$ realizes $\mathrm{germ}(c/bt')$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Now by the minimality of n, $\mathrm{tp}(c/bt)$ is topologically 1-based, so by Lemma 5.6, X_{bt} and $X_{bt'}$ agree on some neighbourhood U of c. By the independent neighbourhood property, we may assume U is definable over a tuple u with $\dim(cbtt'/u) = \dim(cbtt')$. By additivity of (c,btt') over $\emptyset$ (i.e. Lemma 3.13), $\dim(btt'/u) = \dim(btt')$, and thus $\dim(b/tt'u) = \dim(b/tt') = 1$. Thus, there is a neighbourhood V of b such that the fibres X_{vt} and $X_{vt'}$ agree on U for all $v \in V$. This means the fibres X_t and X'_t agree on $U \times V$, so that $\mathrm{germ}_a(X_t) = \mathrm{germ}_a(X_{t'})$.

Now similarly to above, since X is a $\emptyset$-definable d-approximation of $\mathrm{tp}(at)$ at at, we have that X_t realizes $\mathrm{germ}(a/t)$, and thus $X_{t'}$ realizes $\mathrm{germ}(a/t')$. It follows that $\mathrm{germ}(a/t) = \mathrm{germ}(a/t') = g$, say. But then $g \in \mathrm{dcl}(at')$, which gives

$$\dim(t/g) \geq \dim(t/at') = \dim(t/a).$$

So $\mathrm{tp}(a/t)$ is topologically 1-based after all, a contradiction.

So we have $n = 2$ and $d = n - 1 = 1$. Write $a = (a_1,a_2)$. If either $\dim(a_i/t) = 0$, then $\mathrm{tp}(a/t)$ is easily seen to be topologically 1-based. Thus, we have $\dim(a_1/t) = \dim(a_1a_2/t) = \dim(a_2/t) = 1$, as desired. □

Remark 5.12. The proof above actually shows a bit more, in the sense that the statement can be localized. Namely, let $X \subseteq M$ be type-definable over A, and suppose exchange holds for any X-tuples over A. Moreover, suppose there are X-tuples a and t with $\mathrm{tp}(a/t)$ not topologically 1-based over A. Then one can find such tuples a and t so that $a \in X^2$ (and so that the other dimension requirements of the lemma are satisfied). The reason this works is that the extra parameters we had to add to t during the proof came from the original tuple a (so we never introduced anything outside X; similarly, $t' \in X$ as well, though this is not needed).

5.6. Locally linear groups

In the remainder of this section, we give examples of topologically 1-based and non-topologically-1-based structures, with the hope of convincing the reader that we have isolated the correct dividing line.

Recall that a 1-based group in stability theory has only group-like definable sets: namely, every definable set is either a subgroup, a coset of a subgroup or a finite Boolean combination of those. In the topological setting, the analogous notion is *local linearity*:

Definition 5.13. Assume $(M, \cdot, \dots)$ is an expansion of a topological group (with the given topology τ). We say $\mathcal{M}$ is *locally linear* if for all $a \in M^m$ and $b \in M^n$, there is a definable subgroup $H \leq M^m$ such that $\mathrm{germ}(a/b) = \mathrm{germ}_a(a \cdot H)$.

The following are all examples of locally linear t-minimal topological groups with independent neighbourhoods:

- Ordered vector spaces over any ordered division ring.
- The reduct $(K, +, v)$ whenever $(K, +, \cdot, v) \models ACVF$. One can make this more general by taking valued vector spaces (in a certain language) over any fixed valued field, but we omit the details as they will not be needed.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

- If K is a model of ACVF or RCVF, and O is the valuation ring, take K/O with its full induced structure (see [7, Corollary 3.12, §6.3]).

Now we show:

Lemma 5.14. *Assume (in addition to our standing assumption of t-minimality and independent neighbourhoods) that $\mathcal{M} = (M, \cdot, \dots)$ is an expansion of a topological group, and is locally linear. Then $\mathcal{M}$ is topologically 1-based.*

Proof. We use Lemma 5.6. Namely, let $a \in M^m$ and $b \in M^n$, and let $b' \in \mu(b/a)$ – so by Lemma 4.10 applied to the projection $\mu(ab/\emptyset) \to \mu(a/\emptyset)$, we see that $(a,b') \in \mu(ab/\emptyset)$. We will show that $\mathrm{germ}(a/b) = \mathrm{germ}(a/b')$. By Lemma 5.6, this will imply that $\mathrm{tp}(a/b)$ is topologically 1-based.

By local linearity, there is an $\mathcal{M}$-definable subgroup $H \leq M^{m+n}$ so that $\mathrm{germ}(ab/\emptyset) = \mathrm{germ}_{(a,b)}((a,b) \cdot H)$. By Lemmas 3.4 and 2.21, there is an $\mathcal{M}$-definable d-approximation X of $\mathrm{tp}(ab)$ at (a,b) such that:

(1) X is an open subseteq of $(a,b) \cdot H$, and

(2) X is t-definable for some t with $\dim(ab/t) = \dim(ab/\emptyset)$.

Note that (2) implies $\mu(ab/t) = \mu(ab/\emptyset)$ – thus $(a,b') \in \mu(ab/t)$, and thus $\mathrm{tp}(abt) = \mathrm{tp}(ab't)$. It follows that X is also a d-approximation of $\mathrm{tp}(ab')$ at (a,b').

Now let X_b and $X_{b'}$ be the fibres in X above b and b' – so $a \in X_b \cap X_{b'}$. Since X is open in $(a,b) \cdot H$, X_b is open in the corresponding fibre $((a,b) \cdot H)_b$ – and likewise for b'. But since H is a subgroup, the two fibres $((a,b) \cdot H)_b$ and $((a,b) \cdot H)_{b'})$ are equal – namely, since they both contain a, they are both equal to $a \cdot N$ where $N \leq M^m$ is the kernel of the projection $H \to M^n$. It follows that $\mathrm{germ}_a(X_b) = \mathrm{germ}_a(X_{b'})$.

Finally, by Lemma 4.7, we now have

$$\begin{aligned}\mathrm{germ}(a/b) &= \mathrm{germ}(ab/\emptyset)_b = (\mathrm{germ}_{(a,b)}X)_b = \mathrm{germ}_a(X_b)\\ &= \mathrm{germ}_a(X_{b'}) = (\mathrm{germ}_{(a,b')}X)_{b'} = \mathrm{germ}(ab'/\emptyset)_{b'} = \mathrm{germ}(a/b'),\end{aligned}$$

which completes the proof. □

In fact, later on we will show that local linearity is equivalent to topological 1-basedness in the case of topological groups (see Theorem 11.10). This is an analogue of the Hrushovski-Pillay classification of 1-based groups in stability theory [10]). We postpone the proof of the full result as it requires more machinery than we have available at this point.

5.7. 1-basedness and fields

On the opposite end of the previous lemma, let us now check that no infinite field is topologically 1-based:

Lemma 5.15. *Assume $\mathcal{M}$ defines an infinite field – that is, assume there is an infinite definable $F \subseteq M^n$ for some n with a definable field structure. Then $\mathcal{M}$ is not topologically 1-based.*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Proof. Suppose $F \subseteq M^n$ is an infinite definable field. By Lemma 5.7, we may assume F (with its operations) is $\emptyset$-definable. Let $d = \dim(F) > 0$. Let $P = L = F^2$, interpreted as the sets of points and (non-vertical) lines in the plane, and let $I \subseteq P \times L$ be the incidence relation. Then one easily checks the following:

(1) $\dim(P) = \dim(L) = 2d$ and $\dim(I) = 3d$.
(2) If $(p,l) \in I$ with $\dim(pl) = 3d$, then $\dim(p) = \dim(l) = 2d$ and $\dim(p/l) = \dim(l/p) = d$.

Now fix $(p,l) \in I$ as above, so $\dim(pl) = 3d$. Let $g = \mathrm{germ}(p/l)$. If $\mathcal{M}$ were topologically 1-based, we would have

$$\dim(l/gp) = \dim(l/p) = d > 0.$$

So there is $l' \neq l$ with $\mathrm{tp}(plg) = \mathrm{tp}(pl'g)$, and thus $g = \mathrm{germ}(p/l')$. It follows that $\mathrm{tp}(p/l)$ and $\mathrm{tp}(p/l')$ have the same d-approximations at p, which gives that $\mu(p/l) = \mu(p/l')$ and $\dim(\mu(p/l)) = d > 0$. So there is $p' \neq p$ with $p' \in \mu(p/l) \cap \mu(p/l')$. Thus, $\mathrm{tp}(pl) = \mathrm{tp}(p'l)$ and $\mathrm{tp}(pl') = \mathrm{tp}(p'l')$. In particular, p and p' are distinct intersection points of the lines l and l', which is a contradiction. □

Remark 5.16. The above lemma fails if we replace 'definable field' with 'interpretable field'. Indeed, for $K \models ACVF$, we already saw that K/O is locally linear, thus topologically 1-based; but it interprets a copy of the residue field k.

5.8. The o-minimal case

We conclude this section by pointing out that topological 1-basedness successfully captures the dividing line of the trichotomy for o-minimal structures.

Recall that every point in an o-minimal structure is either trivial, an element of a group interval with linear structure or an element of a real closed field interval (see the main result of [18]). Call an o-minimal structure *linear* if only the first two cases occur. In [1], Berenstein and Vassiliev generalize linearity in o-minimal structures to a general notion of *weakly 1-based* for geometric theories (including all t-minimal theories with exchange): a geometric theory is weakly 1-based if for all real tuples a and sets B, there is $a' \models \mathrm{tp}(a/B)$ so that $\dim(a/a') = \dim(a/Ba') = \dim(a/B)$.

Now return to our fixed t-minimal $\mathcal{M}$ (and elementary extension $\mathcal{N}$) with independent neighbourhoods. If $\mathcal{M}$ has exchange, then $\mathcal{M}$ is also geometric. As it turns out, in this case $\mathcal{M}$ is topologically 1-based in our sense if and only if it is weakly 1-based as a geometric structure. We will not fully prove this until the next section; for now, we give one implication in general, and the other implication only in the o-minimal case.

Lemma 5.17. *Suppose $\mathcal{M}$ has exchange. If $\mathcal{M}$ is weakly 1-based, then $\mathcal{M}$ is topologically 1-based.*

Proof. We freely use the extension of the dimension theory to imaginaries, given by [6]). Fix real tuples a and b, let $d = \dim(a/b)$, and let $g = \mathrm{germ}(a/b)$. By Lemma 5.5, it suffices to show that $\dim(g/a) = 0$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Since $\mathcal{M}$ is weakly 1-based, there is $a' \models \mathrm{tp}(a/b)$ with $\dim(a/a') = \dim(a/ba') = d$. Let a'' be an independent realization of $\mathrm{tp}(a'/ab)$ – i.e. $\mathrm{tp}(aba') = \mathrm{tp}(aba'')$ and $\dim(a''/aba') = d$. It follows easily that $\dim(a''/a) = \dim(a''/aa') = d$, so that a' and a'' are independent over a.

Now since $\dim(a/ba') = \dim(a/b) = d$, we have $\mathrm{germ}(a/ba') = g$. Then since $\dim(a/a') = \dim(a/ba')$, we also have $\mathrm{germ}(a/a') = g$. Similarly, $\mathrm{germ}(a/a'') = g$. It follows that $g \in \mathrm{dcl}(aa') \cap \mathrm{dcl}(aa'')$; and since a' and a'' are independent over a, it follows from additivity that $\dim(g/a) = 0$, as desired. □

Corollary 5.18. *Assume $\mathcal{M}$ is o-minimal. Then $\mathcal{M}$ is topologically 1-based if and only if it is linear.*

Proof. If $\mathcal{M}$ is topologically 1-based, Lemma 5.15 gives that it does not define an infinite field, and thus is linear. If $\mathcal{M}$ is linear, then by [1, §3.2, Theorem 2.20], it is weakly 1-based, so by Lemma 5.17, it is topologically 1-based. □

6. The structure theorem for topologically 1-based types

We have now finished gathering definitions and tools, and we are ready to prove our main structure theorem for topologically 1-based types (Theorem 6.4). After that theorem, the rest of the paper will be devoted to type-definable groups and methods for constructing them.

6.1. Summary of tools

The statement and proof of Theorem 6.4 will involve many of the tools we have developed to this point. So, for the reader's convenience, we begin with a quick summary of our work so far.

We are working throughout in a highly saturated t-minimal structure $\mathcal{M}$ (Definition 2.2) with the independent neighbourhood property (Definition 3.1), equipped with a fixed elementary extension $\mathcal{N}$, highly saturated with respect to $\mathcal{M}$. We work throughout with tuples $a \in M^n$ and parameter sets $A \subset \mathcal{M}^{eq}$.

Suppose we are given such data, a and A. Then, to this point, we have worked with the following four main associated objects:

(1) The *dimension* $\dim(a/A)$, defined as the length of some (any) *basis* of a over A (taken from [11]).

(2) The *germ* $\mathrm{germ}(a/A)$, defined as the smallest germ of a formula in $\mathrm{tp}(a/A)$ at a (which exists by Lemma 2.9).

(3) The collection of *d-approximations* of $\mathrm{tp}(a/A)$ at a, defined as those definable sets X (not necessarily over A) satisfying $\dim(X) = \dim(a/A)$ and $\mathrm{germ}_a(X) = \mathrm{germ}(a/A)$.

(4) The *infinitesimal neighbourhood* $\mu(a/A)$, defined as the intersection (inside $\mathcal{N}$) of all $\mathcal{M}$-definable d-approximations of $\mathrm{tp}(a/A)$ at a.

These four objects provide four different perspectives on the topological data of $\mathrm{tp}(a/A)$. The dimension $\dim(a/A)$ is just a natural number, most useful for numerical

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

computations. The germ $\mathrm{germ}(a/A)$ is a topological analogue of a canonical base, and has the main advantage that it can be encoded as an imaginary tuple (Lemma 2.18). D-approximations serve as definable sets approximating $\mathrm{tp}(a/A)$ (analogous to sets of Morley degree 1). Meanwhile, the infinitesimal neighbourhood $\mu(a/A)$ is analogous to a global type – it consists of those elements of $\mathcal{N}^n$ which realize $\mathrm{tp}(a/A)$ and are infinitesimally close to a (Lemma 4.5). Note that μ is type-definable over $\mathcal{M}$; moreover, it serves as a 'canonical type-definable d-approximation', in the sense that (as a type-definable set) one has $\dim(\mu(a/A)) = \dim(a/A)$ and $\mathrm{germ}_a(\mu(a/A)) = \mathrm{germ}(a/A)$.

Using the four objects above, we defined *independence*, as well as the stronger *symmetric* independence. Each of these has an obvious definition in terms of each of (1)–(4) above, and the expected equivalences all hold. For example, a is independent from B over $A \subset B$ if and only if the dimension of a (equivalently, the germ, the infinitesimal neighbourhood or the collection of d-approximations) is the same over A and over B. Note that independence and symmetric independence coincide for *additive* pairs (in particular if $\mathcal{M}$ satisfies exchange), but we are not assuming additivity, and thus we will need to carefully differentiate them in what follows.

Next, we developed *generic continuity*. The most important version of generic continuity for us says the following: Suppose a,b are interalgebraic over A. Then $\mu(ab/A)$ is the graph of a homeomorphism between $\mu(a/A)$ and $\mu(b/A)$. This is the 'infinitesimal' way of saying that finite-to-one functions are local homeomorphisms near any generic point.

Finally, in the previous section, we developed topologically 1-based types. Most of the results of that section are for justification of the definition; going forward, we will mainly focus on equivalent formulations of the definition in terms of infinitesimal neighbourhoods and germs. One such formulation is Lemma 5.6, which will be important below. Then the main goal of the current section will be to develop a second such condition (Theorem 6.4).

6.2. Additive extensions

Let us now proceed with the main material of this section. As stated above, the goal is to prove our main topological structure theorem for topologically 1-based types. Roughly speaking, we show that such types look like actual partitions of an ambient space, at least after zooming into an infinitesimal locus. This result and proof are inspired by arguments from Peterzil's group construction in linear o-minimal theories [16].

Before presenting the main result, we give two technical lemmas in this and the following subsections. The first lemma concerns taking independent extensions of types. Roughly, we want to know that we can always find independent realizations of a type without introducing new failures of additivity:

Lemma 6.1. *Let a and b be real tuples, and A a parameter set.*

(1) *There is $c \models \mathrm{tp}(b/a)$ such that $\dim(c/Aab) = \dim(b/Aa)$ and (c,ab) is additive over A.*

(2) *For any c as in (1), and any $B \supset A$ with $\dim(abc/B) = \dim(abc/A)$, we have that b and c are symmetrically independent over Ba.*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

(3) *If* $\mathrm{tp}(a/Ab)$ *is topologically 1-based over* A*, then we can choose* c *as in (1) so that in addition* $\mathrm{germ}(a/Ab) = \mathrm{germ}(a/Ac)$.

Proof.

(1) Let e be a basis for b over Aa. Then e is independent over Aa, so also over A, and thus $\dim(e/A) = \dim(e/Aa) = \dim(b/Aa)$. Note that this implies $\mu(e/Aa) = \mu(e/A)$, which we use below.

Next, note that $\dim(\mu(ae/A)) = \dim(ae/A) = \dim(ab/A)$, and $\dim(\mu(e/A)) = \dim(e/A) = \dim(b/Aa)$. So $X = \mu(ae/A) \times \mu(e/A)$ is a type-definable set of dimension $\dim(ab/A) + \dim(b/Aa)$. By compactness, we can find elements of X of maximal dimension over A. Our main goal is to show that we can take the first coordinate of such an element to be (a,e) itself.

Claim 6.2. *There is* $f' \in \mu(e/A)$ *such that* $\dim(aef'/A) \geq \dim(X) = \dim(ab/A) + \dim(b/Aa)$.

Proof. Let Y be any d-approximation of $\mathrm{tp}(e/A)$ at e. By compactness, it will suffice to find $y \in Y$ so that $\dim(aey/A) \geq \dim(X)$ (noting that this last inequality is a type-definable condition in y).

