The Lattice of Definable Equivalence Relations in Homogeneous n-dimensional Permutation Structures

Samuel Braunfeld

Department of Mathematics
Rutgers University
New Jersey, U.S.A.

swb52@math.rutgers.edu

Submitted: Mar 2, 2016; Accepted: Dec 7, 2016; Published: Dec 23, 2016
Mathematics Subject Classifications: 03C50, 03C13

Abstract

In *Homogeneous permutations*, Peter Cameron [Electronic Journal of Combinatorics 2002] classified the homogeneous permutations (homogeneous structures with 2 linear orders), and posed the problem of classifying the homogeneous n-dimensional permutation structures (homogeneous structures with n linear orders) for all finite n. We prove here that the lattice of $\emptyset$-definable equivalence relations in such a structure can be any finite distributive lattice, providing many new imprimitive examples of homogeneous finite dimensional permutation structures. We conjecture that the distributivity of the lattice of $\emptyset$-definable equivalence relations is necessary, and prove this under the assumption that the reduct of the structure to the language of $\emptyset$-definable equivalence relations is homogeneous. Finally, we conjecture a classification of the primitive examples, and confirm this in the special case where all minimal forbidden structures have order 2.

Keywords: countable homogeneous; Fraisse theory; infinite permutations

1 Introduction

In [1], Cameron classified the homogeneous permutations, which he identified with homogeneous structures consisting of two linear orders. He then posed the problem of classifying the homogeneous structures consisting of n linear orders for any n [1, §6, Problem 1], which we call n-dimensional permutation structures. The first step toward a classification is the production of a catalog, or census, of examples occurring "in nature." We begin that project here.

The homogeneous permutations are of three kinds. The main distinction is between imprimitive (where there is a non-trivial definable equivalence relation) and primitive, but we subdivide further.

- Imprimitive: Lexicographic orderings (more naturally viewed as equipped with a single order and an equivalence relation)
- Primitive:
  - Degenerate: The orders agree, up to reversal
  - Fully generic in the sense of the Fraïssé theory of amalgamation classes (see Appendix A.1)

The only known primitive homogeneous finite dimensional permutation structures are a straightforward mix of these two types. We note here that Appendix A.1 contains a brief introduction to homogeneity, while Appendix A.2 contains contains information about the amalgamation diagrams that appear throughout this paper.

Conjecture 1 (Primitivity Conjecture)**.** *Every primitive homogeneous finite dimensional permutation structure can be constructed by the following procedure.*

1. *Identify certain orders, up to reversal.*
2. *Take the Fraïssé limit of the resulting amalgamation class, getting a fully generic structure, probably in a simpler language.*

As far as the primitive case is concerned, this is a satisfying catalog of all known examples. As we will show below, the imprimitive case is more complicated. The first question that must be addressed is the following.

Problem. *Describe all possible lattices of $\emptyset$-definable equivalence relations in homogeneous finite dimensional permutation structures.*

We propose the following.

Conjecture 2. *A finite lattice is isomorphic to the lattice of $\emptyset$-definable equivalence relations in some homogeneous finite dimensional permutation structure iff it is distributive.*

In the case of homogeneous permutations ($n = 2$), the only lattices seen are linear of order at most 3. But the following shows the class of possible lattices must be enlarged to include distributive lattices.

Theorem 3.1 (Representation Theorem)**.** *Let Λ be a finite distributive lattice. Then there is a homogeneous finite dimensional permutation structure whose lattice of $\emptyset$-definable equivalence relations is isomorphic to Λ.*

So the question remains, whether distributivity is necessary. In this direction, we have the following.

Theorem 4.6. *Let Λ be the lattice of $\emptyset$-definable equivalence relations in a homogeneous finite dimensional permutation structure $\mathcal{M}$. If the reduct of $\mathcal{M}$ to the language of equivalence relations from Λ is homogeneous, then Λ is distributive.*

A sharper statement may be given in terms of the *infinite index property*, defined as follows.

Definition 4.1. Let $\mathcal{M}$ be a structure with a transitive automorphism group.

1. For $\emptyset$-definable equivalence relations $F \subset E$ on $\mathcal{M}$, set

$$[E : F] = |C/F|$$

for C any E-class, and call this the *index of F in E*.

2. Let Λ be the lattice of $\emptyset$-definable equivalence relations on $\mathcal{M}$. Then Λ has the *infinite index property (IIP)* if whenever $F \subset E$ for $E, F \in \Lambda$, $[E : F]$ is infinite.

Theorem 1.1.

1. *If $\mathcal{M}$ is a homogeneous permutation structure, then the lattice of $\emptyset$-definable equivalence relations on $\mathcal{M}$ has the IIP.*

2. *If $\mathcal{M}$ is a homogeneous structure consisting of a set equipped with a finite lattice Λ of equivalence relations with the IIP, then Λ is distributive.*

The Representation Theorem provides a large variety of new examples of homogeneous permutation structures, but the associated amalgamation classes have a surprisingly simple form: the minimal forbidden substructures all have order 3 or less.

Question. *Can a homogeneous finite dimensional permutation structure have a minimal forbidden substructure of order greater than 3?*

In the last section, we consider homogeneous permutation structures in which all minimal forbidden substructures are of order 2. Such a structure is necessarily primitive, since defining equivalence relations requires forbidding 3-types, and the Primitivity Conjecture predicts its form, which is confirmed in this special case by the following proposition.

Proposition 5.1. *Let $\mathcal{K}$ be an amalgamation class of n-dimensional permutation structures. If no 3-type compatible with the allowed 2-types is forbidden, then the forbidden 2-types collectively specify that certain orders agree up to reversal.*

2 Λ -Ultrametric Spaces

In this section, we set up a language that will be more convenient for our amalgamation arguments than the language of equivalence relations.

Definition 2.1. Let Λ be a lattice. A Λ-*ultrametric space* is a metric space where the metric takes values in Λ and the triangle inequality uses the join rather than addition. Analogous to a pseudometric space, we also define a Λ-*ultrapseudometric space* as a Λ-ultrametric space without the requirement that the metric assign non-zero distance to distinct points.

As in metric spaces, quotienting out the relation $d(x, y) = 0$ in a Λ-ultrapseudometric space yields a Λ-ultrametric space, and we also have a path variant of the the triangle inequality: $d(x, y)$ is no greater than the join of the distances between points on any path from x to y.

Theorem 2.2. *For a given finite lattice Λ, there is an isomorphism between the category of Λ-ultrametric spaces and the category of structures consisting of a set equipped with a family of equivalence relations, closed under taking meets in the lattice of all equivalence relations on the set, and labeled by the elements of Λ in such a way that the map from Λ to the lattice of equivalence relations is a homomorphism. Furthermore, the functors of this isomorphism preserve homogeneity.*

Although we do not prove this theorem here, we will define the functors giving this isomorphism.

Given a system of equivalence relations as specified above, we get the corresponding Λ-ultrametric space by taking the same universe and defining $d(x, y) = \bigwedge \{\lambda \in \Lambda | xE_\lambda y\}$. In the reverse direction, given a Λ-ultrametric space, we get the corresponding structure of equivalence relations by taking the same universe and defining $E_\lambda = \{(x, y)|d(x, y) \leqslant \lambda\}$.

There is a well known amalgamation strategy for metric spaces, and we here give the analog for Λ-ultrametric spaces.

Definition 2.3. Consider an amalgamation diagram of Λ-ultrametric spaces with base B. Let x and y be extension points in different factors, and for each $b_i \in B$ let $d(x, b_i) = e_i$ and $d(y, b_i) = e'_i$. *Pre-canonical amalgamation* is the amalgamation strategy assigning $d(x, y) = \bigwedge_i (e_i \vee e'_i)$. *Canonical amalgamation* is the strategy of pre-canonical amalgamation, followed by identifying x and y if $d(x, y) = \mathbb{0}$.

Two-point pre-canonical amalgamation is shown in Figure 2.1 (for guidance on the interpretation of amalgamation diagrams in this paper, see Appendix A.2). Note that by the triangle inequality, we must have $d(x, y) \leqslant e_i \vee e'_i$ for each i. Thus pre-canonical amalgamation makes $d(x, y)$ maximal while respecting these instances of the triangle inequality. The next proposition provides a condition on when this is sufficient to ensure the resulting diagram satisfies the triangle inequality.

