THE JOURNAL OF SYMBOLIC LOGIC
Volume 78, Number 4, Dec. 2013

DECIDABILITY OF DEFINABILITY

MANUEL BODIRSKY, MICHAEL PINSKER, AND TODOR TSANKOV

Abstract. For a fixed countably infinite structure Γ with finite relational signature τ, we study the following computational problem: input are quantifier-free τ-formulas $\phi_0, \phi_1, \dots, \phi_n$ that define relations $R_0, R_1, \dots, R_n$ over Γ. The question is whether the relation R_0 is primitive positive definable from $R_1, \dots, R_n$, i.e., definable by a first-order formula that uses only relation symbols for $R_1, \dots, R_n$, equality, conjunctions, and existential quantification (disjunction, negation, and universal quantification are forbidden).

We show decidability of this problem for all structures Γ that have a first-order definition in an ordered homogeneous structure Δ with a finite relational signature whose age is a Ramsey class and determined by finitely many forbidden substructures. Examples of structures Γ with this property are the order of the rationals, the random graph, the homogeneous universal poset, the random tournament, all homogeneous universal C-relations, and many more. We also obtain decidability of the problem when we replace primitive positive definability by existential positive, or existential definability. Our proof makes use of universal algebraic and model theoretic concepts, Ramsey theory, and a recent characterization of Ramsey classes in topological dynamics.

§1. Motivation and the main result.

When studying a countably infinite relational structure Θ, we often wish to know what Θ can express by its relations; for example, which other structures it interprets or defines. Concentrating on the latter, it would be pleasant to have an oracle which, given two structures Θ_1, Θ_2 on the same domain, tells us whether they define one another. If all structures we are interested in have finite signature, this is the same as having an oracle which, given a structure Θ and a relation R on the same domain, tells us whether R can be defined from Θ.

In this context, different notions of definability can be considered. The first notion that comes to mind is probably *first-order definability*: an n-ary relation R is first-order definable in Θ iff there is a first-order formula $\phi(x_1, \dots, x_n)$ over the language of Θ such that for all n-tuples a of elements in Θ we have $a \in R$ iff $\phi(a)$ holds. Sometimes, however, other notions of definability, in particular syntactic restrictions of first-order definability, are useful. One notion that is of importance in theoretical computer science is *primitive positive definability*: a first-order formula is called *primitive positive* iff it is of the form $\exists y_1 \dots \exists y_m. \psi$, where ψ is a conjunction of atomic formulas; and an n-ary relation R is primitive positive definable over Θ iff it is first-order definable in Θ by means of a primitive positive formula $\phi(x_1, \dots, x_n)$.

Received March 5, 2012.

The research leading to these results has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013 Grant Agreement no. 257039). The second author is grateful for support through an APART-fellowship of the Austrian Academy of Sciences.

0022-4812/13/7804-0002/$2.90
DOI:10.2178/jsl.7804020

Primitive positive definability is of importance in the study of the *constraint satisfaction problem of* Θ, denoted by CSP(Θ), in theoretical computer science. In such a problem, the input consists of a primitive positive sentence ψ (that is, a primitive positive formula without free variables), and the question is whether ψ is true in Θ. Primitive positive definability of relations in Θ is important in the study of CSP(Θ) because the CSP for an expansion of Θ by relations that are primitive positive definable in Θ can be reduced (in linear time) to CSP(Θ).

We will present here conditions under which the oracle which is to tell us whether a relation R has a primitive positive definition from a finite language structure Θ can be a computer, i.e., under which the problem is decidable. In order to make the problem suitable for an algorithm, we need a finite representation of the input of the problem, that is, the relation R and the structure Θ. Our approach is to fix a base structure Γ with finite relational language, and to assume that both R and Θ have a quantifier-free definition in Γ. We then represent R and Θ as quantifier-free formulas over Γ. Therefore, the input of our problem are quantifier-free formulas $\phi_0, \ldots, \phi_n$ in the language of Γ, of which ϕ_0 defines the relation R, and $\phi_1, \ldots, \phi_n$ define the relations $R_1, \ldots, R_n$ of Θ; the question is whether there is a primitive positive definition of ϕ_0 that uses only relation symbols for $R_1, \ldots, R_n$. We denote this computational problem by $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$.

An algorithm for primitive positive definability has theoretical consequences in the study of the computational complexity of CSPs. It turns out that hardness of CSP(Θ) can usually be shown by presenting primitive positive definitions of relations for which it is known that the CSP is hard. Therefore, a procedure that decides primitive positive definability of a given relation can be a useful tool to determine the computational complexity of CSPs. For the simplest of countable structures, namely the structure $(X; =)$ having no relations but equality, the decidability of $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$ has been stated as an open problem in [3].

We will show here decidability of $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$ for a large class of structures Γ which we will now define. Let τ be a finite relational signature. The *age* of a τ-structure Δ is the class of all finite τ-structures that embed into Δ. We say that a class $\mathscr{C}$ of finite τ-structures, and similarly a structure with age $\mathscr{C}$, is

- *finitely bounded* (in the terminology of [17]) iff there exists a finite set of finite τ-structures $\mathscr{F}$ such that for all finite τ-structures A we have $A \in \mathscr{C}$ iff no structure from $\mathscr{F}$ embeds into A;
- *Ramsey* iff for all $k \geq 1$ and for all $H, P \in \mathscr{C}$ there exists $S \in \mathscr{C}$ such that for all colorings of the copies of P in S with k colors there exists a copy of H in S on which the coloring is constant (for background in Ramsey theory see [12]);
- *ordered* iff the signature τ contains a binary relation that is interpreted as a total order in every $A \in \mathscr{C}$.

A structure is called *homogeneous* iff all isomorphisms between finite induced substructures[1] extend to automorphisms of the whole structure. A structure Γ is called a *reduct* of a structure Δ with the same domain iff all relations in Γ are first-order definable in Δ. We will prove the following.

[1] In this article, substructures are always meant to be *induced*; see [15].

THEOREM 1. *Let Δ be a structure which is ordered, homogeneous, Ramsey, finitely bounded, and has a finite relational signature. Then for any reduct Γ of Δ with finite relational signature the problem* $\mathrm{Expr_{pp}}(\Gamma)$ *is decidable.*

We remark that for *finite* structures Γ the problem $\mathrm{Expr_{pp}}(\Gamma)$ is in co-NEXPTIME (and in particular decidable). For the variant where the finite structure Γ is part of the input, the problem has recently been shown to be also co-NEXPTIME-hard [23].

Note that since Δ is homogeneous, it has *quantifier elimination*, i.e., every relation which is first-order definable in Δ can be defined by a quantifier-free formula. Hence, choosing $\Gamma = \Delta$, we see that our requirement for the relations in $\mathrm{Expr_{pp}}(\Gamma)$ to be given by quantifier-free formulas does not restrict the range of relations under consideration.

Examples of structures Δ that satisfy the assumptions of Theorem 1 are $(\mathbb{Q}; <)$, the Fraïssé limit of ordered finite graphs (or tournaments [19]), the Fraïssé limit of finite partial orders with a linear extension [19], and the homogeneous universal 'naturally ordered' C-relations. (For definition and basic properties of C-relations, see [1], in particular Theorem 14.7. The fact that the homogeneous universal naturally ordered C-relations have the Ramsey property follows from Theorem 4.3 in [18]; an explicit and elementary verification of the Ramsey property for the binary branching case can be found in [7].) CSPs of reducts of such structures are abundant in particular for qualitative reasoning calculi in Artificial Intelligence. For instance, our result shows that it is decidable whether a given relation from Allen's Interval Algebra [2, 22] is primitive positive definable in a given fragment of Allen's Interval Algebra.