By Lemma 3.5, we may assume Y is definable over At where $\dim(ae/At) = \dim(ae/A)$, and thus $\mu(ae/At) = \mu(ae/A)$. Now since Y is a d-approximation, we have $\dim(Y) = \dim(e/A) = \dim(\mu(e/A))$. So $\dim(\mu(ae/A) \times Y) = \dim(X)$, and thus there is $(a',e',y') \in \mu(ae/A) \times \mu(e/A)$ with $\dim(a'e'y'/A) \geq \dim(X)$. So $(a',e') \in \mu(ae/A) = \mu(ae/At)$, and thus $\mathrm{tp}(a'e't/A) = \mathrm{tp}(aet/A)$. So there is y with $\mathrm{tp}(a'e'y't/A) = \mathrm{tp}(aeyt/A)$. Then $\mathrm{tp}(y/At) = \mathrm{tp}(y'/At)$, implying $y \in Y$; and moreover $\dim(aey/A) = \dim(a'e'y'/A) \geq \dim(X)$. □

Fix f' satisfying the conclusion of the claim: i.e. $f' \in \mu(e/A)$ and $\dim(aef'/A) \geq \dim(X)$. Since $\mathcal{M}$ is sufficiently saturated, there is f from $\mathcal{M}$ so that $\mathrm{tp}(aef/A) = \mathrm{tp}(aef'/A)$.

Now since $f' \in \mu(e/A) = \mu(e/Aa)$, we have $\mathrm{tp}(ae/A) = \mathrm{tp}(af'/A) = \mathrm{tp}(af/A)$. So there is c from $\mathcal{M}$ so that $\mathrm{tp}(abe/A) = \mathrm{tp}(acf/A)$. Thus, $c \models \mathrm{tp}(b/Aa)$. Moreover, f is also a basis for c over Aa; it follows that bc is interalgebraic with ef over Aa, so that

$$\begin{aligned}\dim(abc/A) = \dim(aef/A) = \dim(aef'/A) &\geq \dim(ab/A) + \dim(b/Aa) \\ &= \dim(ab/A) + \dim(c/Aa) \geq \dim(ab/A) + \dim(c/Aab).\end{aligned}$$

The reverse inequality is automatic by sub-additivity; thus all inequalities above are equalities. We conclude that (c,ab) is additive over A, and that $\dim(c/Aab) = \dim(c/Aa) = \dim(b/Aa)$. Thus, we have now established (1).

(2) Let c be as in (1), and let $B \supset A$ so that $\dim(abc/B) = \dim(abc/A)$. By Lemma 4.10(3), we get

$$\dim(c/Bab) = \dim(c/Aab) = \dim(b/Aa) = \dim(c/Aa) \geq \dim(c/Ba).$$

So c is independent from b over Ba.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

On the other hand, we compute:

$$\begin{aligned}\dim(abc/B) &= \dim(abc/A) = \dim(ab/A) + \dim(c/Aab)\\ &= \dim(ac/A) + \dim(b/Aa) \geq \dim(ac/B) + \dim(b/Bac)\\ &\geq \dim(abc/B).\end{aligned}$$

So all inequalities above must be equalities, and this implies

$$\dim(b/Bac) = \dim(b/Aa) \geq \dim(b/Ba).$$

Thus, b is also independent from c over Ba, which proves (2).

(3) Now assume in addition that $\mathrm{tp}(a/Ab)$ is topologically 1-based over A. We revisit the construction in (1) slightly to prove (3). First, choose e, B and f' as in the original construction. Now let $g = \mathrm{germ}(a/Ab)$, so by assumption $\dim(b/Aa) = \dim(b/Aga)$. It follows that e is also a basis for b over Aga, so that $\dim(e/Aa) = \dim(e/Aga)$, and thus $\mu(e/Aa) = \mu(e/Aga)$. Since $f' \in \mu(e/Aa)$, we conclude that $f' \in \mu(e/Aga)$, and thus $\mathrm{tp}(aeg/A) = \mathrm{tp}(af'g/A)$.

Now we make two quick edits to the construction. Namely, we can first choose f to satisfy the stronger property that $\mathrm{tp}(aefg/A) = \mathrm{tp}(aef'g/A)$, and thus $\mathrm{tp}(afg/A) = \mathrm{tp}(af'g/A) = \mathrm{tp}(aeg/A)$. Then we can in turn choose c to satisfy the stronger property that $\mathrm{tp}(abeg/A) = \mathrm{tp}(acfg/A)$ (so in particular $\mathrm{tp}(abe/A) = \mathrm{tp}(acf/A)$ as in (1)). Then the conclusions of (1) still follow, while we in addition have $\mathrm{tp}(abg/A) = \mathrm{tp}(acg/A)$, and thus $g = \mathrm{germ}(a/Ab) = \mathrm{germ}(a/Ac)$. This proves (2).

□

6.3. Topological characterization of 0-dimensionality

We also need one more lemma before the main argument. Recall that topologically 1-based types require a 'low-complexity germ', in the sense that knowing the germ does not lower the dimension of the parameter (and assuming exchange, this is equivalent to the germ having dimension zero). The following lemma gives a topological characterization of this phenomenon for general imaginaries. This result is likely well known, but we include a proof for completeness.

Lemma 6.3. *Let $a \in M^n$, $b \in \mathcal{M}^{eq}$ and A a parameter set. Assume that $b \in \mathrm{dcl}(Aa)$. Then the following are equivalent:*

(1) $\dim(a/Ab) = \dim(a/A)$.

(2) *There is an A-definable function $f : X \to Y$, where $X \subseteq M^n$ and $Y \subset \mathcal{M}^{eq}$, which takes constant value b on a neighbourhood of a.*

Proof. First, assume (1). Since $b \in \mathrm{dcl}(Aa)$, there is an A-definable function $f : X \to Y$, where $X \subseteq M^n$ and $Y \subset \mathcal{M}^{eq}$, such that $f(a) = b$. Shrinking X if necessary, we may assume it is a d-approximation of $\mathrm{tp}(a/A)$ at a. Now (1) implies that $\mathrm{germ}(a/Ab) = \mathrm{germ}(a/A) = \mathrm{germ}_a(X)$. In particular, the Ab-definable condition '$f(x) = b$' must hold on a neighbourhood of a in X. This proves (2).

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Now assume (2), and let $f : X \to Y$ be the given function. Let U be an open set containing a such that f takes constant value b on $X \cap U$. By the independent neighbourhood property, we may assume U is t-definable where $\dim(a/At) = \dim(a/A)$. But now b is At-definable, since b is the unique element of $f(X \cap U)$. Thus, also $\dim(a/Ab) = \dim(a/A)$, which proves (1). □

6.4. The main structure result

Finally, we give the main result of this section:

Theorem 6.4. *Let a and b be real tuples, and A a parameter set. Then the following are equivalent:*

(1) $\mathrm{tp}(a/Ab)$ *is topologically 1-based over A.*

(2) *Any two fibres in the projection $\mu(ab/A) \to \mu(b/A)$ are either equal or disjoint.*

Proof. First, assume $\mathrm{tp}(a/b)$ is topologically 1-based. To prove (2), we choose tuples $(a',b'),(a'',b'),(a',b'') \in \mu(ab/A)$. We want to show that also $(a'',b'') \in \mu(ab/A)$. To show this, let X be an $\mathcal{M}$-definable d-approximation of $\mathrm{tp}(ab/A)$ at (a,b); we show that $(a'',b'') \in X$. By Lemma 2.21, we may assume X is t-definable where $\dim(ab/At) = \dim(ab/A)$.

Let $g = \mathrm{germ}(a/Ab)$. Since $\dim(ab/At) = \dim(ab/A)$, Lemma 4.10(3) gives that also $g = \mathrm{germ}(a/Abt)$. Now by Lemma 6.1(3), there is $c \models \mathrm{tp}(b/Aat)$ such that $\dim(c/Aabt) = \dim(b/Aat)$, (c,ab) is additive over At and $g = \mathrm{germ}(a/Act)$. Then $\mathrm{tp}(abgt/A) = \mathrm{tp}(acgt/A)$, since g is definable over $Aabt$ and $Aact$ by the same formula.

Since $\dim(ab/At) = \dim(ab/A)$, X is also a d-approximation of $\mathrm{tp}(ab/At)$ at (a,b). Then by Lemma 4.10(2), the fibre X_b realizes $g = \mathrm{germ}(a/Abt)$. Since $\mathrm{tp}(abgt/A) = \mathrm{tp}(acgt/A)$, X_c also realizes g. So X_b and X_c agree on some neighbourhood of a, say U. By Lemma 3.2, we may assume U is u-definable where $\dim(abc/Atu) = \dim(abc/At) = \dim(abc/A)$. Note that by Lemma 6.1(2), we know that b and c are symmetrically independent over $Aatu$; thus $\dim(b/Aactu) = \dim(b/Aatu)$.

Let e be a code of the set $X_b \cap U = X_c \cap U$. So $e \in \mathrm{dcl}(Aabtu) \cap \mathrm{dcl}(Aactu)$. Thus

$$\dim(b/Aatu) \geq \dim(b/Aaetu) \geq \dim(b/Aactu) = \dim(b/Aatu).$$

So equality holds throughout, and by Lemma 6.3, there is an $Aatu$-definable function $f : Y \to Z$ such that $b \in Y$ and f takes constant value e in a neighbourhood of b, say V. By Lemma 3.2, we may assume V is v-definable where $\dim(abc/Atuv) = \dim(abc/Atu) = \dim(abc/At)$. Note, then, that $e \in \mathrm{dcl}(Aatuv)$, since it is the unique image of f on $Y \cap V$.

Now since (c,ab) is additive over At and $\dim(abc/Atuv) = \dim(abc/At)$, it follows by Lemma 3.13 that $\dim(ab/Atuv) = \dim(ab/At) = \dim(ab/A)$. Thus, $\mu(ab/Atuv) = \mu(ab/A)$. In particular, we have

$$(a',b'),(a'',b'),(a',b'') \in \mu(ab/Atuv),$$

and thus each of these three pairs realizes $\mathrm{tp}(ab/Atuv)$. Let e',e'' be such that

$$\mathrm{tp}(abetuv/A) = \mathrm{tp}(a'b'e'tuv/A) = \mathrm{tp}(a'b''e''tuv/A).$$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Since $e \in \mathrm{dcl}(Aatuv)$ (i.e. the b-coordinate is not needed to define e), it follows that $e' = e''$. Thus, e' codes both $X_{b'} \cap U$ and $X_{b''} \cap U$, and thus $X_{b'} \cap U = X_{b''} \cap U$. But since

$$\mathrm{tp}(abtuv/A) = \mathrm{tp}(a''b'tuv/A),$$

and $a \in X_b \cap U$, it follows that $a'' \in X_{b'} \cap U$ (recall here that X and U are Atu-definable). Thus, we also have $a'' \in X_{b''} \cap U$, which shows that $(a'',b'') \in X$, as desired. We have now proved that (1) implies (2).

Now for the converse, assume that any two fibres of $\mu(ab/A) \to \mu(b/A)$ are equal or disjoint. We show that $\mathrm{tp}(a/Ab)$ is topologically 1-based. Since $\dim(\mu(b/Aa)) = \dim(b/Aa)$, a compactness argument gives an element $c \in \mu(b/Aa)$ with $\dim(c/Aab) \geq \dim(b/Aa)$. By Lemma 4.10, $(a,c) \in \mu(ab/A)$. So by assumption, the fibres above b and c in $\mu(ab/A)$ are equal.

Let X be a $\emptyset$-definable d-approximation of $\mathrm{tp}(ab/A)$ at (a,b). By Lemma 4.10, the fibre X_b realizes $\mathrm{germ}(a/Ab)$. Since $(a,c) \in \mu(ab/A)$, we have $\mathrm{tp}(ab/A) = \mathrm{tp}(ac/A)$, and thus the fibre X_c realizes $\mathrm{germ}(a/Ac)$. But by Lemma 4.6, $\mu(ab/A)$ is open in X; thus the common fibre above b and c in $\mu(ab/A)$ realizes both $\mathrm{germ}_a(X_b)$ and $\mathrm{germ}_a(X_c)$. In particular, $\mathrm{germ}(a/Ab) = \mathrm{germ}(a/Ac) = g$, say. Then $\mathrm{tp}(abg/A) = \mathrm{tp}(acg/A)$, since g is definable over Aab and Aac by the same formula. Now since $g \in \mathrm{dcl}(Aab)$, we have

$$\dim(b/Aag) = \dim(c/Aag) \geq \dim(c/Aab) \geq \dim(b/Aa).$$

This implies that $\mathrm{tp}(a/Ab)$ is topologically 1-based. □

We conclude, as promised earlier, that topological 1-basedness restricts exactly to the prior notion of weak 1-basedness for geometric theories:

Corollary 6.5. *Assume $\mathcal{M}$ satisfies exchange. Then $\mathcal{M}$ is topologically 1-based if and only if it is weakly 1-based.*

Proof. We showed in Lemma 5.17 that weak 1-basedness implies topological 1-basedness. We now show the reverse implication. So assume $\mathcal{M}$ is topologically 1-based.

Recall that to show $\mathcal{M}$ is weakly 1-based, we are given a real tuple a and a small set $B \subseteq M$. We need to find $a' \models \mathrm{tp}(a/B)$ so that $\dim(a/a') = \dim(a/Ba') = \dim(a/B)$. By compactness (since $\mathcal{M}$ is sufficiently saturated), it is enough to prove the result in the case where $B := b$ is a finite tuple.

First, since $\dim(\mu(a/b)) = \dim(a/b)$, there is $c \in \mu(a/b)$ with $\dim(c/ab) \geq \dim(a/b)$. In particular, $\mathrm{tp}(c/b) = \mathrm{tp}(a/b)$, and $\dim(c/ab) = \dim(a/b)$.

Similarly, there is $d \in \mu(b/a)$ with $\dim(d/abc) = \dim(b/a)$. Then the fibres above b and d in $\mu(ab/\emptyset) \to \mu(b/\emptyset)$ both contain a, and thus they are equal by Theorem 6.4. In particular, since $(c,b) \in \mu(ab/\emptyset)$, it follows that $(c,d) \in \mu(ab/\emptyset)$ as well. Thus, $\mathrm{tp}(cd) = \mathrm{tp}(ab)$, and so $\dim(d/c) = \dim(b/a) = \dim(d/abc)$. It follows that d and a are independent over c, so that $\dim(a/c) = \dim(a/cd)$.

Now, since $(a,d) \in \mu(ab/\emptyset)$, we also have $\mathrm{tp}(ad) = \mathrm{tp}(ab)$, and thus

$$\dim(d/a) = \dim(b/a) = \dim(d/abc).$$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

It follows that also $\dim(d/ab) = \dim(abc)$, so d and c are independent over ab. Thus also

$$\dim(c/abd) = \dim(c/ab) = \dim(a/b) = \dim(c/d).$$

It follows that c and a are independent over d. Thus, we also have $\dim(a/cd) = \dim(a/d)$.

To recap, we have now shown that $\dim(a/c) = \dim(a/cd) = \dim(a/d)$. Now, since $(a,d) \in \mu(ab/\emptyset)$, we have $\mathrm{tp}(ab) = \mathrm{tp}(ad)$, and thus there is a' so that $\mathrm{tp}(aba') = \mathrm{tp}(adc)$. Since $\mathcal{M}$ is sufficiently saturated, we may assume a' is a tuple from $\mathcal{M}$. Then $\mathrm{tp}(a'b) = \mathrm{tp}(cd) = \mathrm{tp}(ab)$, so that $a' \models \mathrm{tp}(a/b)$ – and moreover, since $\dim(a/c) = \dim(a/cd) = \dim(a/d)$, it follows that $\dim(a/a') = \dim(a/a'b) = \dim(a/b)$. This shows that $\mathcal{M}$ is weakly 1-based. □

7. Groupoid spines

Later on, we will show that under a suitable non-triviality assumption, any topologically 1-based t-minimal structure with the independent neighbourhood property admits an infinite type-definable group. This amounts to a topological version of the group configuration theorem. Before giving the argument, we first separate its purely combinatorial content. Thus, the current section is independent from model theory and concerns only the recovery of a group (in fact, a groupoid) from a well-behaved family of bijections. The general idea is similar to [5], but with details more suited to our needs. Then in the next section, we add type-definability requirements (but working in an arbitrary theory), before finally returning to our topologically 1-based setting in Sections 9–11.

Recall that a *groupoid* is a category in which every morphism is invertible. Concretely, one thinks of a collection of sets, $\mathcal{X} = \{X_i : i \in I\}$ – in addition to a map $(i,j) \mapsto \mathrm{Mor}(i,j)$ assigning each $i,j \in I$ to a collection of bijections $X_i \to X_j$ (called morphisms) – such that the morphisms contain all identity maps and are closed under composition and inverse in the obvious senses. If all of the morphism sets are non-empty, one further calls the groupoid *connected*.

The main object of study in the next two sections is what we call a *groupoid spine*. This is similar to a connected groupoid in the concrete sense above, except not all pairs (i,j) are assigned a morphism set (and the identity, composition and inverse properties are only required when all of the relevant morphism sets exist – see Definition 7.1). A central theme of our study, then, is the *extension problem* – asking when a groupoid spine admits a unique groupoid extension (this means it can have all of its missing morphism sets added in to make it into a groupoid). The point is that this extension problem is entirely analogous to the group configuration construction in model theory. Note that in order for such a construction to work, we need 'enough' morphism sets to begin with. This is the role played by the linear order $<$ in Definition 7.1 below.

7.1. Basic notions

Definition 7.1. A *groupoid spine* consists of:

- A linear order $(I, <)$,
- An I-indexed family of non-empty sets, $\mathcal{X} = \{X_i : i \in I\}$,
- A non-empty relation $\emptyset \neq R \subseteq I^2$ containing all pairs (i,j) with $i < j$, and

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

- For all $(i,j) \in R$, a non-empty collection $\mathrm{Mor}(i,j)$ of bijections $X_i \to X_j$ (called *morphisms*),

such that the morphisms are relatively closed under the groupoid operations. That is:

(1) If $(i,i) \in R$ for some i, then $\mathrm{id} : X_i \to X_i \in \mathrm{Mor}(i,i)$.
(2) If $(i,j),(j,i) \in R$ and $f \in \mathrm{Mor}(i,j)$, then $f^{-1} \in \mathrm{Mor}(j,i)$.
(3) If $(i,j)(j,k),(i,k) \in R$, $f \in \mathrm{Mor}(i,j)$ and $g \in \mathrm{Mor}(j,k)$, then $g \circ f \in \mathrm{Mor}(i,k)$.

We denote a groupoid spine by the data $(I, <, \mathcal{X}, R, \mathrm{Mor})$, where Mor is the family $\{\mathrm{Mor}(i,j) : (i,j) \in R\}$.

So if $R = I^2$ above, we just obtain a groupoid (on a linearly ordered set of objects). The point is to consider groupoid spines whose set R is smaller.