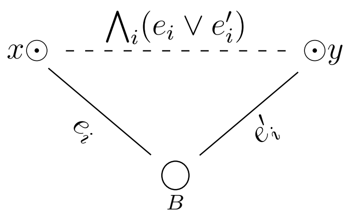

Figure 2.1

Proposition 2.4. *Let Λ be a distributive lattice, and let $\mathcal{K}$ be the class of all finite Λ-ultrametric spaces. Then $\mathcal{K}$ is an amalgamation class, and any amalgamation diagram can be completed by canonical amalgamation.*

Proof. It suffices to check that pre-canonical amalgamation in a 2-point diagram produces a Λ-ultraspeudometric space. In other words, we check that Figure 2.1 satisfies the triangle inequality, given distributivity. Fix some $b_i \in B$ and consider the corresponding triangle.

We have $d(x, y) = \bigwedge_j (e_j \vee e'_j) \leqslant (e_i \vee e'_i) = d(x, b_i) \vee d(b_i, y)$, so this side satisfies the triangle inequality by definition.

The remaining two sides are handled symmetrically, so we only check $d(b_i, y) \leqslant d(b_i, x) \vee d(x, y)$. We have $d(b_i, x) \vee d(x, y) = e_i \vee (\bigwedge_j (e_j \vee e'_j)) = \bigwedge_j (e_i \vee e_j \vee e'_j)$. We now use the path variant of the triangle inequality, which the diagram satisfies even before the distance between x and y is determined. Going from b_i to y, we get that for each j, $e_i \vee e_j \vee e'_j \geqslant e'_i$ (see Figure 2.2), giving $d(b_i, x) \vee d(x, y) \geqslant \bigwedge_j e'_i = d(b_i, y)$.

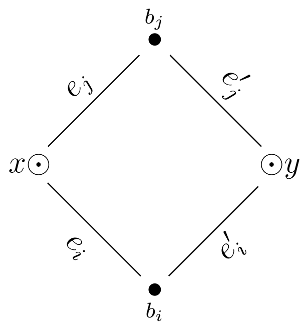

Figure 2.2 □

3 The Representation Theorem

3.1 Preliminaries

Our goal is now the following:

Theorem 3.1 (Representation Theorem). *Let Λ be a finite distributive lattice. Then there is a homogeneous finite dimensional permutation structure whose lattice of $\emptyset$-definable equivalence relations is isomorphic to Λ.*

However, we proceed by stages. The first and main stage is the following:

Theorem 3.2 (Representation Theorem, Part I). *Let Λ be a finite distributive lattice in which $\mathbb{0}$ is meet-irreducible. Then there is a homogeneous finite dimensional permutation structure whose lattice of $\emptyset$-definable equivalence relations is isomorphic to Λ.*

The idea of the proof is to take the homogeneous universal Λ-ultrametric space, viewed as a homogeneous structure in the language $(E_\lambda : \lambda \in \Lambda)$, and to first expand it to a homogeneous structure $\mathcal{M}$ by the addition of finitely many linear orders, in such a way that the meet-irreducible E_λ become $<$-convex for one of the adjoined orders, in the following sense.

Definition 3.3. If E is an equivalence relation and $<$ a linear order, we say that E is $<$-convex if its equivalence classes are convex for the specified order.

Having done so, it is then easy to reinterpret the structure $\mathcal{M}$ obtained as a finite dimensional permutation structure by a definable change of language. Since the second step is much easier than the first, we begin with that.

Lemma 3.4. *Let $\mathcal{M}$ be a homogeneous structure in a finite relational language consisting of symbols of two kinds.*

- *Symbols for equivalence relations $(E_i \mid i \in I)$.*
- *Symbols for linear orders $(<_j \mid j \in J)$.*

Suppose in addition that for every equivalence relation E_i, there is an $\emptyset$-definable linear order $<$ such that E_i is $<$-convex.

Let L be the set of $\emptyset$-definable linear orders on $\mathcal{M}$ and $\mathcal{M}_{lin}$ the reduct to L. Then

(a) $\mathcal{M}_{lin}$ is a homogeneous finite dimensional permutation structure.

(b) The lattice Λ' of $\emptyset$-definable equivalence relations in $\mathcal{M}_{lin}$ contains the E_i for $i \in I$, and is contained in the lattice Λ of $\emptyset$-definable equivalence relations on $\mathcal{M}$.

In particular, if $(E_i \mid i \in I)$ generates Λ, then $\Lambda' = \Lambda$.

Proof. There are finitely many $\emptyset$-definable linear orders on $\mathcal{M}$, so $\mathcal{M}$ is a finite dimensional permutation structure.

(a) To see that $\mathcal{M}_{\text{lin}}$ is homogeneous, it is sufficient to check that the relations given on $\mathcal{M}$ are quantifier-free definable in $\mathcal{M}_{\text{lin}}$. Since the relations $<_j$ are included among the relations of $\mathcal{M}_{\text{lin}}$, we consider an equivalence relation E_i. Fix an $\emptyset$-definable order $<$ such that E_i is $<$-convex. Let $<'$ be the order which agrees with $<$ on M/E_i but is the reversal of $<$ on each E_i-class. Then $<, <'$ are in the language of $\mathcal{M}_{\text{lin}}$, and E_i is $\emptyset$-definable from $<, <'$.

We have proved the first part of (b), and the second part holds simply because $\mathcal{M}_{\text{lin}}$ is a reduct.

The final remark is clear. $\square$

It is an interesting problem to determine the minimum number of additional orders required. We will return to this point in section 3.4.

For the application of this lemma, we have in mind the case in which I is the set of meet-irreducible elements in a given finite distributive lattice Λ, that is, elements which are not the meet of two larger elements. More precisely, we are interested in the meet-irreducible elements of $\Lambda \setminus \{\mathbb{0}, \mathbb{1}\}$, since $\mathbb{0}$ is equality, and $\mathbb{1}$ is trivial, so both are $\emptyset$-definable without quantifiers in any structure.

3.2 The Main Construction

We aim now at the following.

Proposition 3.5 (Main Construction). *Let Λ be a finite distributive lattice in which $\mathbb{0}$ is meet-irreducible, and let $\mathcal{M}_\Lambda$ be the universal homogeneous Λ-ultrametric space, viewed as a structure in the language*

$$(E_\lambda \mid \lambda \in \Lambda)$$

Then there is an expansion $\mathcal{M}$ of $\mathcal{M}_\Lambda$ by a finite set of linear orders with the following properties.

- *$\mathcal{M}$ is homogeneous;*
- *The lattice of $\emptyset$-definable equivalence relations in $\mathcal{M}$ is $(E_\lambda \mid \lambda \in \Lambda)$.*
- *Every equivalence relation E_λ with $\lambda \in \Lambda \setminus \{\mathbb{0}, \mathbb{1}\}$ meet-irreducible is $<$-convex, for some $\emptyset$-definable linear order of $\mathcal{M}$.*

We first complete the proof of Theorem 3.2 modulo Proposition 3.5.

Proof of Theorem 3.2. We begin with a finite distributive lattice Λ in which $\mathbb{0}$ is meet-irreducible and we apply the construction of Proposition 3.5 to the universal homogeneous Λ-ultrametric space $\mathcal{M}_\Lambda$, getting an expansion $\mathcal{M}$ as described, in a language consisting of the equivalence relations $(E_\lambda \mid \lambda \in \Lambda)$, and some additional linear orders.

Let I be the set of meet-irreducible elements of $\Lambda_0 = \Lambda \setminus \{\mathbb{0}, \mathbb{1}\}$, and view $\mathcal{M}$ as a structure in the language $(E_i \mid i \in I)$ together with some linear orders. This structure remains homogeneous since all E_λ are quantifier-free definable from the $(E_i \mid i \in I)$.

The hypotheses of Lemma 3.4 apply to the modified version of $\mathcal{M}$, and by Proposition 3.5 $(E_\lambda \mid \lambda \in \Lambda)$ is the full lattice of $\emptyset$-definable equivalence relations on $\mathcal{M}$, so by Lemma 3.4 the reduct $\mathcal{M}_{\text{lin}}$ is a homogeneous permutation structure with the same lattice of $\emptyset$-definable equivalence relations. □

Thus it will suffice to carry out the main construction. We prepare for this construction with the following.