As mentioned above, for $\Gamma = (X; =)$, the decidability of $\mathrm{Expr_{pp}}(\Gamma)$ has been posed as an open problem in [3]. Our results solve this problem, since $(X; =)$ is definable in $\Delta := (\mathbb{Q}; <)$, which is ordered, homogeneous, Ramsey, and finitely bounded: the Ramsey property for this structure follows from the classical Ramsey theorem, and the other properties are easily verified.

Using similar methods, decidability of the analogous problem for other syntactic restrictions of first-order logic can be shown in the same context. A formula is called *existential* iff it is of the form $\exists y_1 \ldots \exists y_m.\ \psi$, where ψ is quantifier-free. It is called *existential positive* iff it is existential and does not contain any negations. For a τ-structure Γ, we denote by $\mathrm{Expr_{ex}}(\Gamma)$ ($\mathrm{Expr_{ep}}(\Gamma)$) the problem of deciding whether a given quantifier-free τ-formula ϕ_0 has an existential (existential positive) definition over the structure with the relations defined by given quantifier-free τ-formulas $\phi_1, \ldots, \phi_n$ in Γ.

THEOREM 2. *Let Δ be a structure which is ordered, homogeneous, Ramsey, finitely bounded, and has a finite relational signature. Then for any reduct Γ of Δ with finite relational signature the problems* $\mathrm{Expr_{ex}}(\Gamma)$ *and* $\mathrm{Expr_{ep}}(\Gamma)$ *are decidable.*

The assumptions on Δ in our theorems fall into two classes: the conditions of being ordered, homogeneous, Ramsey, and having finite relational signature imposed on Δ generally allow for a relatively good understanding (in a non-algorithmic sense) of the reducts of Δ. The recent survey paper [8] summarizes what we know about reducts of such structures—their exciting feature is that many branches of mathematics, including model theory, combinatorics, universal algebra, and even topological dynamics are employed in their study, and indirectly also in our algorithm. The additional condition of being finitely bounded provides a finite representation of Δ via $\mathscr{F}$ used by our algorithm.

This paper is organized as follows. In Section 2 we show that the assumption of Δ being finitely bounded is necessary for our decidability result. We then turn to the proof of Theorems 1 and 2: in Section 3 we cite preservation theorems of the form "R is definable from Θ (in some syntactically restricted form of first-order logic) if and only if certain functions on the domain of Θ (which functions depends on the syntactic restriction) preserve R". Section 4 is devoted to the use of Ramsey theory in order to standardize functions that do not preserve R—if such functions exist. Our decision procedure, presented in Section 5, then uses this standardization of functions and the preservation theorems to check whether or not R is definable from Θ. The paper ends with two sections containing further discussion and open problems.

§2. Undecidability of definability.

This section demonstrates that the assumption in Theorem 2 of Δ being finitely bounded is necessary. We use a class of homogeneous digraphs introduced by Henson [13]. A *tournament* is a directed graph without self-loops such that for all pairs x, y of distinct vertices exactly one of the pairs (x, y), (y, x) is an arc in the graph. For a set of finite directed graphs $\mathscr{N}$, we write $\text{Forb}(\mathscr{N})$ for the class of all finite directed graphs that do not embed any of the structures from $\mathscr{N}$. For any set $\mathscr{N}$ of finite tournaments there exists a countably infinite homogeneous directed graph Γ with age $\text{Forb}(\mathscr{N})$ (this can be shown by amalgamation, see [15]). Moreover, those properties characterize Γ up to isomorphism. Henson specified an infinite set $\mathscr{T}$ of finite tournaments $\Lambda_1, \Lambda_2, \dots$ with the property that Λ_i does not embed into Λ_j if $i \neq j$; the exact definition of this set is not important in what follows. But note that for two distinct subsets $\mathscr{N}_1$ and $\mathscr{N}_2$ of $\mathscr{T}$ the two sets $\text{Forb}(\mathscr{N}_1)$ and $\text{Forb}(\mathscr{N}_2)$ are distinct as well, and so are the respective homogeneous digraphs with age $\text{Forb}(\mathscr{N}_1)$ and $\text{Forb}(\mathscr{N}_2)$. Since there are 2^ω many subsets of the infinite set $\mathscr{T}$, there are also that many distinct homogeneous directed graphs; they are often referred to as *Henson digraphs*.

Proposition 3. *There exists a ordered directed graph Δ which is homogeneous and Ramsey such that* $\text{Expr}_{\text{pp}}(\Delta)$ *and* $\text{Expr}_{\text{ep}}(\Delta)$ *are undecidable.*

Proof. For any Henson digraph Γ, the class $\mathscr{C}$ of all expansions of the structures in the age of Γ by a linear order is a Ramsey class; this can been shown by the partite method [20]. Moreover, there exists a homogeneous ordered digraph Δ with age $\mathscr{C}$ (again by amalgamation, see [15]), and Γ is a reduct of Δ.

We show that non-isomorphic Henson digraphs Γ_1 and Γ_2 have distinct Expr_{pp} problems. In the following, let E denote a binary relation symbol that we use to denote the edge relation in graphs. In fact, we show the existence of a first-order formula ϕ_1 over digraphs such that the input $\phi_0 := E(x, y)$ and ϕ_1 is a yes-instance of $\text{Expr}_{\text{pp}}(\Gamma_1)$ and a no-instance of $\text{Expr}_{\text{pp}}(\Gamma_2)$, or vice-versa. Since there are uncountably many Henson digraphs, but only countably many algorithms, this clearly shows the existence of Henson digraphs Γ such that $\text{Expr}_{\text{pp}}(\Gamma)$ is undecidable. This finishes the proof since Γ is a reduct of an ordered homogeneous Ramsey structure Δ, as we have seen above, and $\mathrm{Expr}_{\mathrm{pp}}(\Delta)$ must be undecidable as well. The same argument shows undecidability of $\mathrm{Expr}_{\mathrm{ep}}(\Delta)$.

By the definition of Γ_1 and Γ_2, there exists a finite digraph Ω which embeds into Γ_1 but not into Γ_2, or that embeds into Γ_2 but not into Γ_1. Assume without loss of generality the former. Let s be the number of elements of Ω, and denote its elements by $a_1, \ldots, a_s$. Let ψ be the formula with variables $x_1, \ldots, x_s$ that has for distinct $i, j \leq s$ a conjunct $E(x_i, x_j)$ if $E(a_i, a_j)$ holds in Ω, and a conjunct $\neg E(x_i, x_j) \wedge x_i \neq x_j$ otherwise. Let ϕ_1 be the formula $\psi \wedge E(x_{s+1}, x_{s+2})$.

Let D_1 be the domain of Γ_1, and consider the relation $R_1 \subseteq (D_1)^{s+2}$ defined by ϕ_1 in Γ_1. Let R be a relational symbol of arity $s+2$. Let Θ be the structure with signature $\{R\}$, domain D_1, and where R denotes the relation R_1. It is clear that $\exists x_1, \ldots, x_s. R(x_1, \ldots, x_s, x, y)$ is a primitive positive definition of $E(x, y)$ in Θ.

Now consider the relation R_2 defined by ϕ_1 in Γ_2 over the domain D_2 of Γ_2. Since Ω does not embed into Γ_2, the precondition of ϕ_1 is never satisfied, and the relation R_2 is empty. Therefore, the structure $(D_2; R_2)$ is preserved by all permutations of D_2, and the relation $E(x, y)$ is not first-order (and in particular not primitive positive) definable in $(D_2; R_2)$. ⊣

We do not know whether there exists a homogeneous structure Δ such that $\mathrm{Expr}_{\mathrm{pp}}(\Delta)$ is undecidable and whose age equals $\mathrm{Forb}(\mathcal{N})$ for a *decidable* class $\mathcal{N}$ of finite structures.

§3. Preservation theorems.

Let Γ be a reduct of a homogeneous finitely bounded Ramsey structure Δ with finite relational signature. Our algorithm for $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$ is based on the fact that if R is not definable from Θ, then there exists a certain kind of function which violates R; in order to decide whether or not R is definable, the algorithm thus searches for such a function. In this section, we shall formulate this fact in more detail.