Note that Definition 7.1 implies that for any $i,j \in I$, there is a composition of morphisms and their inverses sending X_i to X_j (which may be empty if $i = j$). This is why groupoid spines are analogous to *connected* groupoids (indeed, this condition ensures that the groupoid determined by a groupoid spine, if it exists, is always connected).

The next definition and lemma clarify the 'extension problem' mentioned above.

Definition 7.2. Let $(I, <, \mathcal{X}, R, \mathrm{Mor})$ be a groupoid spine.

(1) Suppose $R \subseteq S \subseteq I^2$. We say that $(I, <, \mathcal{X}, R, \mathrm{Mor})$ *extends to* S if one can define $\mathrm{Mor}(i,j)$ for $(i,j) \in S - R$ in such a way that $(I, <, \mathcal{X}, S, \mathrm{Mor})$ is a groupoid spine.
(2) We say that $(I, <, \mathcal{X}, R, \mathrm{Mor})$ *extends to a groupoid* if it extends to $S = I^2$.

In other words, $(I, <, \mathcal{X}, R, \mathrm{Mor})$ extends to S if it can be grown to a groupoid spine on S without adding any new morphisms to pairs in R.

The following is then clear.

Lemma 7.3. *Suppose $(I, \mathcal{X}, R, \mathrm{Mor})$ is a groupoid spine that extends to a groupoid. Then the extension to a groupoid is uniquely determined. More precisely, the resulting groupoid is precisely the groupoid on I generated by all morphisms between pairs in R.*

7.2. Combinatorial group configuration theorem

We now give a purely combinatorial analogue of the group configuration theorem. The statement will say that a symmetric groupoid spine on at least three objects always extends to a groupoid. First, we define:

Definition 7.4. A groupoid spine $(I, <, \mathcal{X}, R, Mor)$ is *symmetric* if R is symmetric – that is, $(i,j) \in R$ implies $(j,i) \in R$ for all $i,j \in I$.

Equivalently, one easily checks that $(I, <, \mathcal{X}, R, Mor)$ is symmetric if and only if it contains all pairs (i,j) with $i \neq j$. We now show:

Theorem 7.5. *Let $(I, < \mathcal{X}, R, \mathrm{Mor})$ be a symmetric groupoid spine with $|I| \geq 3$. Then $(I, <, \mathcal{X}, R, \mathrm{Mor})$ extends to a groupoid.*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Proof. Let $\mathcal{G}$ be the groupoid on I generated by all maps in the sets $\mathrm{Mor}(i,j)$ for $(i,j) \in R$. It suffices to show that for $(i,j) \in R$, every $\mathcal{G}$-morphism $i \to j$ already belongs to $\mathrm{Mor}(i,j)$.

So let $(i,j) \in R$, and let $f : X_i \to X_j$ be any $\mathcal{G}$-morphism. So f is a composition of maps from Mor and their inverses. On the other hand, by symmetry, inverses are not needed. So in fact, f is directly a composition of maps from Mor.

To be precise, we now obtain $i = i_0, \ldots, i_n = j \in I$ such that each $(i_l, i_{l+1}) \in R$ and f factors as a composition

$$X_{i_0} \xrightarrow{f_1} \ldots \xrightarrow{f_n} X_{i_n},$$

where each $f_{i_l} \in \mathrm{Mor}(i_{l-1}, i_l)$. We may assume that n is minimal among all such decompositions of f. Our goal is to show that $n = 0$ or $n = 1$ (where we interpret $n = 0$ as $i = j$ and $f = \mathrm{id} : X_i \to X_i$). So, assume $n \geq 2$.

Note that if $(i_l, i_{l+2}) \in R$ for any l, then by closure under composition, the chain

$$X_{i_l} \xrightarrow{f_l} X_{i_{l+1}} \xrightarrow{f_{l+1}} X_{i_{l+2}}$$

can be shortened to a single morphism $X_{i_l} \to X_{i_{l+2}}$, contradicting the minimality of n. So each $(i_l, i_{l+2}) \notin R$. Since R contains all distinct pairs, this implies $i_{l+2} = i_l$ whenever $0 \leq l \leq n-2$. So the sequence $i_0, \ldots, i_n$ alternates between the two values $i_0 = i$ and i_1, and $(i,i) \notin R$. Since $(i,j) \in R$ by assumption, this implies $i \neq j$. This is only possible if $j = i_1$. That is, we deduce that i and j are distinct and the sequence $i_0, \ldots, i_n$ alternates between i and j. Since $i_n = j$, it follows that n is odd, so in particular $n \geq 3$. So our decomposition starts with the chain

$$X_i \xrightarrow{f_1} X_j \xrightarrow{f_2} X_i \xrightarrow{f_3} X_j.$$

We now show that this chain of three maps can be shortened to a single morphism $X_i \to X_j$, which will contradict the minimality of n and thus prove the theorem. To do this, recall the assumption that $|I| \geq 3$. Thus (as $i \neq j$), there is $k \in I$ so that i,j,k are all distinct. Thus, all non-diagonal pairs from $\{i,j,k\}$ belong to R.

Now fix any morphism $g : j \to k$. By closure under inverses and composition, there is a morphism $h : k \to i$ so that $h \circ g = f_2$. So the first three maps in our decomposition can be expanded to the chain

$$X_i \xrightarrow{f_1} X_j \xrightarrow{g} X_k \xrightarrow{h} X_i \xrightarrow{f_3} X_j.$$

By closure under composition applied to each of $X_i \xrightarrow{f_1} X_j \xrightarrow{g} X_k$ and $X_k \xrightarrow{h} X_i \xrightarrow{f_3} X_j$, this chain simplifies to a composition of two morphisms $X_i \to X_k \to X_j$ – which then simplifies for the same reason to a single morphism $X_i \to X_j$, as desired. □

7.3. Regularity

Typically, one should not expect a groupoid spine to extend to a groupoid (at least, not without having many pairs $(i,j) \in R$ with $i > j$) – and this is related to group configurations. For example, let $I = \{1,2,3\}$ with the usual order $<$, fix a non-empty set X and fix any non-empty family F of bijections $X \to X$. Now define:

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

- $X_i = X$ for each i.
- $\mathrm{Mor}(1,2) = \mathrm{Mor}(2,3) = F$.
- $\mathrm{Mor}(1,3) = F \circ F$, the set of all compositions of two maps from F.

This data defines a groupoid spine (where $R = \{(1,2),(2,3),(1,3)\}$). One can check that if this groupoid spine extends to a groupoid, then F is a (left and right) coset of a group of permutations of X. In particular, given $f,g \in F$, the composition $f \circ g$ is 'independent' from both f and g (meaning that for any f', there is g' with $f \circ g = f' \circ g'$, and the analogous statement with f and g reversed). This independence of a composition with its factors is also the main idea of group configurations. So one can think of a groupoid spine with this property as an analogue of a group configuration.

Typically, the independence property above is not expected. However, in stable 1-based theories, group configurations arise 'automatically': for example, in a locally modular strongly minimal theory, any composition of one-dimensional families of plane curves gives rise to a group configuration. This is roughly because the composite family cannot be too large, in the sense that a composition $f \circ g$ is generically determined by one point. In the abstract language of groupoid spines, the analogous property is *regularity*:

Definition 7.6. Let $(I, <, \mathcal{X}, R, Mor)$ be a groupoid spine. We say that $(I, <, \mathcal{X}, R, Mor)$ is *regular* if for all $(i,j) \in R$, and all $x \in X_i$ and $y \in X_j$, there is exactly one morphism $f \in \mathrm{Mor}(i,j)$ with $f(x) = y$.

Note the analogy of the above definition with regular group actions. In particular, a group of permutations of a set induces a regular groupoid spine in the obvious way if and only if the corresponding group action is regular.

7.4. The main result on regular groupoid spines

We now show an analogue of the fact that group configurations are automatic in the 1-based case. Namely, we show that for groupoid spines, regularity guarantees extendability.

Theorem 7.7. *Let $(I, <, \mathcal{X}, R, Mor)$ be a regular groupoid spine. Then:*

(1) *There is a symmetric set $S \supset R$ such that $(I, \mathcal{X}, R, Mor)$ extends to S.*

(2) *In particular, if $|I| \geq 3$, then $(I, \mathcal{X}, R, Mor)$ extends to a groupoid.*

Proof. That (1) implies (2) is immediate by Theorem 7.5. We only show (1).

Let S be the symmetric closure of R – so S adds to R all pairs $(i,j) \notin R$ with $i > j$, and in particular S is symmetric. We will use multiple times the fact that S only adds strictly decreasing pairs to R – so if $(i,j) \in S$ and $i \leq j$, then $(i,j) \in R$.

For any pair $(i,j) \in S - R$, note that $(j,i) \in R$, so $\mathrm{Mor}(j,i)$ is well-defined. Now set $\mathrm{Mor}(i,j)$ to be the collection of all inverses of maps in $\mathrm{Mor}(j,i)$. All of the required properties in Definition 7.1 follow immediately, except closure under composition, which will occupy the rest of the proof.

So, let $i,j,k \in I$ with $(i,j),(j,k),(i,k) \in S$, and consider a composition of morphisms $i \to j \to k$. We want to show that the composite map $X_i \to X_k$ is a morphism $i \to k$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

To do this, our first main goal is to convert everything into a problem in the original groupoid spine (by replacing certain maps with their inverses to reduce to a statement about only maps between R-pairs). We then apply regularity in the original spine to prove the resulting statement.

Let us proceed. First, note that if $i \leq j \leq k$, then $(i,j),(j,k),(i,k) \in R$, and the desired statement follows from closure under composition for R-pairs. The case $i \geq j \geq k$ then follows similarly by replacing each map with its inverse.

Otherwise, j is either the maximum or minimum of $\{i,j,k\}$. We will prove the case that $j = \max\{i,j,k\}$, as the other case is similar. Assuming $j = \max\{i,j,k\}$, the problem is now symmetric in i and k – so without loss of generality, we assume $i \leq k$, and thus $(i,k) \in R$ as well.

By replacing the given morphism $j \to k$ with its inverse, we can now restate the problem in terms of R-pairs: we know that $(i,j),(k,j),(i,k) \in R$, and we have morphisms $f : i \to j$ and $g : k \to j$. We want to find a morphism $h : i \to k$ with $g \circ h = f$.

We now find such a morphism h using regularity. First, pick any element $y \in X_k$. Let $z = g(y) \in X_j$, and let $x \in X_i$ with $f(x) = z$. By regularity, there is $h \in \mathrm{Mor}(i,k)$ with $h(x) = y$. Thus, $g \circ h$ is a morphism $i \to j$ sending x to z. By regularity, there is only one such morphism. So since $f(x) = z$ as well, it follows that $g \circ h = f$, as desired. □

Before moving on, we also note that regularity is preserved under groupoid extensions:

Lemma 7.8. *Let $(I,\mathcal{X},R,Mor)$ be a regular groupoid spine. If $(I,\mathcal{X},R,Mor)$ extends to a groupoid, then the resulting extension is also regular.*

Proof. Let $(i,j) \in I^2$, and let $x \in X_i$ and $y \in X_j$. We show that there is exactly one morphism $i \to j$ (in the groupoid extension) sending x to y.

If $i < j$, then $(i,j) \in R$, and the desired statement follows from regularity for R-pairs. The case $i > j$ then follows similarly, by taking inverses and finding a unique morphism $j \to i$ sending y to x. Combining these two cases, the desired statement holds whenever $i \neq j$.

Now assume $i = j$, and from now on write (i,j) as (i,i). If $(i,i) \in R$, then we are done (by regularity for R-pairs), so we may further assume $(i,i) \notin R$. Recall that $R \neq \emptyset$ (by the definition of groupoid spines); so there must be at least one other element of I, say k. Fix any morphism $g : X_i \to X_k$, and let $z = g(y) \in X_k$. By the groupoid axioms, the map $f \mapsto g \circ f$ gives a bijective correspondence between $\mathrm{Mor}(i,i)$ and $\mathrm{Mor}(i,k)$. Moreover, for $f : i \to i$, the condition $f(x) = y$ is equivalent to $g \circ f(x) = z$. Since $i \neq k$ (and this case was covered above), there is a unique morphism $i \to k$ sending x to z. Passing back through our bijective correspondence, this implies there is a unique morphism $i \to i$ sending x to y, as desired. □

8. Type-definability requirements

We now revisit the constructions from the previous section in the context of a first-order theory, adding type-definability hypotheses in each step. The results of this section are valid in any first-order theory. In the following section, we then show that they apply

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

in the setting of topological 1-basedness considered earlier. In this section, we work in a highly saturated first-order structure $\mathcal{M}$ (with no topology assumptions).

8.1. Type-definability and faithful type-definability

Definition 8.1. Let A be a parameter set, and let $(I, <, \mathcal{X}, R, Mor)$ be a groupoid spine. We say that $(I, <, \mathcal{X}, R, Mor)$ is *A-type-definable* if I is small, each X_i is A-type-definable and the morphisms can be represented by A-type-definable families. More precisely, this means that for each $(i,j) \in R$, one can choose A-type-definable sets T_{ij} and $F_{ij} \subseteq X_i \times X_j \times T_{ij}$, so that:

(1) Each fibre $(F_{ij})_t \subseteq X_i \times X_j$ for $t \in T_{ij}$ is the graph of a (necessarily unique) morphism $i \to j$, which we denote $f_t^{ij} : X_i \to X_j$.
(2) For each morphism $f : i \to j$, there is $t \in T_{ij}$ with $f_t^{ij} = f$.

In this case, we call $\{T_{ij}\}$ and $\{F_{ij}\}$ a *type-definition of* $(I, <, \mathcal{X}, R, Mor)$ *over* A.

Remark 8.2. Assume that $(I, <, \mathcal{X}, R, Mor)$ is an A-type-definable groupoid spine, with a type-definition $(\{T_{ij}\}, \{F_{ij}\})$ over A. Then by compactness, each F_{ij} is in fact relatively A-definable in $X_i \times X_j \times T_{ij}$ (this is a general property of type-definable families of functions).

Definition 8.3. Let $(I, <, \mathcal{X}, R, Mor)$ be a groupoid spine, and let A be a parameter set. We say $(I, <, \mathcal{X}, R, Mor)$ is *faithfully A-type-definable* if it is A-type-definable and in addition equality and composition of maps are A-type-definable. More precisely, this means one can find a type-definition $(\{T_{ij}\}, \{F_{ij}\})$ over A with the following extra properties:

(1) For all $(i,j) \in R$, the set

$$\{(t,u) \in T_{ij} : f_t^{ij} = f_u^{ij}\}$$

is A-type-definable.
(2) Whenever $(i,j)(j,k), (i,k) \in R$, the set

$$\{(t,u,v) \in T_{ij} \times T_{jk} \times T_{ik} : f_u^{jk} \circ f_t^{ij} = f_v^{ij}\}$$

is A-type-definable.

Later, we will see that faithfully type-definable groupoid spines can be definably quotiented by equality of maps (see Propositions 8.5 and 8.6). This is the reason for the word *faithful*.

In general, it is not clear that a type-definable groupoid spine should be faithfully type-definable: for example, the obvious definition of the formula $f_t^{ij} = f_u^{ij}$ (for given $(i,j) \in R$) would require quantifying over the type-definable set X_i – and this is only allowed if X_i is in fact definable.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

8.2. Preservation under extensions

The main advantage of ordinary type-definability (versus faithful type-definability) is that it allows for a straightforward extension result:

Lemma 8.4. *Let A be a parameter set, and let $(I, <, \mathcal{X}, R, Mor)$ be an A-type-definable groupoid spine which extends to a groupoid. Then the resulting extension is also A-type-definable.*

Proof. Fix $(\{T_{ij}\},\{F_{ij}\})$, a type-definition of $(I, <, \mathcal{X}, R, Mor)$ over A. We show how to define T_{ij} and F_{ij} when $(i,j) \notin R$.

So let $(i,j) \notin R$. First, suppose $i \neq j$. Then necessarily $i > j$, so that $(j,i) \in R$ and $\mathrm{Mor}(i,j)$ is the collection of inverses of maps in $\mathrm{Mor}(i,j)$. Now define $T_{ij} = T_{ji}$, and set $(x,y,t) \in F_{ij}$ if and only if $(y,x,t) \in F_{ji}$; clearly this gives an A-type-definable family representing all maps in $\mathrm{Mor}(i,j)$.

Now suppose $i = j$. By definition, the relation $R \subseteq I^2$ is non-empty, which implies there must be some $k \in I$ with $k \neq i$. By the previous paragraph, we may assume $(i,k),(k,i) \in R$. Note, then, that $\mathrm{Mor}(i,i)$ is the collection of compositions $g \circ f$ for $f \in \mathrm{Mor}(i,k)$ and $g \in \mathrm{Mor}(k,i)$. So, we let $T_{ii} = T_{ik} \times T_{ki}$. And for $(x,y,(t,u)) \in X_i \times X_i \times T_{ii}$, we set $(x,y,(t,u)) \in F_{ii}$ if and only if there is z with $(x,z,t) \in F_{ik}$ and $(z,y,u) \in F_{ki}$. Then F_{ii} is a projection of an A-type-definable set, so is A-type-definable. □

8.3. An elimination result

Our main goal is to interpret, under appropriate assumptions, a group from a type-definable groupoid spine. The group we hope to construct is the group of morphisms from an object to itself in an extension to a groupoid. Notice, though, that our definitions of type-definability and faithful type-definability have a flaw in this regard: a type-definition is allowed to repeat one map several times (i.e. the natural map $T_{ij} \to \mathrm{Mor}(i,j)$ is surjective but might not be injective). Thus, one cannot hope to build a group structure on the parameter space T_{ii} directly – instead, one needs to build a quotient of it. Forming the necessary quotient object (in a type-definable manner) is the main use of faithful type-definability – in this case, it essentially amounts to an elimination result for certain hyperimaginaries, which we treat separately below.