Lemma 3.6. *Let $A \subseteq B$ be structures for a relational language specifying a set L of nested equivalence relations. Let $<_A$ be an ordering on A with respect to which all the relations $E \in L$, as interpreted in A, are convex (i.e., their classes are $<_A$-convex). Then there is an ordering $<_B$ of B with respect to which all the relations $E \in L$, as interpreted in B, are also convex.*

Proof. By a logical compactness argument, we may suppose that the set L is finite, and proceed by induction on $|L|$.

Let $E \in L$ be minimal. By induction, the induced relation $<_{A/E}$ on A/E can be extended to an ordering $<_{B/E}$ of B/E making each of the induced relations F/E for $F \in L$ convex (this is trivially true for E/E).

Then the relation $R(x, y)$ on B defined by

$$x <_A y \text{ or } x/E <_{B/E} y/E$$

is easily seen to be a partial order, and any extension to a linear order on B will suffice. □

Proof of Proposition 3.5. Recall Λ is a finite distributive lattice, and $\mathbb{0}$ is meet-irreducible.

Let Λ_0 be the set of meet-irreducible elements of $\Lambda \setminus \{\mathbb{0}, \mathbb{1}\}$. Let $\mathcal{L}$ be a set of chains in Λ_0 (i.e., linearly ordered subsets) which covers Λ_0. Although $\mathcal{L}$ might consist simply of the sets $\{\lambda\}$ for $\lambda \in \Lambda_0$, the number of orders required by our construction at this stage will be $|\mathcal{L}|$, and when Λ is finite we will prefer (in practice) to minimize this number, as well as the additional number of orders needed in the proof of Lemma 3.4.

Let $\mathcal{M}_\Lambda$ be the Fraïssé limit of all finite Λ-ultrametric spaces. The expansion $\mathcal{M}$ of $\mathcal{M}_\Lambda$ will also be obtained as a Fraïssé limit, as follows.

Let $\mathscr{L}$ be the language $\{E_\lambda \,|\, \lambda \in \Lambda\}$ and let $\mathcal{A} = \mathcal{A}_\Lambda$ be the amalgamation class of all finite Λ-ultrametric spaces. Via the correspondence with relational structures, we view each $A \in \mathcal{A}$ as an $\mathscr{L}$-structure.

Expand $\mathscr{L}$ to a language $\mathscr{L}^*$ by adding binary relations $<_L$ indexed by $L \in \mathcal{L}$. Consider the class $\mathcal{A}^*$ of $\mathscr{L}^*$-expansions A^* of structures A in $\mathcal{A}$ which satisfy the following conditions.

- Each $<_L$ is a linear order;
- For $L \in \mathcal{L}$ and $\lambda \in L$, the equivalence relation E_λ is $<_L$-convex.

We will show that $\mathcal{A}^*$ is an amalgamation class and that its Fraïssé limit $\mathcal{M}$ is the desired structure.

We restate the points to be proved.

(a) $\mathcal{A}^*$ is an amalgamation class; then its Fraïssé limit may be denoted $\mathcal{M}$.

(b) The reduct of $\mathcal{M}$ to the language $\mathscr{L} = (E_\lambda \,|\, \lambda \in \Lambda)$ is isomorphic to $\mathcal{M}_\Lambda$; so $\mathcal{M}$ may be viewed as an expansion of $\mathcal{M}_\Lambda$.

(c) If $\lambda \in \Lambda \setminus \{\mathbb{0}, \mathbb{1}\}$ is meet-irreducible and $\lambda \in L$ (with $L \in \mathcal{L}$) then E_λ is $<_L$-convex.

(d) The lattice of $\emptyset$-definable equivalence relations in $\mathcal{M}$ is $(E_\lambda \,|\, \lambda \in \Lambda)$.

As point (a) requires a detailed verification, we prefer to treat this point as an independent technical lemma, below. For now we assume this point and complete the proof of points (b, c, d).

(b) The issue here is similar to the one treated in Proposition 5.2 of [2] in a very similar setting.

In order to show that the reduct $\mathcal{M}_\mathscr{L}$ of $\mathcal{M}$ to the language $\mathscr{L}$ is isomorphic to $\mathcal{M}_\Lambda$, it suffices to show that $\mathcal{M}_\mathscr{L}$ is homogeneous, and that $\mathcal{M}_\mathscr{L}$ and $\mathcal{M}_\Lambda$ have the same finite substructures.

The finite substructures of $\mathcal{M}_\Lambda$ are the finite Λ-ultrametric spaces. Given such a structure $A \in \mathcal{A}$, if we apply Lemma 3.6 to the pair $\emptyset \subseteq A$ and each set L, it gives an expansion of A which lies in $\mathcal{A}^*$.

It remains to check the homogeneity of $\mathcal{M}_\mathscr{L}$, which is the point at which we rejoin Proposition 5.2 of [2]. It suffices to check the extension property for $\mathcal{M}_\mathscr{L}$:

If A is a finite substructure of $\mathcal{M}_{\mathscr{L}}$, and B a finite extension of A in $\mathcal{A}$, then there is an embedding of B into $\mathcal{M}_{\mathscr{L}}$ over A.

We let A^* be the substructure of $\mathcal{M}$ whose reduct is A. Applying Lemma 3.6 to A and all $L \in \mathcal{L}$ we get an expansion B^* of B containing A^*, with $B^* \in \mathcal{A}^*$. Then B^* embeds into $\mathcal{M}$ over A^*, so B embeds into $\mathcal{M}_{\mathscr{L}}$ over A.

(c) Given the definition of $\mathcal{M}$ as a Fraïssé limit, the last point is built into the construction.

(d) Calculating the lattice of $\emptyset$-definable equivalence relations in $\mathcal{M}$ requires closer attention.

We consider a 2-type p realized in $\mathcal{M}$. This is the orbit of an ordered pair (a, b) under the automorphism group of $\mathcal{M}$; by homogeneity, it may also be viewed as the isomorphism type of that pair taken in order, and is encoded by the data (E_p, p_{lin}) where $E_p = E_{d(a,b)}$ and p_{lin} is the type of (a, b) in the language restricted to the linear orders $<_L$, which records whether $a <_L b$ or $b <_L a$, for each $L \in \mathcal{L}$.

We may consider such a 2-type as a minimal nontrivial $\emptyset$-definable binary relation on $\mathcal{M}$. Let E'_p denote the smallest equivalence relation containing the relation p: i.e., the transitive and reflexive closure of the symmetrized type $p \cup p^{\text{op}}$, where p^{op} is the type of (b, a).

Claim. Let p be a 2-type realized in $\mathcal{M}$. Then

$$E'_p = E_p$$

Given the claim, consider an arbitrary $\emptyset$-definable equivalence relation E on $\mathcal{M}$. This is the union of the 2-types p contained in E, and hence is the join of the equivalence relations E'_p generated by those types. Since $E'_p = E_p$, this join lies in Λ.

So to verify point (d) it suffices to check the claim.

If $a = b$ then E_p and E'_p are both equality. So suppose $a \neq b$, so that for each linear order L we have either $a <_L b$ or $b <_L a$.

Since the pair (a, b) satisfies E_p, it follows that $E'_p \subseteq E_p$.

Conversely, suppose that we have a pair c, d satisfying $E_p(c, d)$. Let q be the type of (c, d). We extend (c, d) to a triangle (a, c, d) by setting $type(a, c) = type(a, d) = p$. If this triangle belongs to $\mathcal{A}^*$, then by homogeneity it embeds into $\mathcal{M}$ over (c, d), so $E'_p(c, d)$ holds and we are done. (And the proof shows $E_p = p \circ p^{\text{op}}$.)

So let us check on the triangle (a, c, d) belongs to $\mathcal{A}^*$, which amounts to checking the following three conditions.

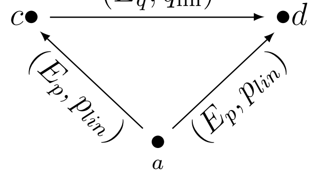

Figure 3.1

1. The Λ-metric triangle inequality holds.

   The metric labels are (E_p, E_p, E_q) where $E_q = d(c, d) \leqslant E_p$. The triangle inequality requires that $E_q \leqslant E_p \vee E_p$ and $E_p \leqslant E_q \vee E_p$, which are both immediate.