A structure is called *ω-categorical* iff its first-order theory has exactly one countable model up to isomorphism. For an n-tuple a of elements of a structure Δ, the *type* of a is the set of all first-order formulas over the signature of Δ which have n free variables $x_1, \ldots, x_n$ and which are satisfied by a. By a theorem of Ryll-Nardzewski (see for example the textbook [15]), a structure is ω-categorical iff it has only finitely many different types of n-tuples (called *n-types*), for each $n \geq 1$. From this characterization it is straightforward to see that structures which are homogeneous and have a finite relational signature are ω-categorical; in particular, this is true for the structure Δ of Theorems 1 and 2. For an n-tuple a of elements of a structure Δ, the *orbit* of a is the set $\{\alpha(a) \mid \alpha \in \mathrm{Aut}(\Delta)\}$, where $\mathrm{Aut}(\Delta)$ denotes the *automorphism group* of Δ. It is well-known that a structure is ω-categorical iff it has for every $n \geq 1$ only finitely many orbits of n-tuples (called *n-orbits*). Moreover, in ω-categorical structures two n-tuples have the same type iff they have the same orbit (see again [15]). In particular, every n-ary relation definable over an ω-categorical structure is a finite union of orbits of n-tuples.

Clearly, when Θ is a reduct of a structure Δ, then $\mathrm{Aut}(\Theta) \supseteq \mathrm{Aut}(\Delta)$. Hence, if Δ is ω-categorical, then so is Θ; therefore, all structures that appear in this paper are ω-categorical.

If R is an m-ary relation on a set D, and $f\colon D^n \to D$ is a finitary operation on D, then we say that f *preserves* R iff $f(r_1, \dots, r_n)$ (calculated componentwise) is in R for all m-tuples $r_1, \dots, r_n \in R$. In other words, when $r_i = (r_i^1, \dots, r_i^m) \in R$ for all $i \leq n$, we require that $(f(r_1^1, \dots, r_n^1), \dots, f(r_1^m, \dots, r_n^m)) \in R$. Otherwise, we say that f *violates* R. Observe that a permutation α acting on the domain of a structure Θ is an automorphism iff both α and its inverse preserve all relations of Θ. An *endomorphism* of a structure Θ with domain D is a unary operation $f\colon D \to D$ which preserves all relations of Θ. A *self-embedding* of Θ is an injective unary operation $f\colon D \to D$ which preserves all relations of Θ and all complements of relations in Θ. A *polymorphism* of Θ is a finitary operation $f\colon D^n \to D$ which preserves all relations of Θ.

We can now state the preservation theorem used by our algorithm. Statement (1) is well-known in model theory and follows from the standard proof of the theorem of Ryll-Nardzewski. Items (2) and (3) are consequences of the Theorem of Łos–Tarski and the Homomorphism Preservation Theorem; for these theorems, see [15], for the (straightforward) proofs of statements (2) and (3) see [9]. Item (4) is due to Bodirsky and Nešetřil [6].

THEOREM 4. *Let Θ be an ω-categorical structure, and let R be a relation on its domain.*

(1) *R has a first-order definition in Θ iff R is preserved by all automorphisms of Θ.*
(2) *R has an existential definition in Θ iff R is preserved by all self-embeddings of Θ.*
(3) *R has an existential positive definition in Θ iff R is preserved by all endomorphisms of Θ.*
(4) *R has a primitive positive definition in Θ iff R is preserved by all polymorphisms of Θ.*

§4. Standardizing functions.

Theorem 4 tells us that if a relation R is not definable in an ω-categorical structure Θ, then this is witnessed by a some finitary function on the domain of Θ; the kind of function depends on the notion of definability. In this section, we show that in the context of Theorems 1 and 2, this is even witnessed by a function which shows a certain regular behavior, making the search for such an (infinite!) function accessible to algorithms. We start by defining what we mean by regular behavior.

4.1. Canonicity.

DEFINITION 5. For a structure Δ and $n \geq 1$, we write S_n^Δ for the set of all n-types in Δ. The cardinality of S_n^Δ is denoted by $o^\Delta(n)$. We write $S^\Delta := \bigcup_{n\geq 1} S_n^\Delta$. For an n-tuple $a \in \Delta$, we write $\mathrm{tp}^\Delta(a)$ for the element of S_n^Δ corresponding to a. We drop the reference to the structure in this notation when the structure is clear from the context.

DEFINITION 6. A *type condition* between two structures Ξ and Ω is a pair (s, t), where $s \in S_n^\Xi$ and $t \in S_n^\Omega$ for the same $n \geq 1$. A function $f\colon \Xi \to \Omega$ *satisfies* a type condition (s, t) iff for all n-tuples $a = (a_1, \dots, a_n)$ in Ξ of type s, the n-tuple $f(a) = (f(a_1), \dots, f(a_n))$ in Ω is of type t.

A *behavior* is a set of type conditions between two structures. A function *has behavior B* iff it satisfies all the type conditions of the behavior B. For $n \geq 1$, a behavior B is called *n-complete* iff for all types $s \in S_n^{\Xi}$ there is a type $t \in S_n^{\Omega}$ such that $(s,t) \in B$. It is called *complete* iff it is n-complete for all $n \geq 1$.

A function $f\colon \Xi \to \Omega$ is *canonical* (n-canonical) iff it has a complete (n-complete) behavior.

For $F \subseteq \Xi$ we say that f *satisfies a type condition* (s,t) *on* F iff for all n-tuples $a = (a_1, \ldots, a_n)$ in F of type s (in Ξ, not in the substructure induced by F), the n-tuple $f(a) = (f(a_1), \ldots, f(a_n))$ in Ω is of type t. The notions of *having a behavior on* F and of *being canonical on* F are then defined naturally.

Observe that a complete behavior is just a function from S^{Ξ} to S^{Ω} which respects the sorts, i.e., n-types are sent to n-types. We remark that not every such function is necessarily the behavior of a canonical function from Ξ to Ω, but every canonical function from Ξ to Ω does define a function from S^{Ξ} to S^{Ω}. A behavior is just a partial function from S^{Ξ} to S^{Ω} respecting the sorts.

Definition 7. For a relational structure Δ, we write $n(\Delta)$ for the supremum of the arities of the relations of Δ.

Suppose that $n(\Xi)$ is finite and that Ξ has *quantifier elimination*, i.e., every first-order formula in the language of Ξ is equivalent to a quantifier-free formula over Ξ; this is in particular the case for the structure Δ of Theorems 1 and 2, since homogeneity implies quantifier elimination. Then the type of any tuple in Ξ is determined by the types of its subtuples of length $n(\Xi)$. If moreover the same condition holds for Ω (in particular, if $\Omega = \Xi$), and we set n to be the maximum of $n(\Xi)$ and $n(\Omega)$, then a total function from S_n^{Ξ} to S_n^{Ω} automatically defines a total function from S^{Ξ} to S^{Ω}. In other words, a function $f\colon \Xi \to \Omega$ is canonical iff it is n-canonical. Note also that S_k^{Ξ} is finite for every $k \geq 1$ since Ξ is ω-categorical (this follows if Ξ has quantifier elimination and finite relational signature, cf. [15]). Therefore, canonical functions can be represented by finite objects, namely by functions from S_n^{Ξ} to S_n^{Ω}. Since Ω is ω-categorical as well, there are only finitely many functions from S_n^{Ξ} to S_n^{Ω}, and hence there exist only finitely many complete behaviors between Ξ and Ω, allowing an algorithm to check all of them. Roughly, our goal in the following is to prove that functions witnessing that a relation R is not definable in Θ can be assumed to be canonical; it will turn out that this is almost true.