Proposition 8.5. *Let A be a parameter set. Let X and T be A-type-definable sets, and $Y \subseteq X \times T$ a relatively A-definable relation. For $t \in T$, denote $Y_t := \{x : (x,t) \in Y\}$, and let $\sim$ be the equivalence relation $Y_t = Y_u$ on T. If $\sim$ is A-type-definable in T^2, then the quotient $T/\sim$ is A-type-definable in $\mathcal{M}^{eq}$: that is, there are an A-type-definable set Z in $\mathcal{M}^{eq}$, and a relatively A-definable function $f : T \to Z$, such that for $t,u \in T$ we have $f(t) = f(u)$ if and only if $t \sim u$.*

Proof. Write $X = \bigcap_{i \in I} X_i$ as a small intersection of A-definable sets. We may assume $\{X_i\}_{i \in I}$ is closed under finite intersections. Also, write $Y = Y_0 \cap (X \times T)$ for some A-definable set Y_0.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

By compactness, for any $t,u \in T$, we have $t \sim u$ if and only if $(Y_0)_t \cap X_i = (Y_0)_u \cap X_i$ for some $i \in I$. By compactness again (and since the X_i form a directed family), there is a single $i \in I$ so that for all $t,u \in T$, $t \sim u$ if and only if $(Y_0)_t \cap X_i = (Y_0)_u \cap X_i$.

Now let $T_0 \supset T$ be an A-definable set, and let E be the A-definable equivalence relation given by $(Y_0)_t \cap X_i = (Y_0)_u \cap X_i$ on T_0. Finally, let $Z = T_0/E$, an A-definable set in $\mathcal{M}^{eq}$. Then we have a relatively A-definable function $f : T \to Z$ (factoring as $T \to T_0 \to Z$). And for $t,u \in T$, we have $f(t) = f(u)$ if and only if t and u are E-equivalent, if and only if $(Y_0)_t \cap X_i = (Y_0)_u \cap X_i$, if and only if $t \sim u$. □

8.4. Interpreting a group in the faithful case

We now use Proposition 8.5 to interpret a group from a faithfully type-definable groupoid:

Proposition 8.6. *Let A be a parameter set, and let $(I, <, \mathcal{X}, R, Mor)$ be a faithfully A-type-definable groupoid spine, with $R = I^2$. (Again, a groupoid spine with $R = I^2$ is just an actual groupoid – or more precisely a connected groupoid of bijections equipped with a linear ordering of its objects.) Then for any i, the group $\mathrm{Mor}(i,i)$, equipped with its action on X_i, is A-type-definable in $\mathcal{M}^{eq}$. More precisely, in $\mathcal{M}^{eq}$, there are an A-type-definable group G_i acting A-type-definably on X_i, and an isomorphism of group actions $(G_i, X_i) \leftrightarrow (\mathrm{Mor}(i,i), X_i)$ fixing X_i pointwise.*

Proof. Let $\{T_{ij}\}$ and $\{F_{ij}\}$ be a faithful type-definition over A, and fix $i \in I$. By Remark 8.2, F_{ii} is relatively A-definable in $X_i \times X_j \times T_{ij}$, so that Proposition 8.5 can apply to F_{ii}. As in the setup of Proposition 8.5, let $\sim$ be the equivalence relation $f_t^{ii} = f_u^{ii}$ on T_{ii}. By faithful A-type-definability, $\sim$ is automatically A-type-definable, so that Proposition 8.5 applies.

Now by Proposition 8.5, after passing to $\mathcal{M}^{eq}$, we obtain an A-type-definable set G and a relatively A-definable function $h : T_{ii} \to G$ such that for $t,u \in T_{ii}$ we have $h(t) = h(u)$ if and only if $t \sim u$. So G is in natural bijective correspondence with $\mathrm{Mor}(i,i)$, and one can now endow G A-type-definably with the action of $\mathrm{Mor}(i,i)$ on X_i (so that for $g \in G$ and $x \in X_i$, gx is now defined). To finish the proof, we just need to show that the group operation and action (of G and X_i) are A-type-definable.

Now for $g \in G$ and $x,y \in X_i$, we have $gx = y$ if and only if there is $t \in T_{ii}$ with $h(t) = g$ and $f_t^{ii}(x) = y$. This shows that the action (i.e. the set defined by $gx = y$) is an image of an A-type-definable set, so is A-type-definable. Similarly, for $g_1, g_2, g_3 \in G$, we have $g_1 g_2 = g_3$ if and only if there are $t_1, t_2, t_3 \in T_{ii}$ with each $h(t_j) = g_j$ and $f_{t_1}^{ii} \circ f_{t_2}^{ii} = f_{t_3}^{ii}$. And again, this shows that the group operation is an image of an A-type-definable set, so is A-type-definable (note that we are using faithful A-type-definability here to get that $f_{t_1}^{ii} \circ f_{t_2}^{ii} = f_{t_3}^{ii}$ is A-type-definable). □

8.5. Regularity revisited

So far, we have an extension result for type-definable groupoid spines, and a group interpretation result for faithfully type-definable groupoid spines. Our goal now is to combine these into a single theorem. The trick is Lemma 8.7 below, which shows that in the regular case, there is no difference between ordinary and faithful type-definability:

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Lemma 8.7. *For every parameter set A, every A-type-definable regular groupoid spine is faithfully A-type-definable.*

Proof. Let $(\{T_{ij}\},\{F_{ij}\})$ be a type-definition of the regular groupoid spine $(I, <, \mathcal{X}, R, Mor)$ over A. We show that equality and composition of maps are A-type-definable:

- Let $(i,j) \in R$. We show that equality of maps is A-type-definable in T^2_{ij}. Now for $(t,u) \in T^2_{ij}$, regularity gives that $f^{ij}_t = f^{ij}_u$ if and only if there are $x \in X_i$ and $y \in X_j$ such that $f^{ij}_t(x) = f^{ij}_u(x) = y$. This shows that the equality relation in T^2_{ij} is a projection of an A-type-definable set, so is A-type-definable.
- Let $(i,j),(j,k),(i,k) \in R$. We show that composition is A-type-definable in $T_{ij} \times T_{jk} \times T_{ik}$. Now for $(t,u,v) \in T_{ij} \times T_{jk} \times T_{ik}$, regularity gives that $f^{jk}_u \circ f^{ij}_t = f^{ik}_v$ if and only if there are $x \in X_i$, $y \in X_j$ and $z \in X_k$ such that $f^{ij}_t(x) = y$, $f^{jk}_u(y) = z$ and $f^{ik}_v(x) = z$. This shows that the graph of composition is a projection of an A-type-definable set, so is A-type-definable.

□

8.6. Main result

Finally, we conclude with the main theorem of this section. The reader should view this as an analogue of the group configuration for 1-based stable theories:

Theorem 8.8. *Let A be a parameter set, and let $(I, <, \mathcal{X}, R, Mor)$ be an A-type-definable regular groupoid spine with $|I| \geq 3$. Then for any i, there is an A-type-definable group in $\mathcal{M}^{eq}$ acting regularly and type-definably on X_i. In particular, for any $e \in X_i$, there is an Ae-type-definable group structure on the set X_i, with e as the identity element.*

Proof. By Theorem 7.7, $(I, <, \mathcal{X}, R, Mor)$ extends to a groupoid. By Lemma 7.8, the resulting groupoid is regular; by Lemma 8.4, the resulting groupoid is A-type-definable; thus by Lemma 8.7, the resulting groupoid is faithfully A-type-definable. Now apply Proposition 8.6 to obtain an A-type-definable group G acting A-type-definably and regularly on X_i.

To build a group structure on X_i itself, fix any $e \in X_i$. There is now an Ae-type-definable bijection $G \to X_i$, given by $g \mapsto ge$ (by regularity). This allows us to transfer the group structure from G to X_i Ae-type-definably. Note that e is identified with the identity, because the identity sends e to e. □

9. Construction of a regular groupoid spine

We now return to the topological setting considered earlier in the paper. Thus, we fix $\mathcal{M}$, a sufficiently saturated t-minimal structure with the independent neighbourhood property; and we fix an elementary extension $\mathcal{N}$ of $\mathcal{M}$, sufficiently saturated with respect to $\mathcal{M}$. As before, we consider infinitesimal neighbourhoods of $\mathcal{M}$-tuples as $\mathcal{M}$-type-definable sets in $\mathcal{N}$.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Our goal now is to show that if $\mathcal{M}$ is non-trivial (see Definition 9.7) and topologically 1-based, then we can construct a type-definable regular groupoid spine in $\mathcal{M}$. Our main tool is the structure theorem for topologically 1-based types, Theorem 6.4.

9.1. Families of bijections

First, we show that from a certain configuration of interalgebraicities, we can construct a regular family of homeomorphisms between infinitesimal neighbourhoods:

Lemma 9.1. *Let (a,b,t) be real tuples. Assume:*

(1) *Any two of a,b,t are symmetrically independent.*

(2) *a and b are interalgebraic over t.*

(3) *(t,ab) is additive.*

(4) *$\mathrm{tp}(ab/t)$ is topologically 1-based.*

Then $\mu(abt/\emptyset)$ is the graph of a regular family of homeomorphisms $\mu(a/\emptyset) \to \mu(b/\emptyset)$. More precisely, there is a family F of homeomorphisms $f : \mu(a/\emptyset) \to \mu(b/\emptyset)$ such that:

(i) *For each $t' \in \mu(t)$, the fibre $\mu(abt/\emptyset)_{t'} \subset \mu(ab/\emptyset)$ is the graph of a homeomorphism $f_{t'} \in F$.*

(ii) *For each $f \in F$, there is $t' \in \mu(t/\emptyset)$ with $f_{t'} = f$.*

(iii) *For each $a' \in \mu(a/\emptyset)$ and $b' \in \mu(b/\emptyset)$, there is exactly one $f \in F$ with $f(a') = b'$.*

Proof. By symmetric independence, Lemma 4.8 gives: $\mu(at/\emptyset) = \mu(a/\emptyset) \times \mu(t/\emptyset)$ and $\mu(bt/\emptyset) = \mu(b/\emptyset) \times \mu(t/\emptyset)$. By Lemma 4.9, since $a \in \mathrm{acl}(bt)$, projecting gives a homeomorphism

$$\mu(abt/\emptyset) \to \mu(bt/\emptyset) = \mu(b/\emptyset) \times \mu(t/\emptyset)$$

(technically this requires that (a,bt) is additive, but additivity follows automatically from $a \in \mathrm{acl}(bt)$). Similarly, the projection $\mu(abt/\emptyset) \to \mu(a/\emptyset) \times \mu(t/\emptyset)$ is a homeomorphism. It follows that $\mu(abt/\emptyset)$ is the graph of a family of homeomorphisms $\mu(a/\emptyset) \to \mu(b/\emptyset)$. Let F be the set of distinct homeomorphisms $\mu(a/\emptyset) \to \mu(b/\emptyset)$ occurring in this way.

It remains to show that given $a' \in \mu(a/\emptyset)$ and $b' \in \mu(b/\emptyset)$, there is exactly one $f \in F$ with $f(a') = b'$. First, the fact that there is at least one such f follows by the additivity of (t,ab): indeed, since a and b are symmetrically independent, we have $\mu(ab/\emptyset) = \mu(a/\emptyset) \times \mu(b/\emptyset)$; then by additivity and Lemma 4.9, the projection

$$\mu(abt/\emptyset) \to \mu(ab/\emptyset) = \mu(a/\emptyset) \times \mu(b/\emptyset)$$

is surjective.

It remains to show that there is at most one $f \in F$ sending a' to b'. But this is exactly the content of Theorem 6.4: if $f,g \in F$ both send a' to b', then there are $t',t'' \in \mu(t/\emptyset)$ so that $\mu(abt/\emptyset)_{t'}$ is the graph of f and $\mu(abt/\emptyset)_{t''}$ is the graph of g. Then both $\mu(abt/\emptyset)_{t'}$ and $\mu(abt/\emptyset)_{t''}$ contain (a',b') – and since $\mathrm{tp}(ab/t)$ is topologically 1-based, Theorem 6.4 gives that these two fibres are equal. That is, f and g have the same graph, and so they are the same function. □

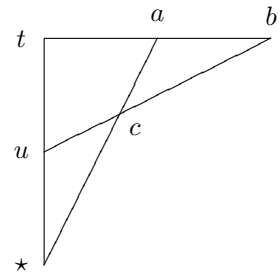
https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

9.2. Pre-group configurations

We now aim towards a topological analogue of a group configuration theorem. Towards that end, we define:

Definition 9.2. A *one-dimensional topological pre-group configuration* in $\mathcal{M}$ is a 5-tuple (a,b,c,t,u) such that all of the following hold:

(1) Each of a,b,c,t,u has dimension 1.
(2) $\dim(abctu) = 3$.
(3) a and b are interalgebraic over t, and b and c are interalgebraic over u.



Thus, a one-dimensional topological pre-group configuration consists of all but one point (denoted $\star$ in the above diagram) in a conventional one-dimensional group configuration. One can easily check that all dimension computations in a one-dimensional pre-group configuration match those in a usual group configuration. Precisely, we leave the following to the reader:

Lemma 9.3. *Let (a,b,c,t,u) be a one-dimensional topological pre-group configuration in $\mathcal{M}$. Separately, let a',t',u' be $\mathbb{Q}$-linearly independent real numbers, let $b' = a' + t'$ and let $c' = b' + u'$. Then the map $f : \{a,b,c,t,u\} \mapsto \{a',b',c',t',u'\}$ – sending each element to its primed counterpart – preserves all dimensions. More precisely, if $A,B \subset \{a,b,c,d,e\}$, then* $\dim(A/B)$ *(in the sense of $\mathcal{M}$) is the same as the pregeometry dimension of $f(A)$ over $f(B)$ in the sense of the strongly minimal group $(\mathbb{R},+)$.*

In particular, all dimension computations in a one-dimensional topological pre-group configuration have the additivity property:

Corollary 9.4. *Let (a,b,c,t,u) be a one-dimensional pre-group configuration. If $A,B \subset \{a,b,c,t,u\}$, then (A,B) is additive.*

Proof. The analogous property holds in $(\mathbb{R},+)$, by exchange. □

9.3. Group configurations

We now define what it means for a one-dimensional topological pre-group configuration (a,b,c,t,u) to be an actual topological group configuration. In a stable theory, the requirement is the existence of a non-algebraic tuple $v \in \operatorname{acl}(ac) \cap \operatorname{acl}(tu)$ (i.e. the sixth point of a conventional group configuration). In the topological setting, this is (a priori)

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

too much to ask for. On the other hand, the same condition can be expressed using 1-basedness: if such a v as above exists in a stable context, then v is interalgebraic with $\mathrm{Cb}(\mathrm{stp}(ac/tu))$, which implies that $\mathrm{stp}(ac/tu)$ is 1-based (and conversely, if $\mathrm{stp}(ac/tu)$ is 1-based, one can take its canonical base for v). Thus, in the topological setting, we simply adapt the version of the definition using 1-basedness:

Definition 9.5. Let (a,b,c,t,u) be a one-dimensional topological pre-group configuration. We say that (a,b,c,t,u) is a *one-dimensional topological group configuration* if $\mathrm{tp}(ac/tu)$ is topologically 1-based.

We now show:

Theorem 9.6. *Let (a,b,c,t,u) be a one-dimensional topological group configuration in $\mathcal{M}$. Then there is an $\mathcal{M}$-type-definable group structure on the set $\mu(a/\emptyset)$ with identity element a (and similarly for b and c).*

Proof. We will build an $\mathcal{M}$-type-definable regular groupoid spine $(I, <, \mathcal{X}, R, \mathrm{Mor})$ in the structure $\mathcal{N}$, where:

- $I = \{1,2,3\}$ with the usual order.
- $X_1 = \mu(a/\emptyset)$, $X_2 = \mu(b/\emptyset)$ and $X_3 = \mu(c/\emptyset)$.
- $R = \{(i,j) : i < j\}$.

The theorem will then follow by Theorem 8.8. Precisely, Theorem 8.8 will provide an $\mathcal{M}a$-type-definable group structure on the set $X_1 = \mu(a/\emptyset)$, with identity element a (and similarly for b and c). But since a, b and c are tuples from $\mathcal{M}$, these groups are actually $\mathcal{M}$-type-definable.

Our task, then, is to define the sets $\mathrm{Mor}(i,j)$ for $i < j$. We will do this by applying Lemma 9.1 to each of (a,b,t), (b,c,u) and (a,c,tu). We leave it to the reader to check the various independence and algebraicity requirements needed to apply Lemma 9.1 in each of these cases; note that there is no need to check the additivity requirements, as Corollary 9.4 gives them for free.

So, in order to apply Lemma 9.1, the main thing to check is that each of $\mathrm{tp}(ab/t)$, $\mathrm{tp}(bc/u)$ and $\mathrm{tp}(ac/tu)$ is topologically 1-based. But indeed:

- By assumption, $t \in \mathrm{acl}(ab)$, so $\dim(t/ab) = 0$, and thus $\dim(t/\mathrm{germ}(ab/t))$ must also be 0. Thus, $\mathrm{tp}(ab/t)$ is topologically 1-based.
- By the same reason as above, $\mathrm{tp}(bc/u)$ is also topologically 1-based.
- Finally, $\mathrm{tp}(ac/tu)$ is topologically 1-based because of the assumption that (a,b,c,t,u) is a one-dimensional topological group configuration.

So, we may legally apply Lemma 9.1 in each of the above instances. Denote the three resulting families of bijections by $\mathrm{Mor}(i,j)$ (in the obvious way, so that $\mathrm{Mor}(i,j)$ is the family of maps from X_i to X_j).

If we can show that the data defined so far gives a groupoid spine, then it will automatically be $\mathcal{M}$-type-definable and regular. That is:

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

- Each X_i is an infinitesimal neighbourhood, so is $\mathcal{M}$-type-definable.
- Similarly, each $\mathrm{Mor}(i,j)$ is represented as in Definition 8.1 by an $\mathcal{M}$-type-definable family – precisely by $\mu(abt/\emptyset)$, $\mu(bcu/\emptyset)$ and $\mu(actu/\emptyset)$, respectively.
- By Lemma 9.1, each $\mathrm{Mor}(i,j)$ has the required property for the definition of regularity.

Thus, our only remaining task is showing that $(I, <, \mathcal{X}, R, \mathrm{Mor})$ indeed gives a groupoid spine. And for this, the only thing to check is closure under composition. So let $f \in \mathrm{Mor}(1,2)$, and let $g \in \mathrm{Mor}(2,3)$. We want to show that $g \circ f \in \mathrm{Mor}(1,3)$.

Since $f \in \mathrm{Mor}(1,2)$, there is $t' \in \mu(t/\emptyset)$ so that the fibre $\mu(abt/\emptyset)_{t'}$ is the graph of f. Similarly, since $g \in \mathrm{Mor}(2,3)$, there is $u' \in \mu(u/\emptyset)$ so that the fibre $\mu(bcu)_{u'}$ is the graph of g. We will show that the fibre $\mu(actu/\emptyset)_{t'u'}$ is the graph of $g \circ f$. Let $h : \mu(a/\emptyset) \to \mu(c/\emptyset)$ be the bijection whose graph is the fibre above (t',u'). So we want to show that $g \circ f = h$.