2. The relations $<_L$ are linear orders.

   We only need to check transitivity. Since *type*$(a, c) =$ *type*(b, d), for each L we have $a <_L c, d$, or $c, d <_L a$, this is clear.

3. For $L \in \mathcal{L}$ and $\lambda \in L$, each relation E_λ is $<_L$-convex.

   If $E_p \leqslant E_\lambda$ then all points (a, c, d) lie in the same E_λ-class and there is no issue.

   If $E_p \not\leqslant E_\lambda$ then the only two points which can lie in the same E_λ-class are (c, d). But because *type*$(a, c) =$ *type*(a, d), a does not lie between these points in the order $<_L$.

Thus $(a, c, d) \in \mathcal{A}^*$ and the claim is proved. Thus condition (d) holds as well. □

We have unfinished business, corresponding to point (a) above. We state this explicitly.

Lemma 3.7 (Amalgamation Lemma)**.** *Let Λ be a finite distributive lattice in which $\mathbb{0}$ is meet-irreducible. Let Λ_0 be the poset of meet-irreducibles of $\Lambda \setminus \{\mathbb{0}, \mathbb{1}\}$, and let $\mathcal{L}$ be a collection of linearly ordered subsets of Λ_0. Let $\mathcal{A}^*$ be the class of finite structures $(A, d, (<_L \mid L \in \mathcal{L}))$ satisfying the following conditions.*

- *(A, d) is a Λ-ultrametric space.*
- *Each $<_L$ is a linear order.*
- *For $L \in \mathcal{L}$ and $\lambda \in L$, the relation E_λ is $<_L$-convex.*

Then $\mathcal{A}^*$ is an amalgamation class.

Since our amalgamation strategy on $\mathcal{A}^*$ will extend canonical amalgamation on $\mathcal{A}$, we will require the following lemma, which is independent of our amalgamation strategy for the orders.

Lemma 3.8. *Under the hypotheses of Lemma 3.7, consider a two-point amalgamation problem in $\mathcal{A}^*$ with extension points a_1 and a_2. If, after canonical amalgamation, a_1 and a_2 are in the same E_λ-class, with $\lambda \in L$, then no point in the base in a distinct E_λ-class lies between them in $<_L$.*

Proof. Fix $L \in \mathcal{L}$ a chain in Λ_0, and let canonical amalgamation assign $d(a_1, a_2) \leqslant \lambda$ for some $\lambda \in L$. Let b be in the base, with $d(a_1, b), d(a_2, b) \not\leqslant \lambda$. Then we must show $a_1 <_L b \Leftrightarrow a_2 <_L b$.

In a distributive lattice, with λ meet-irreducible, the relation

$$\bigwedge \lambda_i \leqslant \lambda$$

entails

$$\lambda_i \leqslant \lambda$$

for some i. Namely, we have

$$\bigwedge(\lambda \vee \lambda_i) = \lambda \vee \bigwedge \lambda_i = \lambda$$
$$\lambda \vee \lambda_i = \lambda \text{ some } i$$

In view of the definition of $d(a_1, a_2)$ in the canonical amalgam, we conclude that for some x in the base, we have

$$d(a_1, x) \vee d(a_2, x) \leqslant \lambda$$
$$d(a_1, x), d(a_2, x) \leqslant \lambda$$

Then a_1, a_2, x lie in the same E_λ-class, and b in another. Therefore

$$a_1 <_L b \iff x <_L b \iff a_2 <_L b$$

as required. □

We also require an analogue of the above lemma for when the relevant equivalence relation is equality.

Lemma 3.9. *Under the hypotheses of Lemma 3.7, pre-canonical amalgamation is canonical amalgamation.*

Proof. Consider a two-point amalgamation problem in $\mathcal{A}^*$, with extension points a_1 and a_2. Suppose $d(a_1, a_2) = \mathbb{0}$ in the pre-canonical Λ-metric amalgam. In this case, since $\mathbb{0}$ is meet-irreducible, the formula for $d(a_1, a_2)$ produces an element $x \in A$ satisfying

$$d(a_1, x) \vee d(a_2, x) = \mathbb{0}$$
$$d(a_1, x) = d(a_2, x) = \mathbb{0}$$
$$a_1 = x = a_2$$

Thus both extension points are actually in the base, and so this situation is impossible. □

Proof of Lemma 3.7. It suffices to show that $\mathcal{A}^*$ contains solutions to all two-point amalgamation problems $A_0^* \subseteq A_1^*, A_2^*$, $A_i^* = A_0^* \cup \{a_i\}$ for $i = 1, 2$.

Let A_0, A_1, A_2 be the underlying Λ-ultrametric spaces of A_0^*, A_1^*, A_2^*, respectively, and let A be their canonical amalgam. We define a structure A^* as follows.

1. Perform canonical amalgamation to determine $d(a_1, a_2)$.

2. For each $L \in \mathcal{L}$, complete all triangles forced by transitivity constraints. There are two types of transitivity constraints. If there is some $x \in A_0^*$ such that:

   (a) $a_1 <_L x <_L a_2$, then set $a_1 <_L a_2$.

(b) $a_2 <_L x <_L a_1$, then set $a_2 <_L a_1$.

3. For each $L \in \mathcal{L}$, and each $\lambda \in L$, complete all triangles forced by convexity constraints. There are four types of convexity constraints. If there is some $x \in A_0^*$ such that:

   (a) $x <_L a_2$, and $d(a_1, x) \leqslant \lambda$ and $d(a_2, x) \nleqslant \lambda$, then set $a_1 <_L a_2$

   (b) $a_2 <_L x$, and $d(a_1, x) \leqslant \lambda$ and $d(a_2, x) \nleqslant \lambda$, then set $a_2 <_L a_1$

   (c) $x <_L a_1$, and $d(a_2, x) \leqslant \lambda$ and $d(a_1, x) \nleqslant \lambda$, then set $a_2 <_L a_1$

   (d) $a_1 <_L x$, and $d(a_2, x) \leqslant \lambda$ and $d(a_1, x) \nleqslant \lambda$, then set $a_1 <_L a_2$

4. If $<_L$ does not yet hold between a_1 and a_2, it may be determined arbitrarily in a single direction.

Note that step (3), together with Lemma 3.8, will ensure that all convexity conditions are satisfied, and Lemma 3.9 ensures that we don't have to worry about points being identified. It remains to check that $<_L$ is still a linear order after amalgamation, i.e. that it is still asymmetric and transitive. It is well known that step (2) produces an asymmetric and transitive relation. Thus we must check that step (3) does not ruin this asymmetry by assigning the opposite direction to $<_L$ on a triangle already completed from step (2) or on a triangle completed earlier in step (3).

Note that, although there are four types of convexity constraints, there is only one up to swapping the extension points and reversing the orders. Since these symmetries simply require a change of notation in our amalgamation proofs, we may always assume that one of the contradictory triangles is of type $(3a)$. In the following cases, we fix an $L \in \mathcal{L}$, and have $\lambda \in L$.

(i) We first check that no triangle from step (3) contradicts a triangle from step (2). Assuming the triangle from step (3) is of type $(3a)$, there is only one possibility, which is shown in Figure 3.2 (in this figure,and all figures for the remainder of the proof,the label on the edge (p_1, p_2) is $d(p_1, p_2)$, and an arrow from a point p_1 to p_2 indicates $p_1 <_L p_2$; if there is no arrowhead, the direction is irrelevant).

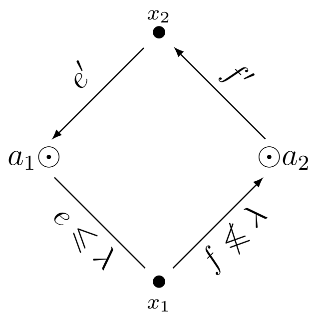

Figure 3.2

Note that transitivity requires $x_1 <_L x_2$. For the first case, assume $d(x_1, x_2) \leqslant \lambda$. Then $x_1 <_L a_2 <_L x_2$ violates convexity. Now assume $d(x_1, x_2) \nleqslant \lambda$. Then $x_1 <_L x_2 <_L a_1$ violates convexity.

(ii) We now check that triangles from step (3) cannot contradict each other. Assuming one of the triangles is of type ($3a$), there are two possibilities. The first case is given in Figure 3.3, where the bottom triangle is of type ($3a$), and the top of type ($3b$).