4.2. Calling Ramsey.

Lemma 8. *Let Ξ be ordered Ramsey, let Ω be ω-categorical, and let $f\colon \Xi \to \Omega$ be a function. Then for all finite substructures $F \subseteq \Xi$ there is a copy of F in Ξ on which f is canonical.*

Proof. Set $n := n(\Xi)$, and let $m := o^{\Omega}(n)$. Now f defines a coloring of the n-tuples in Ξ by m colors: the color of a tuple a is just the type of $f(a)$ in Ω. Note that if P, S are ordered structures, then coloring copies of P in S is the same as coloring tuples of type $\mathrm{tp}(p)$, where p is any tuple which enumerates P—this is because every copy of P in S contains precisely one tuple of type $\mathrm{tp}(p)$, and every tuple of type $\mathrm{tp}(p)$ in S induces precisely one copy of P in S.

Given any finite substructure F of Ξ, enumerate all types of n-tuples that occur in F by $t_1, \ldots, t_k$. There is a substructure S_1 of Ξ such that whenever all tuples of type t_1 in S_1 are colored with m colors, then there exists a substructure H_1 of S_1 isomorphic to F on which the coloring is constant. Further, there is a substructure

S_2 of Ξ such that whenever all tuples of type t_2 in S_2 are colored with m colors, then there exists a substructure H_2 of S_2 isomorphic to S_1 on which the coloring is constant. We iterate this k times, arriving at a structure S_k. Now going back the argument, we find that S_k contains a copy of F on which all colorings are constant. That means that f is canonical on this copy. $\dashv$

We remark that this lemma would be false if one dropped the order assumption: consider for instance $\Xi = (\mathbb{Q}; =)$, $\Omega = (\mathbb{Q}; <)$, and any injection from Ξ to Ω. In that case, f is not canonical on any substructure $F \subseteq \Xi$ with at least two elements.

We will now use Lemma 8 in order to show that for ordered homogeneous Ramsey structures Δ with finite relational signature, arbitrary functions from Δ to Δ *generate* canonical functions from Δ to Δ. To introduce this notion, we make the following observation. The set $\mathrm{End}(\Delta)$ of endomorphisms of a structure Δ forms a transformation monoid, i.e., it is closed under composition $f \circ g$ and contains the identity function id. Moreover, it is *closed* (also called *locally closed* or *local*) in the topological sense, i.e., it is a closed subset of the space D^D, where D is the domain of Δ equipped with the discrete topology. This implies that if a set $\mathscr{F}$ of functions from D to D preserves a set of given relations, then so does the smallest closed monoid containing $\mathscr{F}$. This motivates the following definition.

Definition 9. Let D be a set, $g\colon D \to D$, and let $\mathscr{F}$ be a set of functions from D to D. We say that $\mathscr{F}$ *generates* g iff g is contained in the smallest closed monoid containing $\mathscr{F}$. For a structure Δ with domain D and a function $f\colon D \to D$, we say that f *generates* g *over* Δ iff $\{f\} \cup \mathrm{Aut}(\Delta)$ generates g. Equivalently, for every finite subset F of Δ, there exists a term $\alpha_0 \circ (f \circ \alpha_1 \circ \cdots \circ f \circ \alpha_n)$, where $n \geq 0$ and $\alpha_i \in \mathrm{Aut}(\Delta)$ for $0 \leq i \leq n$, which agrees with g on F.

Note that every operation $f\colon D \to D$ generates an operation g over Δ that is canonical as a function from Δ to Δ, namely the identity operation. What we really want is that f generates over Δ a canonical function g which represents f in a certain sense—it should be possible to retain specific properties of f when passing to the canonical function. For example, when f violates a given relation R, then we would like to have a canonical g which also violates R—this is clearly not the case for the identity function. Unfortunately, f might be such that it violates a relation R without generating any function that is canonical as a function from Δ to Δ and that violates R.

We therefore have to refine our method: we would like to fix constants $c_1, \ldots, c_n \in \Delta$ which witness that f violates R and then have canonical behavior relative to these constants, i.e., on the structure $(\Delta, c_1, \ldots, c_n)$ which is Δ enriched by the constants $c_1, \ldots, c_n$. In order to do this, we must assure that $(\Delta, c_1, \ldots, c_n)$ still has the Ramsey property. This leads us into topological dynamics.

4.3. An escapade in topological dynamics. The goal of this subsection is to show the following propositiofn by using a recent characterization of the Ramsey property in topological dynamics.

Proposition 10. *Let Δ be ordered homogeneous Ramsey, and let $c_1, \ldots, c_n \in \Delta$. Then $(\Delta, c_1, \ldots, c_n)$ is ordered homogeneous Ramsey as well.*

We remark that it is easy to see that the expansion of any homogeneous structure by finitely many constants is again homogeneous, and that the nontrivial part of the proposition concerns the Ramsey property. We do not know if the same proposition holds if one does not assume Δ to be ordered.

To prove the proposition, we use a theorem from [16]. A *topological group* is a group $(G; \cdot)$ together with a topology on G such that $(x, y) \mapsto xy^{-1}$ is continuous from G^2 to G. A group action of G on a topological space X is *continuous* iff it is continuous as a function from $G \times X$ into X.

Definition 11. A topological group is *extremely amenable* iff any continuous action of the group on a compact Hausdorff space has a fixed point.

Theorem 12 (Kechris, Pestov, Todorcevic [16]). *An ordered homogeneous structure is Ramsey iff its automorphism group is extremely amenable.*

Thus the automorphism group of the structure Δ in Proposition 10 is extremely amenable. Note that the automorphism group of $(\Delta, c_1, \ldots, c_n)$ is an open subgroup of $\mathrm{Aut}(\Delta)$. The proposition thus follows from the following fact.

Lemma 13. *Let G be an extremely amenable group, and let H be an open subgroup of G. Then H is extremely amenable.*

Proof. Let H act continuously on a compact space X; we will show that this action has a fixed point. Denote by $H \backslash G$ the set of right cosets of H in G, i.e., $H \backslash G = \{Hg : g \in G\}$. Denote by $\pi : G \to H \backslash G$ the quotient map and let $s : H \backslash G \to G$ be a section for π (i.e., a mapping satisfying $\pi \circ s = \mathrm{id}$) such that $s(H) = 1$. Let α be the map from $H \backslash G \times G \to H$ defined by

$$\alpha(w, g) = s(w) g s(wg)^{-1}.$$

For $w \in H \backslash G$ and $g \in G$, note that $s(w)g$ and $s(wg)$ lie in the same right coset of H, namely wg, and hence the image of α is H. The map α satisfies[2]

$$\begin{aligned} \alpha(w, g_1 g_2) &= s(w) g_1 g_2 (s(w g_1 g_2))^{-1} \\ &= s(w) g_1 s(w g_1) s(w g_1)^{-1} g_2 (s(w g_1 g_2))^{-1} \\ &= \alpha(w, g_1) \alpha(w g_1, g_2). \end{aligned}$$

As H is open, $H \backslash G$ is discrete. Hence, s is continuous, and therefore α is continuous as a composition of continuous maps. The *co-induced action* $G \curvearrowright X^{H \backslash G}$ of G on the product space $X^{H \backslash G}$ is defined by

$$(g \cdot \xi)(w) = \alpha(w, g) \cdot \xi(wg).$$

To check that this action is continuous, it suffices to see that the map $(g, \xi) \mapsto (g \cdot \xi)(w)$ from $G \times X^{H \backslash G}$ to X is continuous for every fixed $w \in H \backslash G$. We already know that α is continuous and that the action $H \curvearrowright X$ is continuous. To see that $(g, \xi) \mapsto \xi(wg)$ is continuous, suppose that $(g_n, \xi_n)_{n \in \omega}$ converges to (g, ξ). Let $w = Hk$. As $(g_n)_{n \in \omega}$ converges to g and $k^{-1}Hk$ is open, we will have that eventually $g_n g^{-1} \in k^{-1}Hk$, giving that $k g_n (kg)^{-1} \in H$, or, which is the same, $Hkg_n = Hkg$. We obtain that $wg_n = wg$ for sufficiently large n. Therefore the sequence of the $\xi_n(wg_n)$ converges to $\xi(wg)$.