To show this, let $a' \in \mu(a/\emptyset)$. Note that $abctu \in \mathrm{acl}(atu)$ – thus $\dim(atu) = 3$, and thus $\mu(atu/\emptyset) = \mu(a/\emptyset) \times \mu(t/\emptyset) \times \mu(u/\emptyset)$. Moreover, by Lemma 4.9, the projection gives a homeomorphism

$$\mu(abctu/\emptyset) \to \mu(atu/\emptyset) = \mu(a/\emptyset) \times \mu(t/\emptyset) \times \mu(u/\emptyset).$$

Since $a' \in \mu(a/\emptyset)$, $t' \in \mu(t/\emptyset)$ and $u' \in \mu(u/\emptyset)$, we conclude that $(a',t',u') \in \mu(atu/\emptyset)$, and thus there are b',c' so that $(a',b',c',t',u') \in \mu(abctu/\emptyset)$. Thus:

- $(a',b',t') \in \mu(abt/\emptyset)$, and so $b' = f(a')$;
- Similarly, $(b',c',u') \in \mu(bcu/\emptyset)$, and so $c' = g(b')$;
- And finally, $(a',c',t',u') \in \mu(actu/\emptyset)$, and so $c' = h(a')$.

It now follows that $g \circ f(a') = h(a')$. Since $a' \in \mu(a/\emptyset)$ was arbitrary, it follows that $g \circ f = h$, and this proves the theorem. □

9.4. Building a topological group configuration

We now show that, combined with a non-triviality assumption, topological 1-basedness implies the type-definability of an infinite group.

Definition 9.7. Say that $\mathcal{M}$ is *non-trivial* if for some parameter set A, there are $a,b,c \in M$ with $\dim(abc/A) = 2$ and each of a,b,c is algebraic over A with the other two.

Corollary 9.8. *Assume $\mathcal{M}$ is non-trivial and topologically 1-based. Then there are $a \in M$ with $\dim(a) = 1$, and an infinite $\mathcal{M}$-type-definable group structure on the set $\mu(a)$, with identity element a.*

Proof. Let A be a parameter set and $(a,b,t) \in M^3$ with $\dim(abt/A) = 2$ and each of a,b,t algebraic over the other two. By Lemma 5.7, we may assume $A = \emptyset$. By Lemma 6.1, there is a pair $(c,u) \models \mathrm{tp}(at/b)$ so that

$$\dim(abctu) = \dim(abt) + \dim(at/b) = 3.$$

It follows easily that (a,b,c,t,u) is a one-dimensional pre-group configuration in $\mathcal{M}$. Since $\mathcal{M}$ is topologically 1-based, (a,b,c,t,u) is then automatically a one-dimensional topological group configuration in $\mathcal{M}$. Now apply Theorem 9.6 to get an $\mathcal{M}$-type-definable group

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

structure on $\mu(a/\emptyset)$ with identity element a, and note that since $\dim(a) = 1$, $\mu(a/\emptyset) = \mu(a)$. □

We can also give some more local analysis:

Definition 9.9. Let $p = \mathrm{tp}(a/A)$ where $\dim(a/A) = 1$.

(1) Say that p is *trivial* if there are $(a,b,c) \models p$ with $\dim(abc/A) = 3$ and each of a,b,c algebraic over A with the other two.

(2) Say that p is *topologically locally modular* if for any $a,b \models p$ and any real tuple c, if $\dim(a/Ac) = \dim(b/Ac) = 1$ and a and b are interalgebraic over Ac, then $\mathrm{tp}(ab/Ac)$ is topologically 1-based over A.

So p is topologically locally modular if, at the level of germs and topological 1-basedness, there are no large families of plane curves inside p^2. Then the exact same argument gives:

Corollary 9.10. *Let $p = \mathrm{tp}(a/A)$ where $\dim(a/A) = 1$. Assume that p is non-trivial and topologically locally modular. Then there is an $\mathcal{M}$-type-definable group structure on the set $\mu(a/A)$, with identity element a.*

10. Topologizing the group

Throughout this section, assume $e \in M$ with $\dim(e) = 1$. Moreover, assume $(G_0,\cdot)$ is an $\mathcal{M}$-type-definable group with universe $\mu(e)$ and identity element e.

Our goal will be to replace the group structure on $\mu(e)$ with a suitable type-definable group over a small set in $\mathcal{M}$. A key feature will be that the group we build is a topological group with the topology inherited from the given t-minimal topology τ.

10.1. Step 1: a type-definable group in $\mathcal{M}$

First, we build a type-definable supergroup of G_0, defined over a small set in $\mathcal{M}$:

Lemma 10.1. *Let A be any parameter set such that $\dim(e/A) = 1$. Then there are a countable set B (not necessarily containing A), and a B-type-definable group $(G,\cdot)$ with $G_0 \leq G$, so that the following hold:*

(1) *G is an open subseteq of M.*

(2) *Every element of G realizes $\mathrm{tp}(e/A)$.*

Proof. Since $\dim(e/A) = 1$, the realizations of any formula in $\mathrm{tp}(e/A)$ form a neighbourhood of e. So by Fact 2.1, there is an open neighbourhood U_0 of e all of whose elements realize $\mathrm{tp}(e/A)$. Shrinking U_0 if necessary, we may assume, by saturation, that it is $\mathcal{M}$-definable.

Since G is type-definable over $\mathcal{M}$, its operation and inverse are given by (restrictions of) $\mathcal{M}$-definable functions, say $(x,y) \models m(x,y)$ and $x \mapsto i(x)$.

Now recall that $G_0 = \mu(e)$ is the intersection of all $\mathcal{M}$-definable open neighbourhoods of e. By compactness, there is such a neighbourhood $U_1 \subseteq U_0$ satisfying the universal group axioms. That is:

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

- If $a,b,c \in U_1$, then $m(m(a,b),c)$ and $m(a,m(b,c))$ are both defined and equal.
- If $a \in U_1$, then $m(a,e) = m(e,a) = a$.
- If $a \in U_1$, then $i(a)$ is defined and $m(a,i(a)) = m(i(a),a) = e$.

We now build a decreasing sequence of definable open neighbourhoods $U_i \subseteq U_1$ (for $i \geq 2$) whose intersection is closed under m and i, and so is a group. To do this, it suffices to check:

Claim 10.2. *Let W be any $\mathcal{M}$-definable open neighbourhood of e. Then there is an $\mathcal{M}$-definable open neighbourhood $V \subseteq U_1$ of e such that $m(V^2),i(V) \subseteq W$.*

Proof. Since W is $\mathcal{M}$-definable, we have $\mu(e) \subseteq W$, and thus all products and inverses of elements of $\mu(e)$ belong to W. By compactness, the same holds after replacing $\mu(e)$ with some $\mathcal{M}$-definable neighbourhood of e. □

Now if we have constructed all U_i, then clearly $G = \bigcap U_i$ is a group with the operations m and i. Thus, G_0 is a subgroup. Since G is an intersection of countably many definable open sets, it is open (by Fact 2.1). Finally, G is defined over a countable set in $\mathcal{M}$ (namely, the parameters used to define each U_i). □

10.2. Step 2: a locally topological group

Recall that a version of generic continuity holds in our setting (see Lemma 4.9). It follows that the group G built above is a 'generically topological group': the operation and inverse are almost everywhere continuous. In [15], Maříková initiated a general method for turning a 'generically topological' group into a genuinely topological group. We recall one version of Maříková's method, following [2]:

Definition 10.3. Let G be a group, and $d \in G$. By G_d, we mean the group on the underlying set G with operation $(x,y) \mapsto xd^{-1}y$ and inverse $x \mapsto dx^{-1}d$.

The group G_d is obtained from G by a relabelling, translating every element by d. Thus, d is the identity of G_d.

Definition 10.4. Let G be a group whose underlying set is also a topological space.

(1) We say G is *generically topological* if there are dense open sets $U \subseteq G^2$ and $V \subseteq G$ such that $(x,y) \mapsto xy$ and $x \mapsto x^1$ are continuous on U and V, respectively.

(2) We say G is *locally topological* if there is an open neighbourhood $U \subseteq G$ of the identity such that $(x,y) \mapsto xy$ and $x \mapsto x^{-1}$ are continuous on U^2 and U, respectively.

The following is from [2, §7]:

Fact 10.5. *Let G be a generically topological group. Then there is a dense open set W such that G_d is locally topological for all $d \in W$.*

We also need the following reinterpretation of generic continuity:

Lemma 10.6. *Let $U \subseteq M^n$ be definable and open, and let $f : U \to M$ be definable. Then f is continuous on a dense open set.*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Proof. Let U' be the set of $u \in U$ so that f is continuous on a neighbourhood of u. We show that U' is dense. If not, there is a definable open $V \subseteq U$ disjoint from U'. Let A be a set so that f and V are A-definable. Since V is open, there is $v \in V$ with $\dim(v/A) = n$. Then by generic continuity, Corollary 3.17, f is continuous on a neighbourhood of v, a contradiction. □

We conclude:

Lemma 10.7. *Let G be a type-definable group over a small set A, whose universe is an open subseteq of M. Then for any $d \in G - \mathrm{acl}(A)$, the group G_d is locally topological with the topology inherited from M.*

Proof. The operation and inverse of G are given by A-definable functions, say $m(x,y)$ and $i(x)$. By compactness, there is an A-definable $X \supset G$ so that m and i are defined on elements of X. By Lemma 10.6, m and i are continuous on dense open subsets of X^2 and X, respectively. But G is open in X, so it follows that m and i are also continuous on dense open subsets of G^2 and G, respectively. Thus, G is generically topological. By Fact 10.5, there is a dense open set $W \subseteq G$ so that G_d is locally topological for all $d \in W$. Now let Y be the set of all $d \in X$ so that the maps $(x,y) \mapsto xd^{-1}y$ and $x \mapsto dx^{-1}d$ are defined and continuous on neighbourhoods of (d,d) and d, respectively. Then Y contains W and is A-definable.

Now let $d \in G - \mathrm{acl}(A)$. If G_d is not locally topological, then $d \notin Y$, and thus by t-minimality, $X - Y$ contains a neighbourhood of d. Since G is open, $X - Y$ contains a neighbourhood of d in G. This contradicts that $W \subseteq Y$ and W is dense. □

10.3. Step 3: A topological group

Finally, we show that from a locally topological group, one can always extract a genuinely topological group:

Lemma 10.8. *Let G be a locally topological type-definable group whose underlying set is an open subseteq of G. Then there is a type-definable topological open subgroup of G defined over a countable set.*

Proof. By Fact 2.1, there is a definable neighbourhood U_0 of the identity contained in G. Since G is locally topological, there is a definable open neighbourhood $U_1 \subseteq U_0$ of the identity such that the operation and inverse are continuous on elements of U_1. By continuity, there is a definable open neighbourhood $U_2 \subseteq U_1$ of the identity such that $U_2 \cdot U_2$ and U_2^{-1} are contained in U_1. Continue in this manner, constructing definable open neighbourhoods $U_2 \supset U_3 \supset U_4 \ldots$ of the identity such that $U_{i+1} \cdot U_{i+t} \subseteq U_i$ and $U_{i+1}^{-1} \subseteq U_i$. Then $\bigcap U_i$ satisfies the requirements of the lemma. □

We conclude the main result of this subsection. Recall that we have a fixed $\mathcal{M}$-type-definable group structure G_0 on $\mu(e)$, with e as the identity. Now we show:

Proposition 10.9. *Let A be any parameter set so that $\dim(e/A) = 1$. Then there are a countable set B, and a B-type-definable group G, such that:*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

(1) *e is the identity of G.*

(2) *G is open in M.*

(3) *G is a topological group with the topology inherited from M.*

(4) *Every element of G realizes* $\mathrm{tp}(e/A)$.

Proof. By Lemma 10.1, we can find B and G satisfying all of the given requirements except being topological. Let $d \in G - \mathrm{acl}(B)$. Then the group G_d is Bd-type-definable and topological. So the pair (G_d, Bd) satisfies all requirements of the proposition except that its identity element is d, not e. But d and e have the same type over A, so there is an automorphism σ of $\mathcal{M}$ fixing A and satisfying $\sigma(d) = e$. Now the pair $(\sigma(G), \sigma(Bd))$ satisfies the requirements of the proposition. □

10.4. Concluding

We now give our main results at the level we have reached thus far. So, we drop the assumptions from the beginning of this section; but we retain our fixed pair $\mathcal{M}$ and $\mathcal{N}$ of t-minimal structures with independent neighbourhoods.

Theorem 10.10. *Assume $\mathcal{M}$ is non-trivial and 1-based. Then there are a countable parameter set B, and a B-type-definable group $(G, \cdot)$ such that:*

(1) *G is a non-empty open subseteq of M.*

(2) *$(G, \cdot)$ is a topological group with the subspace topology inherited from M.*

Proof. By Corollary 9.8, there are $a \in M$ with $\dim(a) = 1$, and an $\mathcal{M}$-type-definable group structure on $\mu(a)$. Now apply Proposition 10.9. □

Theorem 10.11. *Let $a \in M$ and A a parameter set so that* $\dim(a/A) = 1$, *and $p =$* $\mathrm{tp}(a/A)$ *is non-trivial and topologically locally modular. Then there are a countable set B, and a B-type-definable group $(G, \cdot)$ such that:*

(1) *The identity element of G is a.*

(2) *G is a non-empty open subseteq of M, and is contained in the realizations of p.*

(3) *$(G, \cdot)$ is a topological group with the subspace topology inherited from M.*

Proof. By Corollary 9.10, there is an $\mathcal{M}$-definable group structure on the set $\mu(a/A)$, with a as the identity. Now apply Proposition 10.9. □

10.5. Viscerality of t-minimal groups

Before moving on, we point out an isolated fact that falls out of our work in this section and may be of interest. Roughly, we show that *every* t-minimal expansion of a group with independent neighbourhoods is *already* both topological and visceral, at least up to editing the topology at finitely many points.

Recall that $\mathcal{M}$ is a fixed sufficiently saturated t-minimal structure with the independent neighbourhood property.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Theorem 10.12. *Assume that $\mathcal{M} = (M, \cdot, e, ...)$ is an expansion of a group. Then there is a definable topology τ' on $\mathcal{M}$ such that:*

(1) *With respect to τ', $\mathcal{M}$ is a visceral topological group.*

(2) *τ' agrees with τ on a cofinite subseteq of M.*

Proof. It is enough to find τ' satisfying (1). Then by Johnson's appendix (Theorem B.8), $\mathcal{M}$ has the independent neighbourhood property with respect to τ'; and thus by Lemma 3.18, τ' agrees with τ on a cofinite set.

The rest of the proof finds τ' satisfying (1). First, let $d \in M$ be generic, and consider the group M_d as in Definition 10.3. By Lemma 10.7, M_d is locally topological with respect to τ. Moreover, if we construct τ' satisfying (1) for M_d, the same τ' will work for M. Indeed, since (M,τ') and (M_d,τ') are interdefinable, they define the same structure – so (M,τ') will also be visceral; and clearly, if τ' is a group topology on M_d, it is also a group topology on M, by translation invariance. Thus, moving forward, we replace M with M_d and thereby assume M is locally topological with respect to τ.

By Lemma 10.8, there is thus a type-definable open topological subgroup $H \leq M$, defined over a countable set. Absorbing parameters, we assume H is $\emptyset$-definable. We then obtain a definable uniform structure on M by taking translations of basic neighbourhoods of e. Call the resulting topology τ'. So τ' is translation invariant (both left and right), and agrees with τ on H.

We will show that (M,τ') is a topological group, and $\mathcal{M}$ is visceral with respect to τ'. We start with the first of these goals, i.e. showing (M,τ') is topological. By [8, Lemma 2.11], it is enough to show that for each $g \in M$, conjugation by g is a *local homeomorphism*: this means there are τ'-open neighbourhoods $e \in U,V \subseteq M$ so that $u \mapsto gug^{-1}$ induces a homeomorphism $U \to V$. Since τ and τ' agree on H, it is equivalent to find such U,V which are τ-open.

Before proceeding, we collect some easy claims. Throughout, by a *generic* element of M over A, we mean an element g with $\dim(g/A) = 1$.

Claim 10.13. *Every element of M is a product of two generic elements.*

Proof. Given $g \in M$, let h be generic over g, and write $g = h(h^{-1}g)$. □

Claim 10.14. *Let $g,h \in G$ with $\dim(g,h) = 2$. Then the maps $x \mapsto gx$ and $x \mapsto xg$ define homeomorphisms $\mu(h) \to \mu(gh)$ and $\mu(h) \to \mu(hg)$.*

Proof. This follows by Lemma 4.9, noting that $\dim(h/g) = \dim(h) = 1$, so $\mu(h/g) = \mu(h/\emptyset) = \mu(h)$ (and similarly for gh). □

Claim 10.15. *Let $h \in H$ be generic. Then the maps $x \mapsto hx$ and $x \mapsto xh$ give homeomorphisms $\mu(e) \to \mu(h)$.*

Proof. This follows since H is topological, so these maps are homeomorphisms, and thus they preserve all $\mathcal{M}$-definable open sets (see also Lemma 11.2). □

Claim 10.16. *Let $g \in M$ be generic. Then $x \mapsto gx$ and $x \mapsto xg$ each give homeomorphisms $\mu(e) \to \mu(g)$.*

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Proof. Let $h \in H$ with $\dim(g,h) = 2$, and write $g = ah$ where $a = gh^{-1}$. So also $\dim(ah) = 2$. By Claims 10.15 and 10.14, $x \mapsto gx = ahx$ is now a composition of translation homeomorphisms $\mu(e) \to \mu(h) \to \mu(g)$ (translating by h then by a). The argument for $x \mapsto xg$ is similar, using b so that $g = hb$. □

Finally, we conclude:

Claim 10.17. *Let $g \in M$. Then conjugation by g induces a homeomorphism between neighbourhoods of e. Thus, (M,τ') is a topological group.*

Proof. By Claim 10.13, we may assume g is generic. Now using Claim 10.16, write conjugation by g as a composition of translation homeomorphisms $\mu(e) \to \mu(g) \to \mu(e)$, sending $x \mapsto gx \mapsto gxg^{-1}$. So conjugation by g gives a homeomorphism $\mu(e) \to \mu(e)$, and thus (by compactness) also a homeomorphism between neighbourhoods of e. □

We now turn to our second main goal mentioned above:

Claim 10.18. *$\mathcal{M}$ is t-minimal, and thus visceral, with respect to τ'.*

Proof. Viscerality follows from t-minimality by definition, since (M,τ') is uniform. We show t-minimality. Recall this means showing that for definable $X \subseteq M$, X has τ'-interior if and only if it is infinite.