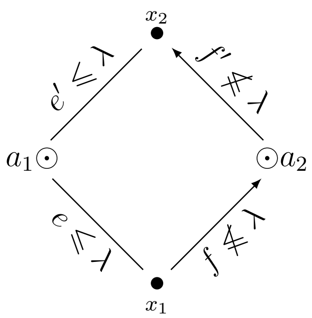

Figure 3.3

Note that transitivity requires $x_1 <_L x_2$. We have $d(x_1, x_2) \leqslant \lambda$. Then $x_1 <_L a_2 <_L x_2$ violates convexity.

The second possible diagram is given in Figure 3.4, where the bottom triangle is of type ($3a$), and the top of type ($3c$).

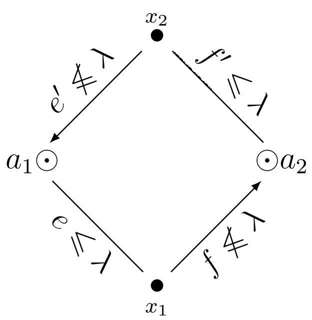

Figure 3.4

We have $e' \leqslant e \vee d(x_1, x_2)$, so we must have $d(x_1, x_2) \nleqslant \lambda$. If $x_1 <_L x_2$, then $x_1 <_L x_2 <_L a_1$ violates convexity. If $x_2 <_L x_1$, then $x_2 <_L x_1 <_L a_2$ violates convexity. $\square$

This concludes Lemma 3.7, hence Proposition 3.5, hence Theorem 3.2.

3.3 The General Case

We now prove the Representation Theorem, reducing the case of a general finite distributive lattice to one where $\mathbb{0}$ is meet-irreducible.

Proposition 3.10. *Let Λ be a finite distributive lattice. Then there is a homogeneous structure of the form*

$$\mathcal{M} = (M, (E_\lambda \mid \lambda \in \Lambda), (<_i \mid i \in I))$$

where I is finite, each E_λ is an equivalence relation, and each $<_i$ is a linear order, satisfying the following conditions.

- *The set of $\emptyset$-definable equivalence relations on M is $(E_\lambda \mid \lambda \in \Lambda)$.*
- *The map $\lambda \mapsto E_\lambda$ is a lattice isomorphism.*
- *For each meet-irreducible $\lambda \in \Lambda$, there is an $i \in I$ such that E_λ is $<_i$-convex.*

We begin with a very general observation.

Lemma 3.11. *Let $\mathcal{M}$ be a structure of the form*

$$(M, (E \mid E \in \Lambda), (<_i \mid i \in I))$$

with Λ a finite sublattice of the lattice of equivalence relations on M and

$$(<_i \mid i \in I)$$

a set of linear orders, such that

- *For each meet-irreducible $E' \in \Lambda$, there is an $i \in I$ such that E' is $<_i$-convex.*

Then for any $E \in \Lambda$, any meet-irreducible $E' \geqslant E$, and any order $<_{E'}$ such that E' is $<_{E'}$-convex, there is a linear order $<_E$ which is quantifier-free definable without parameters, such that E and E' are $<_E$-convex, and $<_E$, $<_{E'}$ induce the same order on M/E'.

Proof. Write E as an intersection $E_1 \cap \cdots \cap E_n$ with $E_i \in \Lambda$ meet-irreducible. Let $i' \in I$ be chosen so that E_i is $<_{i'}$-convex. Let $<_i^*$ be the order induced by $<_{i'}$ on M/E_i, and let $<_{E'}^*$ be the order induced by $<_{E'}$ on M/E'. Via the embedding

$$M/E \hookrightarrow M/E' \times M/E_1 \times \cdots \times M/E_n$$

let $<^*$ be the ordering induced on M/E by the lexicographic product of $<_{E'}^*, <_1^*, \ldots, <_n^*$. Let $<_E$ be a lifting of $<^*$ to M using some fixed order within E-classes (any of the given $<_i$ will do).

Then E is $<_E$-convex and definable from E', the E_i, and the various $<_{i'}$ without quantification. □

Proof of Proposition 3.10. Form $\Lambda' = \Lambda \cup \{\mathbb{0}'\}$ with $\mathbb{0}' < \Lambda$. Then Λ' is a finite distributive lattice whose minimal element is meet-irreducible.

Let $\mathcal{M}' = (M', E'_\lambda, <'_L)_{\lambda \in \Lambda', L \in \mathcal{L}'}$ be the structure afforded by the Representation Theorem for Λ'.

We write $\mathbb{0}$ for the minimal element of Λ (so $\mathbb{0} > \mathbb{0}'$). Let $\mathbb{0} = E_1 \wedge \cdots \wedge E_n$ with E_i meet-irreducible.

Recall that the language of $\mathcal{M}'$ consists of $(E'_\lambda \,|\, \lambda \in \Lambda')$ and a set of orders $(<'_L \;|\, L \in \mathcal{L}')$ with $\mathcal{L}'$ a covering of the meet-irreducible elements other than $\mathbb{0}'$ or $\mathbb{1}$ by linearly ordered sets L. For each $L \in \mathcal{L}'$, let $E'_L = E'_{\inf L}$, and for each $L \in \mathcal{L}'$, using Lemma 3.11, let $<^*_L$ be a $\emptyset$-definable ordering of $\mathcal{M}'$ for which $E'_{\mathbb{0}}$ and E'_L are $<^*_L$-convex, and which agrees with $<'_L$ on $\mathcal{M}'/E'_L$.

Define $\mathcal{M} = \mathcal{M}'/E_{\mathbb{0}}$, $E_\lambda = E'_\lambda/E_{\mathbb{0}}$ for $\lambda \in \Lambda$, and $<_L = <^*_L /E_{\mathbb{0}}$. Note that $\mathcal{M}$ is not a quotient of the structure $\mathcal{M}'$ in its original language, but a quotient of the reduct $\mathcal{M}'_0$ of $\mathcal{M}'$ to the language $(E'_\lambda, <^*_L)_{\lambda \in \Lambda, L \in \mathcal{L}'}$.

Any $\emptyset$-definable relation in $\mathcal{M}$ pulls back to one in $\mathcal{M}'$, so the $\emptyset$-definable equivalence relations on $\mathcal{M}$ are the relations $(E_\lambda \,|\, \lambda \in \Lambda)$, with $E_{\mathbb{0}}$ being equality. For $\lambda \in L$, E_λ is $<_L$-convex.

As the map $\lambda \mapsto E'_\lambda$ $(\lambda \in \Lambda')$ is an isomorphism, and Λ is a sublattice of Λ', the map $\lambda \mapsto E_\lambda$ $(\lambda \in \Lambda)$ is an isomorphism. We are left with one key claim.

Claim 1. The structure $\mathcal{M}$ is homogeneous.

We prefer to rephrase the claim in a more explicit form

Claim 2. Let A be a finite substructure of $\mathcal{M}$, and $<_A$ an arbitrary ordering of A. Then there is a lifting of A to a set of representatives A' in $\mathcal{M}'$ such that the isomorphism type of A', with respect to the relations $(E'_\lambda, <'_L)_{\lambda \in \Lambda', L \in \mathcal{L}'}$, is completely determined by the induced structure on A and the ordering $<_A$.

Let us first note that Claim 2 implies Claim 1. If A, B are isomorphic finite substructures of $\mathcal{M}$, we fix an isomorphism $f : A \to B$. We take an arbitrary ordering $<_A$ of A and the corresponding ordering $<_B$ of B. Then $(A, <_A) \cong (B, <_B)$ via f, so by our claim this isomorphism induces an isomorphism

$$f' : A' \to B'$$

in $\mathcal{M}'$. By homogeneity there is an automorphism α of $\mathcal{M}'$ inducing the map f'. In particular α preserves the relations E'_λ and $<^*_L$, hence induces an automorphism of the quotient M carrying A to B.

So fix $(A, <_A)$ and let $\hat{A} = \{\hat{a} \,|\, a \in A\}$ be an arbitrary set of representatives for A in $\mathcal{M}'$. We wish to specify a structure $\hat{A} \cup A'$ up to isomorphism, meeting the constraints on finite substructures of $\mathcal{M}'$. Then, by homogeneity, we may embed some structure isomorphic to $\hat{A} \cup A'$ into $\mathcal{M}'$ over $\hat{A}$, and the image of the A' part of that structure will give the desired lifting A'. Thus, we specify $\hat{A} \cup A'$ as below.