By the extreme amenability of G, this action has a fixed point ξ_0. Now we check that $\xi_0(H) \in X$ is a fixed point of the action $H \curvearrowright X$. Indeed, for any $h \in H$,

[2] Such maps are called *cocycles*, and the given identity is called the *cocycle identity*.

$h \cdot \xi_0 = \xi_0$ and we have

$$\xi_0(H) = (h \cdot \xi_0)(H) = \alpha(H, h) \cdot \xi_0(Hh) = h \cdot \xi_0(H),$$

finishing the proof. ⊣

We remark that recently, Miodrag Sokic has given a new purely combinatorial proof of Proposition 10.

4.4. Minimal unary functions. Using Proposition 10, we can now prove a 'canonization lemma' that will be central in what follows.

LEMMA 14. *Let* Δ *be ordered homogeneous Ramsey with finite relational signature,* $f\colon \Delta \to \Delta$*, and let* $c_1, \dots, c_n \in \Delta$*. Then* f *generates over* Δ *a function which agrees with* f *on* $\{c_1, \dots, c_n\}$ *and which is canonical as a function from* $(\Delta, c_1, \dots, c_n)$ *to* Δ*.*

PROOF. Let $(F_i)_{i\in\omega}$ be an increasing sequence of finite substructures of $(\Delta, c_1, \dots, c_n)$ such that $\bigcup_{i\in\omega} F_i = (\Delta, c_1, \dots, c_n)$. By Lemma 8, for each $i \in \omega$ we find a copy F_i' of F_i in $(\Delta, c_1, \dots, c_n)$ on which f is canonical. By the homogeneity of $(\Delta, c_1, \dots, c_n)$, there exist automorphisms α_i of $(\Delta, c_1, \dots, c_n)$ sending F_i to F_i', for all $i \in \omega$. Since there are only finitely many type conditions for $n((\Delta, c_1, \dots, c_n))$-tuples, we may assume that if f satisfies a type condition on F_i', then it satisfies the same type condition on F_{i+1}'. Then we can inductively pick automorphisms β_i of $(\Delta, c_1, \dots, c_n)$ such that $\beta_{i+1} \circ f \circ \alpha_{i+1}$ agrees with $\beta_i \circ f \circ \alpha_i$ on F_i, for all $i \in \omega$. The union over the functions $\beta_i \circ f \circ \alpha_i\colon F_i \to \Delta$ is a canonical function from $(\Delta, c_1, \dots, c_n)$ to Δ. ⊣

The set of all closed transformation monoids on a fixed domain D forms a complete lattice with respect to inclusion; it is the lattice of all endomorphism monoids of structures with domain D. Lemma 14 has the following interesting consequence for this lattice.

DEFINITION 15. Let $\mathscr{N}, \mathscr{M}$ be closed monoids over the same domain. We say that $\mathscr{N}$ is *minimal above* $\mathscr{M}$ iff $\mathscr{M} \subsetneq \mathscr{N}$ and $\mathscr{M} \subsetneq \mathscr{R} \subseteq \mathscr{N}$ implies $\mathscr{R} = \mathscr{N}$ for all closed monoids $\mathscr{R}$.

Clearly, every minimal monoid above $\mathscr{M}$ is generated by a single function together with $\mathscr{M}$; such functions are called *minimal* as well (cf. [9]).

LEMMA 16. *Let* Θ *be a structure with a finite relational signature which is a reduct of an ordered homogeneous Ramsey structure* Δ *in a finite relational signature, and let* $\mathscr{N}$ *be a minimal closed monoid above* $\operatorname{End}(\Theta)$*. Then there exist constants* $c_1, \dots, c_{n(\Theta)} \in \Delta$ *and a function* f *which is canonical as a function from* $(\Delta, c_1, \dots, c_{n(\Theta)})$ *to* Δ *such that* $\mathscr{N}$ *is generated by* $\operatorname{End}(\Theta)$ *and* f*.*

PROOF. Pick any $g \in \mathscr{N} \setminus \operatorname{End}(\Theta)$. Since $g \notin \operatorname{End}(\Theta)$, there exist a relation R of Θ and a tuple $c := (c_1, \dots, c_{n(\Theta)})$ such that R is violated on this tuple. By Lemma 14, g generates a function f over Δ which is canonical as a function from $(\Delta, c_1, \dots, c_{n(\Theta)})$ to Δ and which is identical with g on $\{c_1, \dots, c_{n(\Theta)}\}$. Then f and $\operatorname{End}(\Theta)$ generate $\mathscr{N}$. ⊣

PROPOSITION 17. *Let* Θ *be a finite relational signature reduct of an ordered homogeneous finite relational signature Ramsey structure* Δ*. Then there are finitely many minimal closed monoids above* $\operatorname{End}(\Theta)$*, and every closed monoid properly containing* $\operatorname{End}(\Theta)$ *contains a minimal one.*

PROOF. Observe that if c, d are tuples of the same type in Δ, and f, g are canonical functions from (Δ, c) and (Δ, d) to Δ, respectively, and their (complete) behaviors are identical, then f and g generate one another over Δ. Thus, there are only finitely many inequivalent (in the sense of 'do not generate one another') functions generating minimal monoids. The upper bound for minimal monoids is the following: set $j := o^{\Delta}(n(\Theta))$ (there are that many inequivalent choices for the tuple of constants of length $n(\Theta)$ in Δ). For every type of an $n(\Theta)$-tuple c in Δ, set $r_c := o^{(\Delta,c)}(n(\Delta))$. Set r to be the maximum of the r_c. Define moreover $s := o^{\Delta}(n(\Delta))$. Then a bound for the number of inequivalent minimal functions over $\mathrm{End}(\Theta)$ is $j \cdot s^r$.

By the proof of Lemma 16, every closed monoid properly containing $\mathrm{End}(\Theta)$ contains a canonical function from $(\Delta, c_1, \dots, c_{n(\Theta)})$ to Δ which is not an element of $\mathrm{End}(\Theta)$, for constants $c_1, \dots, c_{n(\Theta)} \in \Delta$. Since there are only finitely many inequivalent such functions, we conclude that there exists a finite set $\mathscr{F}$ of functions such that every closed monoid which properly contains $\mathrm{End}(\Theta)$ contains an element of $\mathscr{F}$. It follows that every closed monoid properly containing $\mathrm{End}(\Theta)$ contains a minimal monoid. ⊣

4.5. Minimal higher arity functions. Since primitive positive definability is characterized by finitary functions rather than unary functions (recall Theorem 4), we have to generalize our method to higher arities.

DEFINITION 18. Let $\Xi_1, \dots, \Xi_m$ be a structures. For a tuple x in the product $\Xi_1 \times \cdots \times \Xi_m$ and $1 \leq i \leq m$, we write x_i for the i-th coordinate of x. The *type* of a sequence of tuples $a^1, \dots, a^n \in \Xi_1 \times \cdots \times \Xi_m$, denoted by $\mathrm{tp}(a^1, \dots, a^n)$, is the m-tuple containing the types of $(a^1_i, \dots, a^n_i)$ in Ξ_i for each $1 \leq i \leq m$.

With this definition, also the notions of *type condition*, *behavior*, (n-)*complete behavior*, and (n-)*canonical* generalize in complete analogy from functions $f \colon \Xi \to \Omega$ to functions $f \colon \Xi_1 \times \cdots \times \Xi_m \to \Omega$ whose domain is a product. It is folklore that the Ramsey property is not lost when going to products; for the reader's convenience, we provide a proof here.