First, suppose X has τ'-interior. Then by translation invariance, some translation Y of X contains a τ'-neighbourhood of e. Since τ and τ' agree on H (i.e. near e), Y also contains a τ-neighbourhood of e. By t-minimality of $(\mathcal{M},\tau)$, Y is infinite – thus so is X.

Now suppose X is infinite. Say X is A-definable, and pick $a \in X$ with $\dim(a/A) = 1$. By t-minimality of $(\mathcal{M},\tau)$, X contains a neighbourhood of a, and thus contains $\mu(a)$. By Claim 10.16, $Y = a^{-1}X$ contains $\mu(e)$, so in particular contains a τ-neighbourhood of e. Since τ and τ' agree near e, Y contains a τ'-neighbourhood of e. Then by translation invariance, X contains a τ'-neighbourhood of a. □

We have now shown that $(\mathcal{M},\tau')$ is a visceral topological group, which completes the proof of (1), and thus of the theorem. □

Remark 10.19. One could also prove Theorem 10.12 by quoting upcoming results of Johnson. According to private communication, he shows that any visceral expansion of a group admits a definable visceral group topology. This allows us to skip the proof that (M,τ') is a topological group above, concentrating only on the proof that $(\mathcal{M},\tau')$ is t-minimal. Indeed, once we know $(\mathcal{M},\tau')$ is t-minimal (and thus visceral), we obtain from Johnson a second definable visceral topology τ'' with respect to which $\mathcal{M}$ is a topological group. (In the end, one can show $\tau'' = \tau'$ using independent neighbourhoods and translation invariance, so Johnson's argument really gives an alternative proof that τ' is a group topology.)

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

11. The structure of the group

In this section, we assume that G is a type-definable group on a non-empty open subseteq of M, and is a topological group with the subspace topology inherited from M. Absorbing parameters, we assume G is $\emptyset$-type-definable. We denote the identity of G by e.

Our goal is to prove results analogous to the Hrushovski-Pillay classification of 1-based stable groups [10]). In a very coarse language, we will show:

(1) G is locally linear.

(2) G is locally abelian.

(3) There are no large type-definable families of germs of subgroups of G.

11.1. The infinitesimal subgroup

Let $\mu = \mu(e)$ be the intersection of all $\mathcal{M}$-definable open neighbourhoods of e. So μ is open and type-definable over $\mathcal{M}$. One easily checks:

Lemma 11.1. *μ is a subgroup of G.*

Proof. Let $a,b \in \mu$. We show $ab^{-1} \in \mu$. Indeed, let V be any $\mathcal{M}$-definable open neighbourhood of e. Since G is topological, the set $\{(x,y) : xy^{-1} \in V\}$ is an $\mathcal{M}$-definable open neighbourhood of e^2, so contains a set U^2 where U is an $\mathcal{M}$-definable open neighbourhood of e. By definition of μ, we get $a,b \in U$, and so $ab^{-1} \in V$. Since V was arbitrary, $ab^{-1} \in \mu$. □

Lemma 11.2. *Let $a \in G$. Then $\mu(a) = a \cdot \mu = \mu \cdot a$.*

Proof. The maps $x \mapsto ax$ and $x \mapsto xa$ are $\mathcal{M}$-definable homeomorphisms, so preserve $\mathcal{M}$-definable open sets. □

Lemma 11.3. *Let A be a parameter set, and $H \leq G^n$ an A-type-definable subgroup. Then the following are equivalent:*

(1) *H has non-empty interior.*

(2) *H is open.*

(3) *$\mu^n \subseteq H$.*

(4) *$\dim(H) = n$.*

Proof. (1) $\implies$ (2): Let a be an interior point of H. For any $b \in H$, the map $x \mapsto ba^{-1}x$ sends a to b and is a homeomorphism of both H and G. So b is also an interior point of H.

(2) $\implies$ (3): If H is $\mathcal{M}$-definable and open, it contains the intersection of all $\mathcal{M}$-definable open neighbourhoods of e, which is precisely μ^n.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

(3) $\implies$ (4): If H contains μ^n, then

$$n = \dim(\mu^n) \leq \dim(H) \leq \dim(G^n) = n,$$

so $\dim(H) = n$.

(4) $\implies$ (1): If $\dim(H) = n$, then there is $h \in H$ with $\dim(h/A) = n$. Then every A-definable superset of h contains a neighbourhood of h, and by Fact 2.1, so does H. □

11.2. Characterization of topological 1-basedness

In this subsection, we prove the main geometric result concerning topological 1-basedness in G. This should be seen as analogous to the Hrushovski-Pillay characterization of 1-based stable groups.

We will need the following, which is an elementary group-theoretic exercise:

Fact 11.4. *Let H be a group, and $X \subseteq H$ a non-empty subset. Then the following are equivalent:*

(1) *Any two left translates of X are equal or disjoint.*

(2) *Any two right translates of X are equal or disjoint.*

(3) *X is a left coset of a subgroup of G^n.*

(4) *X is a right coset of a subgroup of G^n.*

(5) *X is closed under the map $(x,y,z) \mapsto xy^{-1}z$.*

Recall that in Section 5.6 we introduced *local linearity* for t-minimal topological groups. We now give this notion a more general definition in the type-definable setting.

Lemma 11.5. *Let $a \in G^n$ and A a parameter set. The following are equivalent:*

(1) *There are a parameter set A, and an A-type-definable subgroup $H \leq G^n$, so that $\mathrm{germ}(a/A) = \mathrm{germ}_a(a \cdot H)$.*

(2) *There are a parameter set A, and an A-type-definable subgroup $H \leq G^n$, so that $\mathrm{germ}(a/A) = \mathrm{germ}_a(H \cdot a)$.*

(3) *$\mu(a/A)$ is a left coset of a subgroup of μ^n.*

(4) *$\mu(a/A)$ is a right coset of a subgroup of μ^n.*

Proof. (1) and (2) are equivalent by Fact 11.4. Similarly, (3) and (4) are equivalent by Fact 11.4.

(1) $\to$ (3): Let A be a parameter set, and $H \leq G^n$ an A-type-definable subgroup, so that $\mathrm{germ}(a/A) = \mathrm{germ}_a(a \cdot H)$. Then $\mu(a/A) = \mu(a) \cap (a \cdot H)$. By Lemma 11.2, this is the same as $a \cdot (\mu \cap H)$. By Lemma 11.1, $\mu \cap H$ is a subgroup of μ^n. This implies (3).

(3) $\to$ (1): Assume $\mu(a/A)$ is a left coset of a subgroup. By Fact 11.4, $\mu(a/A)$ is closed under $(x,y,z) \mapsto xy^{-1}z$. Now using compactness, inductively build $\mathcal{M}$-definable d-approximations $X_1 \supset X_2 \ldots$ of $\mathrm{tp}(a/A)$ at a so that $X_{i+1}X_{i+1}^{-1}X_{i+1} \subseteq X_i$ for all i. Then $X = \bigcap X_i$ is closed under $(x,y,z) \mapsto xy^{-1}z$, so is a left coset of the subgroup $a^{-1}X$. Since each X_i realizes $\mathrm{germ}(a/A)$, Fact 2.1 gives that X also realizes $\mathrm{germ}(a/A)$. Finally,

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

$a^{-1}X$ is defined over countably many parameters from $\mathcal{M}$, so satisfies the requirements of (1). □

Definition 11.6. Let $a \in G^n$ and A a parameter set. We say that $\mathrm{tp}(a/A)$ is *locally linear* if the equivalent conditions of Lemma 11.5 hold.

Definition 11.7. We say that G is *locally linear* if $\mathrm{tp}(a/A)$ is locally linear for all $a \in G^n$ and all parameter sets A.

We also give a general definition of topological 1-basedness for G:

Definition 11.8. We say that G is *topologically 1-based* if for every $a \in G^m$, every $b \in G^n$ and every parameter set A, $\mathrm{tp}(a/Ab)$ is topologically 1-based over A.

Remark 11.9. By Remark 5.12, if exchange holds, then Definition 11.8 is equivalent to the analogous assertion restricted to the case when $n = 2$.

In this language, we now prove our main analogue of the Hrushovski-Pillay theorem on 1-based stable groups:

Theorem 11.10. *The following are equivalent:*

(1) *G is topologically 1-based.*

(2) *G is locally linear.*

In particular, both conditions hold if $\mathcal{M}$ is topologically 1-based.

Proof. (1) $\Longrightarrow$ (2): Let $a \in G^n$ and let A be a parameter set. Let $X = a^{-1} \cdot \mu(a/A) \subset \mu^n$, so that $\mu(a/A) = a \cdot X$. Then let $b \in G^n$ so that $\dim(a,b/A) = \dim(a/A) + n$. So $\dim(b/A) = n$, which implies that $\mu(b/A) = \mu(b) = \mu^n \cdot b$. Also, a and b are symmetrically independent over A, and thus

$$\mu(a,b/A) = \mu(a/A) \times \mu(b/A) = (a \cdot X) \times (\mu^n \cdot b).$$

So for all $u \in \mu^n$, the fibre above $u \cdot b$ in $\mu(a,b/A)$ is $\mu(a/A) = a \cdot X$.

Now (a,b) and $(a \cdot b,b)$ are interdefinable over A by the map $(x,y) \mapsto (x \cdot y,y)$. It follows that the same map gives a homeomorphism $\mu(a,b/A) \to \mu(a \cdot b,b/A)$. Thus, for $u \in \mu^n$, the fibre above $u \cdot b$ in $\mu(a \cdot b,b/A)$ is the right translate by $u \cdot b$ of the corresponding fibre in $\mu(a,b/A)$, i.e. $a \cdot X \cdot u \cdot b$. Now, assuming (1), the type $\mathrm{tp}(a \cdot b/Ab)$ is topologically 1-based over A. Thus, by Theorem 6.4, any two of the sets $a \cdot X \cdot u \cdot b$ for $u \in \mu^n$ are equal or disjoint. The same holds after cancelling a on the left and b on the right. Thus, any two of the sets $X \cdot u$, for $u \in \mu^n$, are equal or disjoint. By Fact 11.4, it follows that X is a left coset of a subgroup of μ^n, and thus so is $a \cdot X = \mu(a/A)$.

(2) $\Longrightarrow$ (1): Let $a \in G^m$, $b \in G^n$ and A a parameter set. Assuming (2), the set $\mu(a,b/A)$ is a left coset of μ^{m+n}. It follows that there is a subgroup $H \leq G^m$ so that each fibre $\mu(a,b/A)_y$ for $y \in \mu(b/A)$ is either empty or a left coset of H. Since any two left cosets of H are equal or disjoint, it follows by Theorem 6.4 that $\mathrm{tp}(a/Ab)$ is 1-based over A. □

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

11.3. Few subgroups

Hrushovski and Pillay also prove that 1-based stable groups have few definable subgroups, in the sense that every connected definable subgroup is defined over $\mathrm{acl}(\emptyset)$. To adapt to our setting, we make two changes:

(1) Instead of connected subgroups, we consider germs of subgroups at the identity.
(2) Instead of asking for the germ to be defined over $\mathrm{acl}(\emptyset)$, we require (as in the definition of topological 1-basedness) that the germ is constant on an open set.

Before proving a precise statement, we need a technical lemma. In the general t-minimal setting (even with independent neighbourhoods), one can have a definable set – say over A – which is not a d-approximation for any of its elements over A. In contrast, Lemma 11.11 essentially says that groups and their cosets are always d-approximations at generic points:

Lemma 11.11. *Let X be an A-type-definable left coset of a subgroup of G^n of dimension d. Let $a \in X$ with $\dim(a/A) = d$. Let b be a basis of a over A, and π the corresponding projection sending a to b. Then:*

(1) *π is finite-to-one on X.*
(2) *X realizes $\mathrm{germ}(a/A)$.*

Proof. Note that (1) implies (2) by compactness, Lemma 3.6 and Fact 2.1 (as cofinally many A-definable supersets of X are d-approximations of $\mathrm{tp}(a/A)$).

We show (1). Since X is a left coset, there is an A-type-definable subgroup $H \leq G^n$ such that all non-empty fibres of π in X are left cosets of H (here we treat fibres as subsets of G^n, not G^{n-d}, as this will make the argument easier). So it suffices to show that H is finite.

Now assume H is infinite. Then there is $h \in H$ with $\dim(a,h/A) = \dim(a/A) + \dim(H) > \dim(X)$.

Since the fibre X_b is a left coset of H and contains a, it follows that $a \cdot h \in X$, and so $\dim(a \cdot h/A) \leq \dim(X) < \dim(a,h/A)$. On the other hand, we claim that $a \cdot h$ is interalgebraic with (a,h) over A, which will give a contradiction. Clearly, $a \cdot h \in \mathrm{acl}(A,a,h)$. For the converse, since $\pi(a \cdot h) = b$, we have $b \in \mathrm{acl}(A,a \cdot h)$. But $a \in \mathrm{acl}(Ab)$ by definition; thus $a \in \mathrm{acl}(A,a \cdot h)$, and thus $h \in \mathrm{acl}(A,a \cdot h)$ as well, since $h = a^{-1} \cdot (a \cdot h)$. □

We now give our 'few subgroups' result:

Theorem 11.12. *Assume G is topologically 1-based. Let $Y \subseteq G^m$ be A-type-definable, let $\{H_y : y \in Y\}$ be an A-type-definable family of d-dimensional subgroups of G^n and let $b \in Y$. Then the germ of H_y at the identity is constant for all $y \in \mu(b/A)$.*

Proof. First, we need:

Claim 11.13. *There is $a \in H_b$ so that $\dim(a,b/A) = \dim(b/A) + d$.*

Proof. Let $u \in \mu^d$ with $\dim(u,b/A) = d + \dim(b/A)$. By Lemma 11.11, there is a finite-to-one projection $\pi : H_b \to G^d$. Then $\pi(H_b)$ has dimension d, so by Lemma 11.3, $\pi(H_b)$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

contains μ^d. In particular, $u \in \pi(H_b)$. Let $a' \in H_b$ with $\pi(a') = u$. Then a' and u are interalgebraic over Ab. So (a',b) and (u,b) are interalgebraic over A, and thus $\dim(a',b/A) = d + \dim(d/A)$. Finally, since $\mathcal{M}$ is sufficiently saturated, there is a tuple a from $\mathcal{M}$ with $\mathrm{tp}(a,b/A) = \mathrm{tp}(a',b/A)$, and thus $\dim(a,b/A) = d + \dim(b/A)$ as well. □

Now fix a as in the claim. Since G is topologically 1-based, Theorem 11.10 gives that $\mu(a,b/A)$ is a left coset of a subgroup. In particular, all of the non-empty fibres $\mu(a,b/A)_y$ for $y \in \mu(b/A)$ are left cosets of the same group, say $L \leq G^n$.

Let $y \in \mu(b/A)$. We show that the germ of H_y at the identity is the same as the germ of L at the identity. As y is arbitrary, this will prove the theorem.

Now by the choice of a, we have that (a,b) is additive over A. So by Lemma 4.9, there is x so that $(x,y) \in \mu(a,b/A)$. By Corollary 4.11, $\mathrm{germ}(x/Ay)$ is realized by the fibre $\mu(a,b/A)_y$, and therefore by a coset of L. On the other hand, by Lemma 11.11, $\mathrm{germ}(a/Ab)$ is realized by H_b; and since $(x,y) \models \mathrm{tp}(a,b/A)$, $\mathrm{germ}(x/Ay)$ is also realized by H_y. So H_y agrees with the coset $x \cdot L$ in a neighbourhood of x; and thus the germ of H_y at the identity is realized by $x^{-1}(x \cdot L) = L$, as desired. □

11.4. Local abelianity

Recall that every 1-based stable group is abelian-by-finite (see [10, Theorem 3.2]). This is shown by applying the 'few subgroups' result in that context to the family of conjugation maps.

In the topological setting, the analogous statement is that G contains an open abelian subgroup – or equivalently, μ is abelian. We now prove this, using essentially the same strategy as in the stable case.

Theorem 11.14. *Assume G is topologically 1-based.*

(1) *There are a countable parameter set A, and an A-type-definable open abelian subgroup $H \leq G$.*

(2) *In particular, μ is abelian.*

Proof. The family of (graphs of) conjugation maps $x \mapsto axa^{-1}$ is a $\emptyset$-type-definable family of subgroups of G^2. Let $b \in G$ with $\dim(b) = 1$. Then by Theorem 11.10, the germ of $x \mapsto yxy^{-1}$ at the identity is constant for $y \in \mu(b/\emptyset)$.

Let H be the *local centre* of G – the set of elements $a \in G$ so that $x \mapsto axa^{-1}$ is the identity on a neighbourhood of e. Then H is a $\emptyset$-type-definable subgroup of G. Since any two elements of $\mu(b/\emptyset)$ induce the same conjugation on a neighbourhood of e, it follows that $\mu(b/\emptyset)\mu(b/\emptyset)^{-1} \subseteq H$. In particular, H is infinite, and so $\dim(H) = 1$. Let $h \in H$ with $\dim(h) = 1$, and then let $u \in \mu$ with $\dim(h,u) = 2$ (so $h \in M$ and $u \in N$). So h commutes with an $\mathcal{M}$-definable neighbourhood of the identity, and so h commutes with u. Now since $\mathcal{M}$ is sufficiently saturated, there is $v \in G$ from $\mathcal{M}$ so that $\dim(h,v) = 2$ and h and v commute.

Since $\dim(h,v) = 2$, the $\emptyset$-definable condition 'x and y commute' holds on all (x,y) in some $\mathcal{M}$-definable neighbourhood of (h,v). In particular, there are $\mathcal{M}$-definable open neighbourhoods, U of h and V of v, so that all elements of U commute with all elements

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

of V. But the centralizer of V is a subgroup, so since it has non-empty interior, it contains an $\mathcal{M}$-definable neighbourhood of e, say U' (see Lemma 11.3). Similarly, the centralizer of U' now contains an $\mathcal{M}$-definable neighbourhood V' of e. So all elements of U' commute with all elements of V'. It follows that μ is abelian, because $\mu \subseteq U' \cap V'$. Thus, we have shown (2).