- $E'_{\mathbb{0}}(a', \hat{a})$.
- $E'_\lambda(a', \hat{b})$ iff $E'_\lambda(\hat{a}, \hat{b})$ for $\lambda \in \Lambda$.

- $E'_\lambda(a', b')$ iff $E_\lambda(a, b)$ for $\lambda \in \Lambda$; $E'_{\mathbb{0}'}(a', b')$ iff $a = b$.
- If $a, b \in A$ and $\neg E_L(a, b)$, then $a' <'_L b'$ iff $a <_L b$, and $a' <'_L \hat{b}$ iff $a <_L b$.
- If $a, b \in A$ and $E_L(a, b)$, then $\hat{b} <'_L a'$, and $a' <'_L b'$ iff $a <_A b$.

We now must check that $\hat{A} \cup A'$ satisfies the constraints on finite substructures of $\mathcal{M}'$, i.e. (a) that the E'_λ are equivalence relations respecting the structure of Λ', (b) that the $<'_L$ are linear orders, and (c) that for every $L \in \mathcal{L}$ and every $\lambda \in L$, the corresponding equivalence relation E'_λ is $<'_L$-convex.

(a) Since $\hat{A}$ is a substructure of $\mathcal{M}'$, the E'_λ are equivalence relations on $\hat{A}$. Then, the E'_λ on A' mirror the E'_λ on $\hat{A}$, and $E'_\lambda(a', \hat{a})$ holds for every $\lambda \in \Lambda$. Thus, the E'_λ are equivalence relations on $\hat{A} \cup A'$.

(b) Fix an $L \in \mathcal{L}$. Between E'_L-classes $<'_L$ agrees with $<_L$. Thus E'_L is $<'_L$-convex, and so any cycle must appear within a single E'_L-class. In a given E'_L class, $\hat{A}$ is already ordered, A' is ordered by $<_A$, and we then put the elements of $\hat{A}$ below those of A'. Thus $<'_L$ will be a linear order.

(c) Fix an $L \in \mathcal{L}$. For $E_\lambda \in L$, $<'_L$ agrees with $<_L$ between distinct E'_λ-classes, so there is nothing to check for convexity. □

3.4 The Number of Orders Needed for the Representation Theorem

Although we have finished the proof of the Representation Theorem, the homogeneous structure we have produced via Lemma 3.4 consisted of all linear orders definable from the main construction of Proposition 3.5. Before continuing to the other direction of Conjecture 2, we take this section to provide a better bound on the number of orders required.

Proposition 3.12. *Let $\{E_i\}$ be a chain of equivalence relations of height at most $2^n - 1$, and let $<$ be a linear order such that each E_i is $<$-convex. Then there exist n linear orders $\{<_j\}$, such that each $<_j$ is quantifier-free definable in $((E_i), <)$, and each E_i is quantifier-free definable in $(<, (<_j))$.*

Proof. We may suppose the chain has height exactly $2^n - 1$, and does not contain equality or the universal relation, since those are already definable. Extend the chain to length $2^n + 1$ by letting E_0 be equality, and E_{2^n} be the universal relation. In the language $(<, (<_j)_{j=1}^n)$, where the $<_j$ are binary relations, enumerate the non-trivial quantifier-free 2-types containing the formula $x < y$ (and so these 2-types merely specify whether $x <_j y$ for each j) as $(p_i : 1 \leqslant i \leqslant 2^n)$. We will use each pairing $p_i \cup p_i^{op}$ to produce an equivalence relation.

Define the relation $R_j(x, y) \Leftrightarrow \bigvee_{\{i | (x <_j y) \in p_i\}} (E_i(x, y) \wedge \neg E_{i-1}(x, y))$. We now define $<_j$ to be the canonical irreflexive, asymmetric extension of $(x < y) \wedge R_j(x, y)$, that is,

$x <_j y \Leftrightarrow ((x < y) \wedge R_j(x,y)) \vee ((y < x) \wedge \neg R_j(y,x))$. Thus each $<_j$ is quantifier-free definable from $((E_i), <)$.

Conversely, we see $x(E_i \backslash E_{i-1})y \Leftrightarrow (x < y \wedge type(x,y) = p_i) \vee (y < x \wedge type(x,y) = p_i^{op})$, so each E_i is quantifier-free definable from $(<, (<_j))$.

It remains to check that the relations $<_j$ are actually linear orders. Clearly $<_j$ is irreflexive and asymmetric, so it suffices to check that it is total and has no cycle. To see it is total, first assume $x < y$. Then if $x \not<_j y$, we must have $\neg R_j(x,y)$. But then $y <_j x$.

We now show there is no cycle. Suppose $x <_j y <_j z <_j x$. Up to a change of notation (cyclically permuting the variables, and reversing $<$ if needed), we may assume $x < y < z$. Let $d(x,y) = \min(E_i : E_i(x,y))$, and define $d(x,z)$ and $d(y,z)$ similarly. By the $<$-convexity of the E_i, we have $d(x,y), d(y,z) \leqslant d(x,z)$, and so the triangle inequality gives either $d(x,z) = d(x,y)$ or $d(x,z) = d(y,z)$. In the first case, our definition of $<_j$ gives $x <_j z$ iff $x <_j y$, since $type(x,z) = type(x,y)$ in the language $(E_i), <)$, and similarly in the second case it gives $x <_j z$ iff $y <_j z$. □

$\mathbb{Q}^n$ with the lexicographic order can naturally be expressed in a language of one order $<$ and a chain of $n-1$ $<$-convex equivalence relations E_i, $1 \leqslant i \leqslant n-1$, given by xE_iy iff x and y agree in the first i coordinates. The lexicographic $\mathbb{Q}^2$ requires two orders to define, and the lexicographic $\mathbb{Q}^3$ requires three. One might expect each new convex equivalence relation to require an additional order, but we already see the exponential growth implied by the above proposition by the lexicographic $\mathbb{Q}^4$, which also only requires three orders.

Corollary 3.13. *Let Λ be a finite distributive lattice, Λ_0 the poset of meet-irreducibles of $\Lambda \backslash \{\mathbb{0}, \mathbb{1}\}$, and $\mathcal{L}$ a set of chains covering Λ_0. Then the dimension of the permutation structure needed for the representation theorem is at most $|\mathcal{L}| + \sum_{L \in \mathcal{L}} \lceil \log_2(|L|+1) \rceil$.*

Proof. First, assume $\mathbb{0}$ is meet-irreducible. Then the intermediate structure produced by the main construction has $|\mathcal{L}|$ linear orders, and each $L \in \mathcal{L}$, considered with the order $<_L$, satisfies the hypotheses of Proposition 3.12. Thus, for each $L \in \mathcal{L}$, the equivalence relations labeled by elements of L are definable after the addition of $\lceil \log_2(|L|+1) \rceil$ linear orders, and so all the meet-irreducibles of Λ_0 are definable after the addition of $\sum_{L \in \mathcal{L}} \lceil \log_2(|L|+1) \rceil$ linear orders, and all of Λ is definable from the elements of Λ_0. The quantifier-free-definability conditions from Proposition 3.10 ensure that the structure obtained by adding these linear orders and removing the equivalence relations is still homogeneous, and so we obtain the bound of the statement.

In the case where $\mathbb{0}$ is not meet irreducible, let Λ' be the lattice obtained by adding a new element $\mathbb{0}'$ below $\mathbb{0}$, and let Λ'_0 be the meet-irreducibles of $\Lambda \backslash \{\mathbb{0}', \mathbb{1}\}$. Then $\Lambda_0 = \Lambda'_0$, and so no additional orders are needed. □

Consider $\mathbb{Q}^2$ with two equivalence relations, each given by equality in one of the coordinates. There are no further non-trivial $\emptyset$-definable equivalence relations, and so the above corollary gives a bound of four orders to define this structure, which is in fact the number needed. We do not know if the bound of Corollary 3.13 is tight in general.

4 Towards the Necessity of Distributivity

We start by recalling some definitions from the introduction.

Definition 4.1. Let $\mathcal{M}$ be a structure with a transitive automorphism group.

1. For $\emptyset$-definable equivalence relations $F \subset E$ on $\mathcal{M}$, set

$$[E : F] = |C/F|$$

   for C any E-class, and call this the *index of F in E.*

2. Let Λ be the lattice of $\emptyset$-definable equivalence relations on $\mathcal{M}$. Then Λ has the *infinite index property (IIP)* if whenever $F \subset E$ for $E, F \in \Lambda$, $[E : F]$ is infinite.