LEMMA 19 (The ordered Ramsey product lemma). *Let $\Xi_1, \dots, \Xi_m$ be ordered and Ramsey, and set $\Xi := \Xi_1 \times \cdots \times \Xi_m$. Let moreover a number $k \geq 1$, an n-tuple $(a^1, \dots, a^n) \in \Xi$, and finite $F_i \subseteq \Xi_i$ be given. Then there exist finite $S_i \subseteq \Xi_i$ with the property that whenever the n-tuples in $S := S_1 \times \cdots \times S_m$ of type $\mathrm{tp}(a^1, \dots, a^n)$ are colored with k colors, then there is a copy of $F := F_1 \times \cdots \times F_m$ in S on which the coloring is constant.*

PROOF. We use induction over m. The base case $m = 1$ is trivial, so assume $m > 1$ and that the lemma holds for $m - 1$. For all $1 \leq i \leq n$, set $c^i := (a^i_1, \dots, a^i_{m-1})$. By the induction hypothesis, there exist finite $S_i \subseteq \Xi_i$ for all $1 \leq i \leq m-1$ such that whenever its n-tuples of type $\mathrm{tp}(c^1, \dots, c^n)$ are colored with k colors, then there is a copy of $F_1 \times \cdots \times F_{m-1}$ in $S_1 \times \cdots \times S_{m-1}$ on which the coloring is constant. Let p be the number of n-tuples of this type in $S_1 \times \cdots \times S_{m-1}$. Also by induction hypothesis, there exists a finite $S_{m,1} \subseteq \Xi_m$ with the property that whenever its n-tuples of type $\mathrm{tp}(a^1_m, \dots, a^n_m)$ are colored with k colors, then it contains a monochromatic copy of F_m. Further, there is a finite $S_{m,2} \subseteq \Xi_m$ with the property that whenever its subsets of this type are colored with k colors, then it contains a monochromatic copy of $S_{m,1}$. Continue constructing finite substructures of Ξ_m like that, arriving at $S_m := S_{m,p}$.

We claim that $S := S_1 \times \cdots \times S_m$ has the desired property. To see this, let a coloring χ of the n-tuples in S of type $\mathrm{tp}(a^1, \ldots, a^n)$ be given. Let $b(1), \ldots, b(p)$ be an enumeration of all the n-tuples in $S_1 \times \cdots \times S_{m-1}$ which have type $\mathrm{tp}(c^1, \ldots, c^n)$. For $1 \leq i \leq p$ and $1 \leq j \leq n$, we write $b(i)^j$ for the j-th component of $b(i)$ (note that this component is an $(m-1)$-tuple in $S_1 \times \cdots \times S_{m-1}$). Now for all $1 \leq i \leq p$, define a coloring χ^i of the n-tuples $t = (t^1, \ldots, t^n)$ in S_m of type $\mathrm{tp}(a^1_m, \ldots, a^n_m)$ by setting $\chi^i(t) := \chi(b(i)^1 * t^1, \ldots, b(i)^n * t^n)$, where $r * s$ denotes the concatenation of two tuples r, s. By thinning out S_m p times, we obtain a copy F'_m of F_m in S_m on which each coloring χ^i is constant with color q^i. Now by that construction, all n-tuples $b(i)$ have been assigned a color q^i, the assignment thus being a coloring of all the n-tuples of type $\mathrm{tp}(c^1, \ldots, c^n)$ in $S_1 \times \cdots \times S_{m-1}$. By the choice of that product, there is a copy $F'_1 \times \cdots \times F'_{m-1}$ of $F_1 \times \cdots \times F_{m-1}$ in $S_1 \times \cdots \times S_{m-1}$ on which that coloring is constant, say with value q. But that means that if a tuple $(d^1, \ldots, d^n) \in F'_1 \times \cdots \times F'_m$ has type $\mathrm{tp}(a^1, \ldots, a^n)$, then $\chi(d^1, \ldots, d^n) = q$, proving our statement. ⊣

We now generalize the notion of a transformation monoid to higher arities. Denote the set of all polymorphisms of Δ by $\mathrm{Pol}(\Delta)$. Irrespectively of the structure Δ, this set contains all finitary projections and is closed under composition. Sets of finitary functions with these two properties are referred to as *clones*—for a survey of clones on infinite sets, see [11]. In addition, the clone $\mathrm{Pol}(\Delta)$ is a closed subset of the sum space of the spaces D^{D^n}, where D is again taken to be discrete; such clones are called *closed*, *local*, or *locally closed* (cf. the corresponding terminology for monoids before). This means that if a set $\mathscr{F}$ of finitary functions on a domain D preserves a set of given relations, then so does the smallest closed clone containing $\mathscr{F}$, motivating the following extension of Definition 9.

Definition 20. Let D be a set, $g\colon D^m \to D$, and let $\mathscr{F}$ be a set of finitary operations on D. We say that $\mathscr{F}$ *generates* g iff g is contained in the smallest closed clone containing $\mathscr{F}$. For a structure Δ with domain D and a function $f\colon D^n \to D$, we say that f *generates* g *over* Δ iff $\{f\} \cup \mathrm{Aut}(\Delta)$ generates g. Equivalently, for every finite subset F of Δ^m, there exists an m-ary term built from f, $\mathrm{Aut}(\Delta)$, and projections, which agrees with g on F.

As before, finitary functions on ordered homogeneous Ramsey structures generate canonical functions, and we can add constants to the language.

Lemma 21. *Let Δ be ordered homogeneous Ramsey with finite relational signature, and let $f\colon \Delta^m \to \Delta$. Let moreover finite tuples $c_1 = (c_1^1, \ldots, c_1^{n_1}), \ldots, c_m = (c_m^1, \ldots, c_m^{n_m})$ of constants in Δ be given. Then f generates over Δ an m-ary operation g on Δ which is canonical as a function from $(\Delta, c_1) \times \cdots \times (\Delta, c_m)$ to Δ and which agrees with f on all tuples $(c_1^{j_1}, \ldots, c_m^{j_m})$.*

Proof. Since the proof is almost identical with the proof of Lemma 14, we only sketch it.

The first step is to see that whenever $F_i \subseteq (\Delta, c_i)$ are finite for $1 \leq i \leq n$, then there exist copies $F'_i \subseteq (\Delta, c_i)$ of F_i such that f is canonical on $F'_1 \times \cdots \times F'_m$; this is proven exactly the same way as Lemma 8, using Lemma 19 instead of the classical Ramsey property.

From here on, the proof is exactly the same as the proof of Lemma 14. ⊣

The set of all closed clones on a fixed domain D forms a complete lattice with respect to inclusion; it is the lattice of all polymorphism clones of structures with domain D. This lattice has been investigated in universal algebra (see [21]).

Definition 22. For closed clones $\mathscr{C}, \mathscr{D}$ on the same set, we say that $\mathscr{D}$ is *minimal* above $\mathscr{C}$ iff $\mathscr{C} \subsetneq \mathscr{D}$ and $\mathscr{C} \subsetneq \mathscr{E} \subseteq \mathscr{D}$ implies $\mathscr{E} = \mathscr{D}$ for all closed clones $\mathscr{E}$. Every minimal closed clone above $\mathscr{C}$ is generated by $\mathscr{C}$ plus a single function f outside $\mathscr{C}$; we call such a function f *minimal* above $\mathscr{C}$ if there is no function of smaller arity which generates (together with $\mathscr{C}$) the same closed clone as f.

Lemma 21 allows us to find the minimal clones above a closed clone on an ordered homogeneous Ramsey structure. The main difference here compared with monoids is that the arities of minimal canonical functions are not bounded a priori, which means that there could be infinitely many minimal clones. The following lemma, which has been observed in [5], yields a bound on the arities of minimal functions.