To show (1), inductively construct a sequence $U_1 \supset U_2 \supset \ldots$ of $\mathcal{M}$-definable open neighbourhoods of e so that $U_1 \subseteq U' \cap V'$ and each $U_{i+1} \cdot U_{i+1}, U_{i+1}^{-1} \subseteq U_i$. Then $\bigcap U_i$ is an open type-definable abelian subgroup defined over a countable set. □

11.5. Summarizing

Finally, we end the paper by giving our main theorem in full. Thus, we now drop all underlying assumptions that were present in this section until now.

Theorem 11.15. *Let $\mathcal{M}$ be a sufficiently saturated t-minimal structure with the independent neighbourhood property. Assume $\mathcal{M}$ is non-trivial and topologically 1-based. Then there are a countable parameter set A, and an A-type-definable abelian group G, such that the following hold:*

(1) *G is open in $\mathcal{M}$.*

(2) *G is a topological group with the topology inherited from M.*

(3) *G is locally linear.*

Proof. By Theorem 10.10, we can satisfy all requirements other than abelianity and local linearity. But local linearity follows by Theorem 11.10; and by Theorem 11.14, after further shrinking we can arrange that G is abelian. □

The following may also be worth noting. We leave the details for the interested reader to verify:

Remark 11.16. Let $\mathcal{G}$ be a t-minimal structure with independent neighbourhoods. Assume that $\mathcal{G}$ expands a group, $(G, \cdot, e)$. Let τ denote the given definable topology on G. By Lemma 10.7, G is definably isomorphic to a group G_d locally topological with respect to the topology of the t-minimal structure $\mathcal{G}$. By (the proof of) Lemma 11.1, the intersection μ_d of all G-definable neighbourhoods of d (the identity element of G_d) is a type-definable topological subgroup of G_d. A short argument shows that μ_d is invariant under conjugation by any element of G_d (in the same model over which μ_d was defined). By compactness, this implies that conjugation by $g \in G_d$ is continuous at d (for g in any model) – see e.g. the proof of [8, Proposition 5.17(3)]. By [8, Lemma 2.11], there is an (automatically uniform) definable group topology τ_d on G_d extending the topology on μ_d. Since $\mathcal{G}$ is t-minimal, we get that G_d with τ_d is visceral. Indeed, any infinite definable set S contains a τ-open set U, and for any generic $d' \in U$, we get that $\mu_{d'} \subseteq U$ is open and definably τ_d-homeomorphic to μ_d, so τ_d open. So S has non-empty τ_d-interior. By the main result of Appendix B, this implies that (G_d, τ_d) has the independent neighbourhood property, so by Lemma 3.18, it coincides with τ away from finitely many points. Since G_d is definably isomorphic to G, this isomorphism induces a definable group topology

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

ρ on G, and by the same argument, ρ and τ coincide outside finitely many points. In particular, a topological t-minimal group admits exactly one t-minimal group topology.

Appendix A. Proving the independent neighbourhood property

Here we include a brief appendix to demonstrate that various well-known structures have the independent neighbourhood property. It is possible that the facts we present here were already known. In this case, our contribution is to give a unified treatment.

Assumption A.1. Throughout the appendix, $\mathcal{M}$ is a sufficiently saturated t-minimal structure. We do not assume $\mathcal{M}$ has the independent neighbourhood property.

A.1. Broadness at a point

In his paper, Johnson defines the class of *broad* subsets of M^n. Here we give a similar notion but applying locally near one point. This is not quite a direct translation, because Johnson's notion of broadness is invariant under coordinate permutations and ours is not. Nevertheless, our notion will better suit the argument at hand. So we define what it means for a set to be broad at a point a:

Definition A.2. For each $a \in M$, we inductively define the class of *a-broad* subsets of M^n as follows:

- If $n = 1$, then $B \subseteq M$ is a-broad if a is an accumulation point of B, i.e. every neighbourhood of a contains a point of $B - \{a\}$.
- If $n \geq 1$, then $B \subseteq M^{n+1} = M^n \times M$ is a-broad if there is an a-broad set $C \subseteq M$ such that for each $c \in C$, the fibre $B_c \subseteq M^n$ is a-broad.

Remark A.3. By induction, note that a-broad sets are infinite, and that the assertion 'X is a-broad' is $\emptyset$-definable in families (one only needs the $\emptyset$-definability of the topology).

The main elementary fact we need about a-broad sets is the following:

Lemma A.4. *Let $a \in M$, and let $\{B_i : i \in I\}$ be a small collection of subsets of M^n. If $B = \bigcup_{i \in I} B_i$ is a-broad, then some B_i is a-broad.*

Proof. We induct on n. If $n = 1$, this is just Fact 2.1: If no B_i is a-broad, we can find open U_i containing a and disjoint from B_i. Then $\bigcap U_i$ is open and disjoint from B, so B is not a-broad, a contradiction.

Now suppose $n = m + 1$ where $m \geq 1$. Since B is a-broad, there is an a-broad set $C \subseteq M$ so that for all $c \in C$, the fibre $B_c \subseteq M^m$ is a-broad. For each $i \in I$, let C_i be the set of $c \in C$ so that the fibre $(B_i)_c$ is a-broad. For each $c \in C$, the fibres $(B_i)_c$ union to the a-broad set B_c, so by induction, one of the $(B_i)_c$ is a-broad. Thus, $\bigcup C_i = C$, and so by the base case, some C_{i_0} is a-broad. Then by definition, B_{i_0} is a-broad. □

A.2. Algebraicity over broad sets, Part 1

There are many examples of t-minimal theories admitting locally constant definable functions $f : M \to M$ with infinite image (indeed, every failure of exchange in a t-minimal

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

theory looks something like this). However, no such function can be a choice function: that is, one cannot hope that the constant value taken on an open set $U \subseteq M$ always belongs to U (otherwise, the image of f would be infinite and discrete, which is impossible). Lemma A.5 below captures this idea on a general level:

Lemma A.5. *Let $a \in M$, let A be a parameter set and let $B \subseteq M^n$ be a-broad. If $a \in \mathrm{acl}(Ab)$ for all $b \in B$, then $a \in \mathrm{acl}(A)$.*

Proof. We induct on n. The inductive step is straightforward: Assume $n = m+1$ for some $m \geq 1$, and further assume the statement is true for all values less than n. Let $C \subseteq M$ be a-broad so that B_c is a-broad for each $c \in C$. For each $c \in C$, applying the inductive hypothesis to B_c over the parameter set Ac gives $a \in \mathrm{acl}(Ac)$. Now applying the base case to C gives that $a \in \mathrm{acl}(A)$.

The bulk of the work is the case $n = 1$. So assume $B \subseteq M$ is a-broad and $a \in \mathrm{acl}(Ab)$ for all $b \in B$. By Lemma A.4, this algebraicity is witnessed uniformly on a broad set. That is, after potentially replacing B with a smaller (still a-broad) subset, we can find an A-definable $X \subseteq M^2$ with the following properties:

(1) $(a,b) \in X$ for all $b \in B$.

(2) The second projection $\pi_2 : M^2 \to M$ is at most k-to-1 on X for some integer k.

Let $\pi_1 : M^2 \to M$ be the first projection. Let Y be the set of $y \in M$ so that the fibre $\pi_1^{-1}(y) \cap X$ is y-broad. Note that Y is A-definable and contains a (as $\pi^{-1}(a) \supseteq B$). If a is a boundary point of Y, then by t-minimality, $a \in \mathrm{acl}(A)$ (Remark 2.3), and we are done. We show that a cannot be an interior point of Y. Assume otherwise, and let U be an open neighbourhood of a with $U \subseteq Y$.

We now inductively construct elements $a_1, a_2, \cdots \in U$ and non-empty open sets $U = U_0 \supset U_1 \supset U_2 \ldots$ with the following properties:

(1) $a_i \in U_j$ whenever $i > j$.

(2) Each $\{a_i\} \times U_i \subseteq X$ for $i \geq 1$.

Suppose $i \geq 1$ and we have built all previous a_j and U_j. First, let a_i be any element of U_{i-1}, noting that this will ensure (1) is satisfied.

Next, by definition, $a_i \in U_0 \subseteq Y$, so the fibre $F_i = \pi_1^{-1}(a_i) \cap X$ is a_i-broad. In particular, $F_i \cap U_{i-1}$ is infinite, so has non-empty interior by t-minimality. Now let U_i be any non-empty open subseteq of $F_i \cap U_{i-1}$, and note that this ensures we satisfy (2).

Finally, assume we have built a_i and U_i as above for all i. Then by (1) and (2), we have:

(3) $(a_i, a_j) \in X$ whenever $i < j$.

Indeed, if $i < j$, then (1) gives $a_j \in U_i$, and so (2) gives $(a_i, a_j) \in X$. In particular, (3) implies that the fibres $\pi_2^{-1}(a_j) \cap X$ become arbitrarily large as $j \to \infty$, which contradicts that π_2 is k-to-1 on X. □

A.3. Algebraicity over broad sets, Part 2

Next, we consider a counterpart to Lemma A.5 where we consider two points. That is, suppose B is a-broad, and $c \in \mathrm{acl}(Ab)$ for all $b \in B$. Does it follow that $c \in \mathrm{acl}(Aa)$? In

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

fact, we do not know how to prove this, but we also do not know of a t-minimal theory where it fails. Instead, we make it a definition:

Definition A.6. $\mathcal{M}$ has the *few broad sets property* if the following holds: Let $X \subseteq M^2$ be definable, $\pi_1,\pi_2 : M^2 \to M$ the projections onto the first and second factors, respectively. Assume that π_2 is finite-to-one on X. Then for each $a \in M$, there are only finitely many $b \in M$ so that the fibre $\{y : (b,y) \in X\}$ is a-broad.

We give some examples:

Example A.7. Suppose $\mathcal{M}$ satisfies exchange. Then $\mathcal{M}$ has the few broad sets property. Indeed, let X and a be as in Definition A.6. If infinitely many fibres in X under π_1 are a-broad, then in particular, infinitely many fibres in X under π_1 are infinite. By exchange, this implies $\dim(X) = 2$, which contradicts that π_2 is finite-to-one on X.

Example A.8. Say that $\mathcal{M}$ has the *few germs property* if for all $a \in M$, there are only a small (i.e. bounded) number of definable germs at a. For example:

- Every weakly o-minimal theory has the few germs property. Indeed, in this case, there are only four possibilities for the germ of a definable set $X \subseteq M$ at $a \in X$: (i) the single point a; (ii) an interval $[a,b)$ with $b > a$; (iii) an interval $(b,a]$ with $b < a$; or (iv) an interval (b,c) with $b < a < c$.
- Every C-minimal theory has the few germs property. Indeed, suppose $\mathcal{M}$ is C-minimal. Then given a definable $X \subseteq M$ and $a \in X$, either a is isolated in X or X contains a ball around a. Thus, there are only two definable germs at a.

Now we show that if $\mathcal{M}$ has the few germs property, then it has the few broad sets property. Indeed, let $X \subseteq M^2$ and a be as in Definition A.6. For each $y \in M$, let $B_y = \{a\} \cup \{z : (y,z) \in X\}$, i.e. the fibre above y in X under the first projection (with a added to ensure that $\mathrm{germ}_a(B_y)$ is well-defined; note that adding or removing a does not affect a-broadness).

Let Y be the (definable) set of $y \in M$ so that B_y is a-broad. Towards a contradiction, assume Y is infinite. Then Y is not small, so by the few germs property, there is an infinite $Z \subseteq Y$ so that $\mathrm{germ}_a(B_y)$ is constant for $y \in Z$. Thus, any finitely many B_y for $y \in Z$ have infinite intersection (here we use that all such B_y are a-broad). In particular, for any positive integer k, we can find sets $S_k,T_k \subseteq M$ with $|S_k| = |T_k| = k$ and $S_k \times T_k \subseteq X$. By compactness, we can find infinite sets $S,T \subseteq M$ with $S \times T \subseteq X$. This contradicts that π_2 is finite-to-one on X.

Example A.9. For those well-versed in invariant types, one can generalize the previous example as follows: Say that a global 1-type is a-broad if every formula in it defines an a-broad set. Then say $\mathcal{M}$ is *locally invariant* if for each $a \in M$, every a-broad global 1-type is a-invariant (in fact, it is enough to find a small set B so that every a-broad 1-type is Ba-invariant). Weakly o-minimal and C-minimal theories are easily seen to be locally invariant. Now suppose $\mathcal{M}$ is locally invariant; we show $\mathcal{M}$ has the few broad sets property. Let X and a be as in Definition A.6, where X is A-definable. In the case that every a-broad type is only Ba-invariant for a fixed B, we include B in A; in any case, every a-broad type is now Aa-invariant.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

Now assuming that infinitely many π_1-fibres in X are a-broad, one finds $b \in M - \mathrm{acl}(Aa)$ so that $\pi_1^{-1}(b) \cap X$ is a-broad. A compactness argument (using Lemma A.4) produces a global a-broad type p inside the fibre $\pi^{-1}(b) \cap X$. By Aa-invariance, and the fact that $b \notin \mathrm{acl}(Aa)$, any realization $x \models p$ satisfies $x \times T \subseteq X$, where T is the infinite set of realizations of $\mathrm{tp}(b/Aa)$ in $\mathcal{M}$. But since p is a-broad, it is not algebraic. Let S be an infinite set of realizations of p in an elementary extension. Then $S \times T \subseteq X$, and we again contradict that π_2 is finite-to-one on X.

Now using a similar argument to Lemma A.5, we show:

Lemma A.10. *Assume $\mathcal{M}$ has the few broad sets property. Let $a,c \in M$, A a parameter set and $B \subseteq M^n$ an a-broad set. If $c \in \mathrm{acl}(Ab)$ for all $b \in B$, then $c \in \mathrm{acl}(Aa)$.*

Proof. Similar to Lemma A.5. We induct on n, the inductive step being essentially identical to that of Lemma A.5. For the base case, assume $n = 1$. As in Lemma A.5, we can reduce to a single A-definable set $X \subseteq M^2$ witnessing all algebraicities. So if π_1, π_2 are the two projections, then π_2 is finite-to-one on X and $\pi_1^{-1}(c) \cap X$ is a-broad. By the few broad sets property, c is one of only finitely many points with a-broad fibre in X. Thus, $c \in \mathrm{acl}(Aa)$. □

A.4. Independent parameters from a broad set

Assume $\mathcal{M}$ has the few broad sets property. We now prove that, given a tuple and a broad set at any of its coordinates, we can choose parameters from the broad sets without decreasing the dimension of the whole tuple. In many cases, this is sufficient to prove the independent neighbourhood property.

We begin by proving the result at a single coordinate. We then conclude the full version inductively.

Lemma A.11. *Assume $\mathcal{M}$ has the few broad sets property. Let $a = (a_1, \ldots, a_n) \in M^n$, let A be a parameter set and let $B \subseteq M^m$ be a_i-broad for some i. Then there is $b \in B$ so that $\dim(a/Ab) = \dim(a/A)$.*

Proof. Fix i such that B is a_i-broad. Let e be a basis for a over A. Choose e so that if a_i belongs to any basis at all, then it belongs to e.

Assume $\dim(a/Ab) < \dim(a/A)$ for each $b \in B$. Then for each $b \in B$, there is a coordinate $x \in e$ so that $x \in \mathrm{acl}(Ab \cup (e - \{x\}))$. By Lemma A.4, after shrinking B, we can assume there is a single coordinate $x \in e$ so that $x \in \mathrm{acl}(Ab \cup (e - \{x\}))$ for all $b \in B$. We have two cases:

- Case 1: $x = a_i$. Then applying Lemma A.5 over the parameter set $A \cup (e - \{x\})$, we conclude that $a_i \in A \cup (e - \{a_i\})$. This contradicts that e is a basis for a over A.
- Case 2: $x \neq a_i$. Then applying Lemma A.10 over the parameter set $A \cup (e - \{x\})$, we conclude that $x \in \mathrm{acl}(Aa_i \cup (e - \{x\}))$. In particular, let e' be e with x replaced by a_i; then $e \in \mathrm{acl}(Ae')$, and thus $a \in \mathrm{acl}(Ae')$. Since e and e' have the same length, e' is also a basis for a over A. Then by the choice of e, $a_i \in e$. Thus, $Aa_i \cup (e - \{x\})$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

is just $A \cup (e - \{x\})$. So $x \in \mathrm{acl}(A \cup (e - \{x\}))$, and this again contradicts that e is a basis for a over A. □

Finally, we conclude:

Theorem A.12. *Assume $\mathcal{M}$ has the few broad types property. Fix $a = (a_1, \ldots, a_n) \in M^n$ and a parameter set A. For each i, let B_i be an a_i-broad subseteq of M^{m_i}. Then there is $b = (b_1, \ldots, b_n) \in B_1 \times \cdots \times B_n$ so that $\dim(a/Ab) = \dim(a/A)$.*

Proof. Choose the b_i inductively, at each stage working over A together with all previous b_j's. □

A.5. Concluding

Finally, we connect Theorem A.12 to the topology:

Definition A.13. Say that $\mathcal{M}$ is *locally broad* if the following holds: for every $a \in M$ and every open neighbourhood U of a, there are $m \geq 1$ and an a-broad set $B \subseteq M^m$ so that for every $b \in B$, there is a b-definable open set V_b with $a \in V_b \subseteq U$.

Roughly speaking, Definition A.13 says that one can define small neighbourhoods of $a \in M$ using parameters near a – and in addition, a-broad many perturbations of those parameters should not move the neighbourhood too much (or move a out of it). Weakly o-minimal and C-minimal theories easily have this property, as do most natural examples of t-minimal theories. For example, if $\mathcal{M}$ is weakly o-minimal and (b,c) is a neighbourhood of a, then one can define subneighbourhoods using parameters x and y whenever $x \in (b,a)$ and $y \in (a,c)$ (and the set of such pairs (x,y) is a-broad). The C-minimal case follows by a similar argument.

On the other hand, a real closed field with the Sorgenfrey topology is *not* locally broad: for example, given a and a neighbourhood $[a,b)$, the set of (x,y) with $a \in [x,y) \subset [a,b)$ is $\{a\} \times (a,b]$, which is not a-broad.