Lemma 4.2. *Given a homogeneous permutation structure $\mathcal{M}$, let Λ be the lattice of $\emptyset$-definable equivalence relations. Then Λ satisfies the IIP.*

Proof. Let $E, F \in \Lambda$ with $F < E$, and let p be a 2-type of orders that is realized in $E \backslash F$. Then p is an intersection of linear orders, and so gives a partial order on a given E-class. Since the structure's automorphism group is transitive, this partial order has no maximal elements, and so contains an infinite linear order L. The 2-type between any pair of elements of L is p, and thus every pair is E-related but not F-related. Thus $[E : F] \geqslant |L|$ is infinite. □

The following lemma is reminiscent of Neumann's lemma that a group cannot be covered by finitely many cosets of subgroups of infinite index. We generalize from the group-theoretic setting, replacing the equivalence relations induced by subgroups with equivalence relations in some lattice, but impose the stronger condition that this lattice must satisfy the IIP.

Lemma 4.3. *Let Λ be a finite lattice of equivalence relations satisfying the IIP, and let $E \in \Lambda$ with C an E-class. Let $\{B_i\}_{i \in I}$ be a finite set of equivalence classes of certain equivalence relations in Λ such that for each of the corresponding equivalence relations $E_i \in \Lambda$, $E_i \ngeqslant E$. Then there exists some $c \in C \backslash \bigcup_{i \in I} B_i$.*

Proof. We proceed by induction on the height of E in the lattice. In the base case, E is equality, and the claim is vacuous.

Now assume E is higher up. We wish to work entirely below E, so we replace each B_i with $B_i \cap C$, and replace each E_i with $E_i \cap E$. Let E' be a maximal equivalence relation strictly below E. Then, for any E_i, we cannot have $E_i \geqslant E'$ unless $E_i = E'$. Since we are trying to avoid finitely many equivalence classes, and by the IIP there are infinitely many E'-classes in C, we may pick an E'-class $C' \subset C$ that is not equal to any of the B_i. Then, letting $I' = \{i \in I | E_i \neq E'\}$, by induction we can find a $c \in C' \backslash \bigcup_{i \in I'} (B_i \cap C) \subseteq C \backslash \bigcup_{i \in I} B_i$. □

We now use the above lemma to prove a one-point extension property for homogeneous Λ-ultrametric spaces where Λ satisfies the IIP. However, we restate the property using amalgamation classes and diagrams.

Lemma 4.4. *Let $\mathcal{K}$ be an amalgamation class of Λ-ultrametric spaces with Fraïssé limit $\mathcal{M}$, and suppose the lattice of $\emptyset$-definable equivalence relations in $\mathcal{M}$ satisfies the IIP. Let $S \in \mathcal{K}$, $b \in S$, $e \in \Lambda$, and $B = \{b, y\}$ with $d(b, y) = e$. Then the canonical amalgam of S and B is in $\mathcal{K}$. (Alternatively, the below amalgamation diagram, with arbitrary first factor, a single point in the base, and a single extension point in the second factor, can be completed by canonical amalgamation.)*

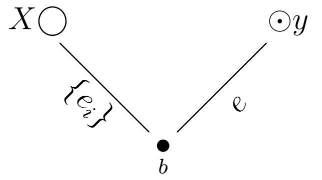

Figure 4.1

Proof. Let $X = S \setminus \{b\}$, and identify the elements of Λ with the corresponding $\emptyset$-definable equivalence relations in $\mathcal{M}$. We choose $y \in \mathcal{M}$ using Lemma 4.3 with $E = e$, C the e-class containing b, and $\{B_i\}$ the set of equivalence classes containing b for every equivalence relation below e, as well as, for every $x_i \in X$, the equivalence classes containing x_i for equivalence relations not above e. Note that the first group of B_i ensures that $d(b, y) = e$.

Now fix an $x_i \in X$, let $d(x_i, b) = e_i$, and let $d(x_i, y) = e'$. From the second group of B_i, we have $e \leqslant e'$. Thus, using the triangle inequality for the upper bound, we have $e \leqslant e' \leqslant e \vee e_i$. Then, $e_i \leqslant e' \vee e = e'$, so $e_i \leqslant e'$ as well. Thus $e' = e \vee e_i$. $\square$

Lemma 4.5. *Let $\mathcal{M}$ be a homogeneous structure consisting of a set of equivalence relations from a finite lattice Λ of equivalence relations satisfying the IIP. Then Λ is distributive.*

Proof. Let $\mathcal{K}$ be the amalgamation class corresponding to $\mathcal{M}$, viewed as a class of Λ-ultrametric spaces.

Claim 1. Suppose both factors of the amalgamation diagram shown in Figure 4.2 are contained in an amalgamation class of Λ-ultrametric spaces, for every $e, f, g \in \Lambda$. Then Λ is distributive.

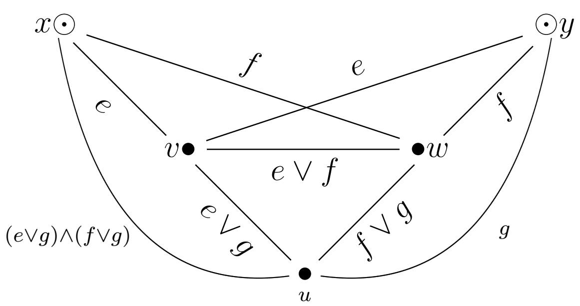

Figure 4.2

Let $d(x, y) = h$ in the completed diagram. Then $h \leqslant e$ and $h \leqslant f$ by the triangle inequality. Going from x to u via y, the triangle inequality gives $(e \vee g) \wedge (f \vee g) \leqslant h \vee g \leqslant (e \wedge f) \vee g$, and so the claim follows.

Claim 2. $\mathcal{K}$ contains both factors of the amalgamation diagram shown in Figure4.2, for every $e, f, g \in \Lambda$.

Because Λ satisfies the IIP, we may use the one-point extension property in Lemma 4.4 to build the factors of Figure 4.2 one point at a time. For the second factor (omitting x), we start with y as a base point, and proceed as in Figure 4.3, adding v, w, and u, in order.

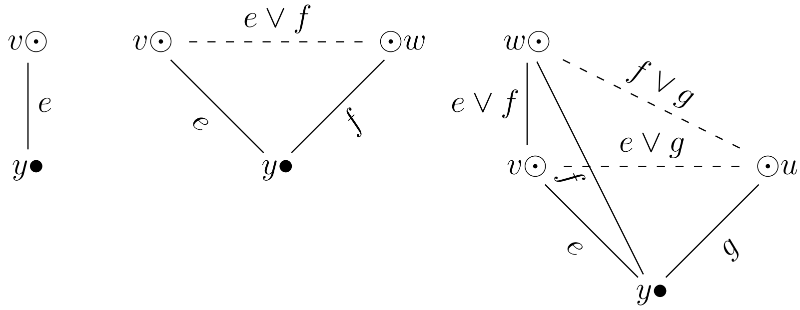

Figure 4.3

The construction of the first factor proceeds similarly, starting with x as a base point, but in the step corresponding to the last diagram of Figure 4.3, we put $d(x, u) = (e \vee g) \wedge (f \vee g)$ instead of $d(x, u) = g$. Since the diagram is then completed by canonical amalgamation, in we must check that $f \vee ((e \vee g) \wedge (f \vee g)) = f \vee g$ and $e \vee ((e \vee g) \wedge (f \vee g)) = e \vee g$. Since the arguments are identical, we will only consider the first identity.