Lemma 23. *Let Θ be a structure, $m \geq 1$, and let $R \subseteq \Theta^n$ be a relation which intersects precisely m n-orbits of Θ. If a function $f\colon \Theta^p \to \Theta$ violates R, then f generates over Θ a function of arity m which violates R, too.*

Proof. Let $O_1, \ldots, O_m$ be the orbits of Θ that are intersect R, and fix arbitrary tuples $s_i \in O_i$. Since f violates R, there exist $r_1, \ldots, r_p \in R$ such that $f(r_1, \ldots, r_p) \notin R$. Say that $b_i \in O_{j_i}$, for all $1 \leq i \leq p$, and choose for all $1 \leq i \leq p$ an automorphism α_i of Θ sending s_{j_i} to r_i. The function $g(x_1, \ldots, x_m) := f(\alpha_1(x_{i_1}), \ldots, \alpha_p(x_{i_p}))$ has arity m and violates R since $g(s_1, \ldots, s_m) = f(r_1, \ldots, r_p)$ is not in R. ⊣

Proposition 24. *Let Θ be a finite relational signature reduct of an ordered homogeneous Ramsey structure Δ with finite relational signature. Then there are finitely many minimal closed clones above* $\mathrm{Pol}(\Theta)$, *and every closed clone containing* $\mathrm{Pol}(\Theta)$ *contains a minimal one.*

Proof. Let $R_1, \ldots, R_n$ be the relations of Θ. If f is a minimal operation above $\mathrm{Pol}(\Theta)$, then it violates a relation R_i. By Lemma 23, it generates over Θ a function of arity $o^{\Theta}(k_i)$, where k_i is the arity of R_i, which still violates R_i. Setting m to be the maximum of the $o^{\Theta}(k_i)$ where $1 \leq i \leq n$, we get that every minimal clone above $\mathrm{Pol}(\Theta)$ is generated by a function of arity at most m. By Lemma 21, such functions can be made canonical—the rest of the proof is just like the proof of Proposition 17. ⊣

If one wishes to determine the minimal clones above the endomorphism monoid of a structure Θ, then there is a bound on the arities of minimal functions which only depends of the number of 2-orbits of the structure Θ, rather than the number of orbits of possibly longer tuples as in the preceding proof.

Definition 25. Let D be a set, and let $f\colon D^m \to D$ be an operation on D. Then f is called *essentially unary* iff there exist $1 \leq i \leq m$ and $F\colon D \to D$ such that $f(x_1, \ldots, x_m) = F(x_i)$. Conversely, f is called *essential* iff it is not essentially unary.

PROPOSITION 26. *Let Θ be any relational structure for which $o^\Theta(2)$ is finite. Then every minimal closed clone above* $\text{End}(\Theta)$ *is generated by a function of arity at most* $2 \cdot o^\Theta(2) - 1$ *together with* $\text{End}(\Theta)$.

PROOF. Let $\mathscr{D}$ be a minimal closed clone above $\text{End}(\Theta)$. If all the functions in $\mathscr{D}$ are essentially unary, then $\mathscr{D}$ is generated by a unary operation together with $\text{End}(\Theta)$ and we are done. Otherwise, let f be an essential operation in $\mathscr{D}$. Then one can verify that f violates the 3-ary relation P_3 defined by the formula $(x = y) \vee (y = z)$. The assertion then follows from Lemma 23: the 3-ary subrelation of P_3 defined by the formula $x = y$ clearly consists of $o^\Theta(2)$ orbits in Θ; similarly, the 3-ary subrelation defined by $y = z$ consists of the same number of orbits. Since P_3 is the union of these two subrelations, and since the intersection of the two subrelations consists of exactly one orbit (namely, the triples with three equal entries), we obtain $2 \cdot o^\Theta(2) - 1$ different orbits for tuples in P_3. ⊣

Consider the situation where the structure Θ in Proposition 26 is a reduct of a structure Δ, a situation we are often interested in. Then since Δ has at least as many 2-orbits as Θ, we can also write $2 \cdot o^\Delta(2) - 1$ for the arity bound in the proposition. This gives us a uniform bound for all reducts Θ of Δ which is independent of Θ.

§5. **The algorithm.** We now present the algorithm proving Theorem 1; the proof of the two statements of Theorem 2 is a subset. The input to the algorithm are quantifier-free formulas $\phi_0, \dots, \phi_n$ over Γ which define relations $R_0, \dots, R_n$ on the domain D of Γ. Set Θ to be the reduct $(D; R_1, \dots, R_n)$ of Γ, and write $R := R_0$. We will decide whether there is a primitive positive definition of R in Θ.

5.1. Operationalization. If there is no such definition, then since Θ is ω-categorical, by Theorem 4 there is a polymorphism f of Θ which violates R; we call f a *witness*. Our algorithm will now try to build a witness. If it fails to do so, then R is primitive positive definable in Θ; otherwise, it is not.

5.2. Arity reduction. Let k be the arity of R. By Lemma 23, if there exists a witness, then there exists also a witness of arity equal to the number N of those k-orbits in Θ that intersect R. By assumption, Γ has a first-order definition in an ordered homogeneous structure Δ that is finitely bounded, Ramsey, and has finite relational signature. The number N is not larger than $o^\Theta(k)$, which is not larger than $o^\Gamma(k)$ since $\text{Aut}(\Delta) \subseteq \text{Aut}(\Gamma) \subseteq \text{Aut}(\Theta)$. Since Δ is homogeneous, $m := o^\Delta(k)$ equals the number of non-isomorphic substructures of size k. Since the maximal arity of Δ is $n = n(\Delta)$, we have $m \leq 2^{O(k^n)}$. The algorithm now tries to detect a witness of arity m.

5.3. Ramseyfication. If f is a witness of arity m, then there are k-tuples $c_1, \dots, c_m \in R$ such that $f(c_1, \dots, c_m) \notin R$. By Lemma 21, f generates over Δ an m-ary function g which is canonical as a function from $(\Delta, c_1) \times \cdots \times (\Delta, c_m)$ to Δ and which agrees with f on all m-tuples whose i-th component is taken from the k-tuple c_i for all $1 \leq i \leq m$. In particular, g still violates R and preserves Θ, and hence is a witness, too. Our algorithm thus tries to find a witness of this form.

5.4. Finite representation. Let $n := \max(s, n(\Delta), 3)$, where s is the maximal size of the finitely many finite forbidden substructures of Δ. Since $n \geq n(\Delta)$, a function from $(\Delta, c_1) \times \cdots \times (\Delta, c_m)$ to Δ is canonical iff it is n-canonical. Such functions can thus be represented as functions from $S_n^{(\Delta,c_1)} \times \cdots \times S_n^{(\Delta,c_m)}$ to S_n^{Δ}. Note that the type space $S_n^{(\Delta,c_i)}$ only depends on the type of c_i in Δ. Since $o^{\Delta}(k)$ is finite, there are only finitely many choices of types for each tuple c_i—our algorithm tries all such choices. Since Δ has a finite relational signature, and is homogeneous, the types of a tuple are given by the substructure induced by the tuple in Δ, and hence those choices can be made effectively. For each choice for the types of the c_i, and for each function σ from $S_n^{(\Delta,c_1)} \times \cdots \times S_n^{(\Delta,c_m)}$ to S_n^{Δ}, the algorithm checks whether σ is the behavior of a witness.