Now we show:

Theorem A.14. *Suppose $\mathcal{M}$ is locally broad and has the few broad sets property. Then $\mathcal{M}$ has the independent neighbourhood property.*

Proof. Combine the definition of local broadness with Theorem A.12. □

Corollary A.15. *If $\mathcal{M}$ is weakly o-minimal, or if $\mathcal{M}$ is C-minimal, then $\mathcal{M}$ has the independent neighbourhood property.*

Corollary A.16. *If $\mathcal{M}$ is locally broad and satisfies exchange, then $\mathcal{M}$ has the independent neighbourhood property.*

Appendix B. Independent neighbourhoods in visceral theories, by Will Johnson

We show that the independent neighbourhood property (Definition 3.1) holds in any visceral theory. We remind that a t-minimal theory is *visceral* if it admits a definable

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

uniform structure (see, e.g. [11, Definition 1.17]). So a dp-minimal visceral theory is an SW-uniformity (in the sense of [20]). For the latter, the independent neighbourhood property was proved in [8, Corollary 3.12].

Work in a monster model $\mathbb{M}$ of a visceral theory T in a language $\mathcal{L}$. Let $\mathcal{B}$ be a definable basis for the uniformity, chosen so that basic entourages are symmetric and open. If E is a basic entourage, and $a \in \mathbb{M}$, then $E[a]$ denotes the 'open ball' $\{x \in \mathbb{M} : (a,x) \in E\}$. The family $\{E[a] : E \in \mathcal{B}\}$ is a neighbourhood basis of a.

Remark B.1. The arguments below generalize to the case of a definable set D with a visceral uniformity, not necessarily definable in the induced structure on D.

Lemma B.2. *For any small cardinal κ, there is a small cardinal λ such that if $S_1,\dots,S_n \subseteq \mathbb{M}$ are sets of size at least λ and $A \subseteq \mathbb{M}$ is a set of size at most κ, then there is $\bar{b} \in \prod_{i=1}^n S_i$ such that $\mathrm{tp}(\bar{b}/A)$ is broad.*

Proof. Fix κ as in the statement. By the Erdős-Rado theorem, choose λ such that if we colour the n-element subsets of λ with $2^{|\mathcal{L}|+\kappa}$-many colours, then there is a homogeneous set of size $\aleph_0$.

Now let $S_1,\dots,S_n,A \subseteq \mathbb{M}$ be given. Shrinking the S_i, we may assume $|S_i| = \lambda$ for each i. Let $\{c_{i,\alpha} : \alpha < \lambda\}$ be an enumeration of S_i for each i. Consider the map from n-element subsets of λ to the type space over A given by

$$\{\alpha_1,\dots,\alpha_n\} \mapsto \mathrm{tp}(c_{1,\alpha_1},\dots,c_{n,\alpha_n}/A)$$

for $\alpha_1 < \alpha_2 < \cdots < \alpha_n$. Let $J \subseteq \lambda$ be an infinite subseteq which is homogeneous of colour p, so that

$$\mathrm{tp}(c_{1,\alpha_1},\dots,c_{n,\alpha_n}/A) = p$$

for all $\alpha_1 < \cdots < \alpha_n$ in J. Note that $(c_{1,\alpha_1},\dots,c_{n,\alpha_n}) \in \prod_{i=1}^n S_i$, so it remains to show that p is broad, i.e. the set $p(\mathbb{M})$ of realizations is broad. Given any k, we can find an increasing sequence

$$\beta_{1,1} < \beta_{1,2} < \cdots < \beta_{1,k} < \beta_{2,1} < \cdots < \beta_{2,k} < \cdots < \beta_{n,1} < \cdots < \beta_{n,k}$$

in J, because J is infinite. Then every element of

$$\begin{aligned}&\prod_{i=1}^n \{c_{i,\beta_{i,1}}, c_{i,\beta_{i,2}},\dots,c_{i,\beta_{i,k}}\}\\ &= \{c_{1,\beta_{1,1}}, c_{1,\beta_{1,2}},\dots,c_{1,\beta_{1,k}}\} \times \cdots \times \{c_{n,\beta_{n,1}}, c_{n,\beta_{n,2}},\dots,c_{n,\beta_{n,k}}\}\end{aligned}$$

has type p. As k was arbitrary, $p(\mathbb{M})$ is broad. □

Lemma B.3. *Let $(P,\leq)$ be an A-definable preorder. Suppose $(P,\leq)$ is downward directed, in the sense that for any $x,y \in P$, there is $z \in P$ with $z \leq x$ and $z \leq y$. If $\mathrm{tp}(\bar{b}/A)$ is broad and $c \in P$, there is $c' \leq c$ such that $\mathrm{tp}(\bar{b}/Ac')$ is broad.*

Proof. Let $\kappa = |A| + \aleph_0$. Let λ be as in Lemma B.2. Let $p = \mathrm{tp}(\bar{b}/A)$. Since p is broad, the set $p(\mathbb{M})$ contains a product $S := \prod_{i=1}^n S_i$, with $|S_i| = \lambda$ for each i. Note that $\bar{e} \equiv_A \bar{b}$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

for every $\bar{e} \in S$. By the axiom of choice, there is a function $f : S \to P$ such that

$$(\bar{e}, f(\bar{e})) \equiv_A (\bar{b}, c) \text{ for all } \bar{e} \in S.$$

By saturation and directedness, there is some $c_0 \in P$ such that $c_0 \leq f(\bar{e})$ for every $\bar{e} \in S$. By Lemma B.2, there is some $\bar{e} \in S$ such that $\mathrm{tp}(\bar{e}/Ac_0)$ is broad. Take $c' \in P$ such that

$$(\bar{e}, f(\bar{e}), c_0) \equiv_A (\bar{b}, c, c').$$

The fact that $c_0 \leq f(\bar{e})$ implies $c' \leq c$. The fact that $\mathrm{tp}(\bar{e}/Ac_0)$ is broad implies that $\mathrm{tp}(\bar{b}/Ac')$ is broad. □

Corollary B.4. *Let $(P, \leq)$ be an A-definable downward-directed preorder. For any $\bar{b} \in \mathbb{M}^n$ and any $c \in P$, there is $c' \leq c$ such that $\dim(\bar{b}/A) = \dim(\bar{b}/Ac')$.*

Proof. Let $\bar{e}$ be a subtuple of $\bar{b}$ which is acl-independent over A, and interalgebraic with $\bar{b}$ over A, so

$$\dim(\bar{b}/A) = |\bar{e}|.$$

Moreover, $\mathrm{tp}(\bar{e}/A)$ is broad. Take $c' \leq c$ such that $\mathrm{tp}(\bar{e}/Ac')$ is broad. Then $\bar{e}$ is acl-independent over Ac', and interalgebraic with $\bar{b}$ over Ac', and so

$$\dim(\bar{b}/Ac') = |\bar{e}| = \dim(\bar{b}/A).$$

□

Lemma B.5. *Let $p = \mathrm{tp}(\bar{a}/C)$, and let $p(\mathbb{M})$ be the set of realizations. Let $k = \dim(\bar{a}/C) = \dim(p) = \dim(p(\mathbb{M}))$. If $U \ni \bar{a}$ is any neighbourhood, then $\dim(p(\mathbb{M}) \cap U) = k$.*

Proof. Otherwise, there is a C-definable set $D \supseteq p(\mathbb{M})$ such that $\dim(D \cap U) < k$ by [11, Corollary 2.42]. Then $\dim_{\bar{a}}(D) < k$, where $\dim_{\bar{a}}(D)$ is the local dimension of D at $\bar{a}$. Let $X = \{\bar{b} \in D : \dim_{\bar{b}}(D) < k\} \subseteq D$. The set X is C-definable and contains $\bar{a}$. If $\bar{b} \in X$, then $\dim_{\bar{b}}(X) \leq \dim_{\bar{b}}(D) < k$. So the local dimension of X is at most $k-1$ at any point. By [11, Proposition 3.10],

$$\dim(X) = \max_{\bar{b} \in X} \dim_{\bar{b}}(X) \leq k - 1.$$

Then $\bar{a}$ belongs to the $(k-1)$-dimensional C-definable set X, contradicting the fact that $\dim(\bar{a}/C) = k$. □

Lemma B.6. *Let C be a set. Let E be a C-definable basic entourage. For any $(\bar{a}, \bar{b})$, there is $\bar{c} \in E[\bar{b}] \subseteq \mathbb{M}^n$ such that $\dim(\bar{a}, \bar{b}, \bar{c}/C) = \dim(\bar{a}, \bar{b}/C) + \dim(\bar{c}/C)$.*

Proof. Take E_1 a basic entourage such that $E_1 \circ E_1 \subseteq E$. Let $C_1 \supseteq C$ be a small set defining E_1. Let $k = \dim(\bar{a}, \bar{b}/C)$. Let $n = |\bar{b}| = \dim E[\bar{b}] = \dim E_1[\bar{b}]$. Let X be the type-definable set of pairs $(\bar{a}', \bar{b}')$ such that $(\bar{a}', \bar{b}') \equiv_C (\bar{a}, \bar{b})$ and $\bar{b}' \in E_1[\bar{b}]$. By Lemma B.5, $\dim(X) = \dim(\bar{a}, \bar{b}/C) = k$. Then $\dim(X \times E_1[\bar{b}]) = k + n$. Since the set $X \times E_1[\bar{b}]$ is type-definable over $C_1\bar{b}$, there are $(\bar{a}', \bar{b}', \bar{c}') \in X \times E_1[\bar{b}]$ with $\dim(\bar{a}', \bar{b}', \bar{c}'/C_1\bar{b}) = k + n$. Then

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

$\mathrm{tp}(\bar{a}',\bar{b}'/C) = \mathrm{tp}(\bar{a},\bar{b}/C)$, so $\dim(\bar{a}',\bar{b}'/C) = k$. And $\dim(\bar{c}'/C) \leq n$ because $\bar{c}'$ is an n-tuple. Then

$$\begin{aligned} k+n = \dim(\bar{a}',\bar{b}',\bar{c}'/C_1\bar{b}) &\leq \dim(\bar{a}',\bar{b}',\bar{c}'/C) \\ &\leq \dim(\bar{a}',\bar{b}'/C) + \dim(\bar{c}'/C) \leq k+n. \end{aligned}$$

Equality must hold, so $\dim(\bar{a}',\bar{b}',\bar{c}'/C) = \dim(\bar{a}',\bar{b}'/C) + \dim(\bar{c}'/C)$. By choice of $(\bar{a}',\bar{b}',\bar{c}') \in X \times E_1[\bar{b}]$, we have $\bar{b}',\bar{c}' \in E_1[\bar{b}]$. By choice of E_1, it follows that $\bar{c}' \in E[\bar{b}']$. To summarize, we have $\bar{a}',\bar{b}',\bar{c}'$ such that

$$\begin{aligned} \bar{a}'\bar{b}' &\equiv_C \bar{a}\bar{b} \\ \bar{c}' &\in E[\bar{b}'] \\ \dim(\bar{a}',\bar{b}',\bar{c}'/C) &= \dim(\bar{a}',\bar{b}'/C) + \dim(\bar{c}'/C). \end{aligned}$$

Taking $\bar{c}$ such that $(\bar{a}',\bar{b}',\bar{c}') \equiv_C (\bar{a},\bar{b},\bar{c})$, and moving by an automorphism, we get the desired configuration. □

One might conjecture the following cleaner form of Lemma B.6:

Non-Theorem B.7. *Given a type* $\mathrm{tp}(\bar{a}/C)$ *and a non-empty open set* $U \subseteq \mathbb{M}^n$*, there is* $\bar{b} \in U$ *with* $\dim(\bar{a},\bar{b}/C) = \dim(\bar{a}/C) + \dim(\bar{b}/C)$.

If the exchange property holds, this is true: Take $\bar{b} \in U$ with $\dim(\bar{b}/C,\ulcorner U \urcorner,\bar{a}) = n$. Then $\dim(\bar{b}/C) = \dim(\bar{b}/C\bar{a}) = n$, so $\dim(\bar{a},\bar{b}/C) = \dim(\bar{a}/C) + \dim(\bar{b}/C)$ by additivity.

On the other hand, if the exchange property fails, so does Non-Theorem B.7. Specifically, take $a,b_0 \in \mathbb{M} \setminus \mathrm{acl}(C)$ with $a \in \mathrm{acl}(Cb_0)$ and $b_0 \notin \mathrm{acl}(Ca)$. Take $\phi(x,y) \in \mathrm{tp}(a,b_0/C)$ algebraizing a over b_0. The boundary $\mathrm{bd}(\phi(a,\mathbb{M}))$ is finite and Ca-definable, so it does not contain b_0, and instead $b_0 \in \mathrm{int}(\phi(a,\mathbb{M})) =: U$. If $b \in U$, then a is in the finite set $\phi(\mathbb{M},b)$, so $a \in \mathrm{acl}(Cb)$ and

$$\dim(a,b/C) = \dim(b/C) < \dim(a/C) + \dim(b/C).$$

So we cannot find any $b \in U$ which is 'dimension independent' from a over C.

Theorem B.8. *Let* $(\bar{a},\bar{b})$ *be a finite tuple in* $\mathbb{M}$*. Let* C *be a small set of parameters. Let* U *be a neighbourhood of* $\bar{b}$*. Then there is* $C' \supseteq C$ *such that* $\dim(\bar{a},\bar{b}/C') = \dim(\bar{a},\bar{b}/C)$*, and a* C'*-definable neighbourhood* $U' \ni \bar{b}$*, with* $U' \subseteq U$*.*

Proof. Let $\{E_p : p \in P\}$ be the definable basis for the uniformity. Regard P as a downward-directed preorder by defining

$$p \leq q \iff E_p \subseteq E_q.$$

Let p_0 be so small that $E_{p_0}^2[\bar{b}] \subseteq U$. By Corollary B.4, there is $p \leq p_0$ such that $\dim(\bar{a},\bar{b}/Cp) = \dim(\bar{a},\bar{b}/C)$. By Lemma B.6, there is $\bar{c} \in E_p[\bar{b}]$ such that

$$\dim(\bar{a},\bar{b},\bar{c}/Cp) = \dim(\bar{a},\bar{b}/Cp) + \dim(\bar{c}/Cp).$$

By sub-additivity of dimension [11, Proposition 2.31(5)],

$$\dim(\bar{a},\bar{b},\bar{c}/Cp) \leq \dim(\bar{a},\bar{b}/Cp\bar{c}) + \dim(\bar{c}/Cp) \leq \dim(\bar{a},\bar{b}/Cp) + \dim(\bar{c}/Cp),$$

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

and so equality holds and

$$\dim(\bar{a},\bar{b}/Cp\bar{c}) = \dim(\bar{a},\bar{b}/Cp) = \dim(\bar{a},\bar{b}/C).$$

Because $E_p^2[\bar{b}] \subseteq E_{p_0}^2[\bar{b}] \subseteq U$ and $\bar{c} \in E_p[\bar{b}]$, we have $\bar{b} \in E_p[\bar{c}] \subseteq U$. Take $U' = E_p[\bar{c}]$ and $C' = Cp\bar{c}$. Then U' is C'-definable and $\dim(\bar{a},\bar{b}/C') = \dim(\bar{a},\bar{b}/C)$. □

Theorem B.8 is the independent neighbourhood property, or perhaps something slightly stronger on account of the additional parameter $\bar{a}$.

Competing interests. None of the authors have competing interests.

Data availability statement. There is no data associated with this paper.

References

[1] Berenstein A and Vassiliev E (2012) Weakly one-based geometric theories. *J. Symbolic Logic* **77**(2), 392–422.
[2] Castle B and Hasson A (October 2024) Strongly Minimal Relics of T-convex Fields. *arXiv e-prints*, page arXiv:2410.22442.
[3] Castle B, Hasson A and Ye J (May 2024) Zilber's Trichotomy in Hausdorff Geometric Structures. *arXiv e-prints*, page arXiv:2405.02209.
[4] Cherlin G, Harrington L and Lachlan AH (1985) $\aleph_0$-categorical, $\aleph_0$-stable structures. *Ann. Pure Appl. Logic* **28**(2), 103–135.
[5] Chernikov A, Peterzil Y and Starchenko S (2024) Model-theoretic Elekes-Szabó for stable and o-minimal hypergraphs. *Duke Math. J.* **173**(3), 419–512.
[6] Gagelman J (2005) Stability in geometric theories. *Ann. Pure Appl. Logic* **132**(2–3), 313–326.
[7] Halevi Y, Hasson A and Peterzil Y (2022) Interpretable fields in various valued fields. *Adv. Math.* **404**, Paper No. 108408.
[8] Halevi Y, Hasson A and Peterzil Y (2024) On groups interpretable in various valued fields. *Selecta Math. (N.S.)* **30**(4), Paper No. 59.
[9] Hrushovski E (1987) Locally modular regular types. In: *Classification Theory (Chicago, IL, 1985)*, Lecture Notes in Math. 1292, 132–164. Berlin: Springer.
[10] Hrushovski U and Pillay A (1987) Weakly normal groups. In: *Logic Colloquium '85 (Orsay, 1985)*, Stud. Logic Found. Math. 122, 233–244. Amsterdam: North-Holland.
[11] Johnson W (April 2024) Visceral theories without assumptions. *arXiv e-prints*, page arXiv:2404.11453.
[12] Loveys J and Peterzil Y (1993) Linear o-minimal structures. *Israel J. Math.* **81**(1–2), 1–30.
[13] Maalouf F (2008) Construction d'un groupe dans les structures C-minimales. *J. Symbolic Logic* **73**(3), 957–968.
[14] Mathews L (1995) Cell decomposition and dimension functions in first-order topological structures. *Proc. London Math. Soc. (3)* **70**(1), 1–32.
[15] Maříková J (2007) Type-definable and invariant groups in o-minimal structures. *J. Symbolic Logic* **72**(1), 67–80.
[16] Peterzil Y (1994) Constructing a group-interval in o-minimal structures. *J. Pure Appl. Algebra* **94**(1), 85–100.
[17] Peterzil Y (September 2019) An o-minimalist view of the group configuration. *arXiv e-prints*, page arXiv:1909.09994.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press

[18] Peterzil Y and Starchenko S (1998) A trichotomy theorem for o-minimal structures. *Proc. London Math. Soc. (3)* **77**(3), 481–523.
[19] Pillay A (1996) *Geometric Stability Theory*. Oxford Logic Guides, vol. 32. New York: The Clarendon Press, Oxford University Press. Oxford Science Publications.
[20] Simon P and Walsberg E (2019) Tame topology over dp-minimal structures. *Notre Dame J. Form. Log.* **60**(1), 61–76.
[21] Zilber B (1993) *Uncountably Categorical Theories*. Translations of Mathematical Monographs, vol. 117. Providence, RI: American Mathematical Society. Translated from the Russian by D. Louvish.

https://doi.org/10.1017/S1474748026102035 Published online by Cambridge University Press