Clearly, $f \vee ((e \vee g) \wedge (f \vee g)) \leqslant f \vee g$, and $f \vee ((e \vee g) \wedge (f \vee g)) \geqslant f \vee g$, since $(e \vee g) \wedge (f \vee g) \geqslant g$. $\square$

Theorem 4.6. *Let Λ be the lattice of $\emptyset$-definable equivalence relations in a homogeneous finite dimensional permutation structure $\mathcal{M}$. If the reduct of $\mathcal{M}$ to the language of equivalence relations from Λ is homogeneous, then Λ is distributive.*

Proof. By Lemma 4.2, Λ satisfies the IIP. Thus, we may apply Lemma 4.5 to conclude. $\square$

5 Forbidden 2-Types

In Cameron's homogeneous permutations, whenever a 2-type is forbidden it forces one order to be equal to another, up to reversal. This need not always be the case in higher dimensional homogeneous permutation structures. Consider $\mathbb{Q}^4$ as a lexicographic order. As discussed following Proposition 3.12, this structure can be defined using 3 orders. Since each $\emptyset$-definable equivalence relation is separated by a single 2-type and its opposite, forbidding a 2-type and its opposite causes one equivalence relation to collapse to one beneath it, and the resulting structure is the lexicographic $\mathbb{Q}^3$. Since we only forbid one pair of 2-types, we cannot have made one order equal to another, up to reversal.

We know two ways to forbid 2-types in homogeneous n-dimensional permutation structures: collapse one order to another, up to reversal, or collapse one equivalence relation to another. In the second case, we must forbid some 3-types compatible with the allowed 2-types, i.e. 3-types such that the restriction to any 2 variables is an allowed 2-type, and transitivity constraints are respected. The following result suggests these two constructions may be typical to some degree.

Proposition 5.1. *Let $\mathcal{K}$ be an amalgamation class of n-dimensional permutation structures. If no 3-type compatible with the allowed 2-types is forbidden, then the forbidden 2-types collectively specify that certain orders agree up to reversal.*

Proof. First, assume that no two orders are equal up to reversal, since otherwise we could pass to a reduct in which this is the case. We will use the notation $[n] = \{1, \dots, n\}$ and let $\binom{[n]}{k}$ denote the set of k-subsets of $[n]$. Given 2-types p and q and $X \subset [n]$, we say p is an X-approximation to q if p and q agree on the orders indexed by elements of X. Given a 2-type t, we will prove t is realized using induction on the size of approximations to t. By assumption, given any $X \in \binom{[n]}{2}$, there is a non-forbidden X-approximation to t; otherwise the orders in X would have to be equal up to reversal.

Before proceeding to the inductive step, we consider the following amalgamation diagram, where p, q, r are 2-types realized in $\mathcal{K}$, and the type of (x_3, x_1) is q. (In this diagram and the next, the label on an arrow from a point p_1 to p_2 gives the 2-type of (p_1, p_2).)

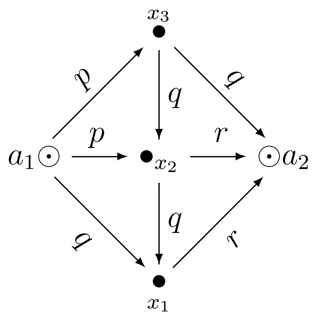

Figure 5.1

Claim 1. Both factors of Figure 5.1 are in $\mathcal{K}$.

We consider only the first factor, since there is a symmetric argument for the second factor. First, note that by transitivity of 2-types, the first factor is the unique amalgam of two triangles, as shown in the following diagram:

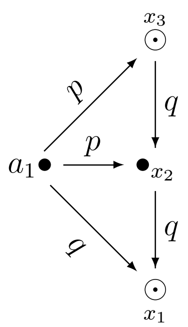

Figure 5.2

Both factors of this diagram are in $\mathcal{K}$ because no 3-types compatible with the allowed 2-types are forbidden, so the only constraint is transitivity. However, because each triangle has two equal sides pointing from or to the same point, all transitivity constraints are satisfied.

Claim 2. There is a unique solution to the diagram in Figure 5.1, given by $a_1 <_i a_2$ iff $<_i$ is true in at least two of p, q, and r.

Note that every pair of p, q, and r appears on a path of length two from a_1 to a_2. Thus if $<_i$ is true in some pair, the path containing that pair will force $a_1 <_i a_2$, and vice versa if $<_i$ is false in some pair.

We are now ready to treat the inductive step of our argument. Suppose the 2-type t has a non-forbidden X-approximation for every $X \in \binom{[n]}{k}$. Fix $Y \in \binom{[n]}{k+1}$. Without loss of generality, we may assume $Y = [k+1]$. Let p be a $\{1, \dots, k\}$-approximation to t, q a $\{1, \dots, k-1, k+1\}$-approximation to t, and r a $\{k, k+1\}$-approximation to t. Then, by Claim 2, the solution to the corresponding diagram in Figure 5.1 will be a Y-approximation to t. $\square$

As mentioned in the introduction, this confirms a special case of the Primitivity Conjecture.

A Homogeneity Background

A.1 Fraïssé's Theorem

For this appendix, let L be a countable language of relations and M a countable L-structure, although the notions introduced here can be defined in greater generality.

Definition A.1. M is *homogeneous* if any isomorphism between finite substructures of M extends to an automorphism of M.

Let $Age(M)$ be the class of finite L-structures isomorphic to a substructure of M. Note $Age(M)$ satisfies the following properties:

(i) $Age(M)$ is closed under isomorphism and substructure

(ii) $Age(M)$ has countably many isomorphism types

(iii) Given, $B_1, B_2 \in Age(M)$, there is a $C \in Age(M)$ such that B_1, B_2 embed in C

Definition A.2. A class $\mathcal{K}$ of finite L-structures has the *amalgamation property* (and will be called an *amalgamation class*) if, given $A, B_1, B_2 \in \mathcal{K}$ with embeddings $f_i : A \to B_i$, there exist a $C \in \mathcal{K}$ and embeddings $g_i : B_i \to C$ such that $g_1 \circ f_1 = g_2 \circ f_2$.

Theorem (Fraïssé)**.** *(a) Let M be homogeneous. Then $Age(M)$ has the amalgamation property.*

(b) Let $\mathcal{K}$ be a collection of finite L-structures satisfying $(i) - (iii)$ from above as well as the amalgamation property. Then, up to isomorphism, there is a unique countable, homogeneous L-structure M with $Age(M) = \mathcal{K}$.

The structure M with $Age(M) = \mathcal{K}$ from part (b) is called the Fraïssé limit of $\mathcal{K}$. Part (b) of the above theorem provides a way to construct homogeneous structures.

Example. Let $\mathcal{K}$ be the class of all finite n-dimensional permutation structures, for some n. Since linear orders can be amalgamated, and the amalgamation strategies can be carried out independently, $\mathcal{K}$ is an amalgamation class. Thus it has a Fraïssé limit, called the fully generic n-dimensional permutation structure.

For more about homogeneity, see Macpherson's survey [3], from which the above material is largely taken.

A.2 Amalgamation Problems and Amalgamation Diagrams

In an amalgamation problem, one is asked to verify the amalgamation condition for specific structures A, B_1, B_2, and embeddings $f_i : A \hookrightarrow B_i$. This problem can be represented by an amalgamation diagram.

In these diagrams, the points of the base A are represented by points, while the points of $B_i \backslash A$, which we call "extension points", are represented by circled points. We sometimes wish to depict an arbitrary finite set, in which case we use a large circle instead of individual points. The extension points of the first factor B_1 are placed on the left side of the diagram, while those of the second factor B_2 are placed on the right side.

Since we are only considering binary languages, the relations are given by putting an edge between any pair of points in one of the B_i and labeling it with the 2-type between these points; a solution to the problem consists of determining the 2-types between the extension points in distinct B_i, which may then be placed on a dotted line between the points.

Examples of amalgamation diagrams, both with and without solutions, may be found throughout the paper.

It is worth noting that in order to verify that some class satisfies the amalgamation property, it suffices to verify a weaker form called 2-point amalgamation, in which each B_i contains one extension point. Although this is true in general, the proof is particularly simple for the cases in this paper since:

- The languages we consider are binary.
- In amalgamation strategies we consider, to determine the 2-type between two extension points, only those extension points and the base are used.

By the first point, a general amalgamation problem only requires determining the 2-types between extension points in separate factors. By the second point, each of these 2-types can be determined independently by solving the 2-point amalgamation problem containing the same base and the two relevant extension points.

Acknowledgements

I am grateful to Gregory Cherlin for our many discussions concerning the material of this paper.

References

[1] Peter J. Cameron. Homogeneous permutations. *The Electronic Journal of Combinatorics*, 9(2):#R2, 2002.

[2] Alexander S. Kechris, Vladimir G. Pestov, and Stevo Todorcevic. Fraïssé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis*, 15(1):106–189, 2005.

[3] Dugald Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311(15):1599–1634, 2011.