5.5. Verification. Given σ, we verify the following.

- (Compatibility.) If σ is a behavior of a canonical operation, then for all $1 \leq k \leq n$ it must also be extendible to a function from $S_k^{(\Delta,c_1)} \times \cdots \times S_k^{(\Delta,c_m)}$ to S_k^{Δ}. This is possible in the following situation: if s is an n-type, then it has certain k-*subtypes* t, i.e., projections of tuples of type s onto k coordinates satisfy t. Now products of k-subtypes are automatically sent to a k-subtype under σ: if $s_1, \ldots, s_m$ are n-types and $I \subseteq \{1, \ldots, n\}$ is a set of size k inducing k-subtypes t_i of s_i, then I induces a k-subtype of $\sigma(s_1, \ldots, s_m)$. Our algorithm checks for n-types $p_1, q_1, \ldots, p_m, q_m$ and all $I, J \subseteq \{1, \ldots, n\}$ that if I and J induce identical k-subtypes in p_i and q_i, respectively, then they induce identical k-subtypes in $\sigma(p_1, \ldots, p_m)$ and $\sigma(q_1, \ldots, q_m)$—otherwise, σ is rejected as a candidate. If on the other hand σ satisfies this condition, then it naturally extends to a function from $S^{(\Delta,c_1)} \times \cdots \times S^{(\Delta,c_m)}$ to S^{Δ} respecting arities, and we can compute the value of this function for every argument. In the following, we write σ for this extended function.
- (Violation.) Since R has a first-order definition in Δ, and automorphisms of Δ preserve first-order formulas, it follows that R is a union of orbits, i.e., if a, b are of the same type, then $a \in R$ iff $b \in R$. Set $t := \sigma(\mathrm{tp}^{(\Delta,c_1)}(c_1), \ldots, \mathrm{tp}^{(\Delta,c_m)}(c_m))$. Our algorithm checks that t is not a type in R, since we only want to accept σ if it is the behavior of an operation which violates R on $c_1, \ldots, c_m$.
- (Preservation.) For every relation R_i from Θ, we check that σ "preserves" Θ as follows: write p for the arity of R_i. For all p-types $t_1, \ldots, t_m$ of tuples in R_i, we verify that $\sigma(t_1, \ldots, t_m)$ is the type of a tuple in R_i; otherwise we reject σ.

We now argue that the algorithm finds a σ satisfying our three conditions if and only if there is an m-ary polymorphism of Θ that violates R. It is clear that the type function of a witness will satisfy all the conditions, so one direction is straightforward. For the opposite direction, suppose that σ is accepted by our algorithm. We build a canonical operation from $(\Delta, c_1) \times \cdots \times (\Delta, c_m)$ to Δ in three steps. Let τ be the signature of Δ.

We first construct an infinite structure Π with domain D^m and signature $\tau \cup \{\sim\}$, where $\sim$ is a new binary relation symbol, as follows. For all $(a_1, b_1), \ldots, (a_m, b_m) \in D^2$ with types $t_1, \ldots, t_m$ in $(\Delta, c_1), \ldots, (\Delta, c_m)$, respectively, if the 2-type $\sigma(t_1, \ldots, t_m)$ contains $x_1 = x_2$ then we set $(a_1, \ldots, a_m) \sim (b_1, \ldots, b_m)$. Note that since $n \geq 3$ and because of the compatibility constraints and transitivity of equality, $\sim$ then denotes an equivalence relation on D^m. The other relations of Π are defined as follows. Let R be a k-ary relation from τ. We add the k-tuple $((a_1^1, \ldots, a_m^1), \ldots, (a_1^k, \ldots, a_m^k))$ to the relation R of Π if and only if $R(x_1, \ldots, x_k)$ is contained in $\sigma(t_1, \ldots, t_m)$, where

t_i is the type of the tuple $(a_i^1, \ldots, a_i^k) \in D^k$ in (Δ, c_i). Since $n \geq n(\Delta) \geq k$, this is well-defined by the compatibility item of our algorithm.

The quotient structure $\Pi/_\sim$ is defined to be the τ-structure whose domain is the set $D^m/_\sim$ of all equivalence classes of $\sim$, and where $R(E_1, \ldots, E_p)$ holds for a p-ary $R \in \tau$ and $E_1, \ldots, E_p \in D/_\sim$ if and only if there are $b_1 \in E_1, \ldots, b_p \in E_p$ such that $R(b_1, \ldots, b_p)$ holds in Π. The final step is to show that there exists an embedding f of $\Pi/_\sim$ into Δ. By ω-categoricity of Δ and a standard compactness argument (see, e.g., Lemma 2 in [4]), it suffices to show every finite substructure Ω of $\Pi/_\sim$ embeds into Δ. This follows from the fact that none of the forbidden substructures embeds into Ω, since $n \geq s$, where s is the size of the largest obstruction.

Finally, observe that the mapping g from D^m to D that maps every u in D^m to $f(u/_\sim)$ (where $u/_\sim$ denotes the $\sim$-equivalence class of u in Π) is a polymorphism of Θ by the preservation item of the algorithm, and that g violates R by the violation item of the algorithm. ⊣

Note that our algorithm is even uniform in Δ, that is, we could make Δ part of the input by specifying the finitely many forbidden substructures in the signature of Δ, and quantifier-free definitions of the relations of Γ in Δ.

§**6. Application illustration: Allen's Interval Algebra.** Our algorithm can be applied in many situations that are of interest in temporal and spatial reasoning. In this section we illustrate this by verifying the assumptions of our decidability result in a well-known setting from temporal reasoning, known as *Allen's Interval Algebra*.

Many computational problems studied in temporal and spatial reasoning can be formulated as CSPs for ω-categorical structures. One of the examples is the network satisfaction problem for Allen's Interval Algebra, and for all its fragments. Allen's Interval Algebra can be viewed as a relational structure Γ with domain $D = \{(a, b) \in \mathbb{Q}^2 \mid a < b\}$ and which contains a binary relation $\{((a, b), (c, d)) \mid (\mathbb{Q}; <) \models \phi(a, b, c, d)\}$ for every first-order formula ϕ; we refer to [10] for a formal definition of relation algebras, and their use in temporal and spatial reasoning. A *fragment* of Allen's Interval Algebra is a binary structure obtained from Γ by dropping some of the relations of Γ. Clearly, Γ and all its fragments have finitely many orbits of n-tuples, for all n, and are therefore ω-categorical by the theorem of Ryll-Nardzewski.

We verify that Γ is ordered, homogeneous, Ramsey, and finitely bounded. Theorem 1 then shows that primitive positive definability in fragments of Allen's Interval Algebra is decidable. Note that the lexicographic ordering on D is first-order definable over $(\mathbb{Q}; <)$, and therefore Γ is ordered. It has been observed by Hirsch (Corollary 18 in [14]) that Γ is even homogeneous. The example given after Corollary 45 in [8] discusses how to derive the fact that Γ is Ramsey from the fact that $(\mathbb{Q}; <)$ is Ramsey. We finally have to show the following.

PROPOSITION 27. *Allen's Interval Algebra* Γ *is finitely bounded.*

PROOF. Let τ be the signature of Γ. Let $\mathscr{F}$ be the set of all 3-element τ-structures that do not embed into Γ. We claim that the age of Γ is precisely Forb($\mathscr{F}$), i.e., the set of finite τ-structures that do not embed any of the structures from $\mathscr{F}$. Clearly, the age of Γ is contained in Forb($\mathscr{F}$). Conversely, let Λ be a finite τ-structure that does not embed into Γ.

We construct a structure Λ' with signature $\leq$ from Λ in the following way. For each element v of Λ we have two elements v_1 and v_2 in Λ'. If $R(u, v)$ holds in Λ for some $R \in \tau$, let $\phi(a, b, c, d)$ be the first-order formula that defines R over $(\mathbb{Q}; <)$. If ϕ implies $a \leq c$ over $(\mathbb{Q}; <)$ then we set $u_1 \leq v_1$ in Λ'. Similarly, if ϕ implies $a \leq d$, $b \leq c$, or $b \leq d$, then we add $u_1 \leq v_2$, $u_2 \leq v_1$, or $u_2 \leq v_2$ to Λ', respectively. Observe that when x, y are elements of Λ' such that $x \leq y$ and $y \leq x$ holds in Λ', then for any other element z of Λ' we have $x \leq z$ if and only if $y \leq z$, and $z \leq x$ if and only if $z \leq y$. Let Λ'' be the structure obtained from Λ' by successively removing elements x when there exists another vertex y such that $x \leq y$ and $y \leq x$.

We claim that Λ'' does not embed into $(\mathbb{Q}; \leq)$. Suppose otherwise that h is such an embedding. Let $g : \Lambda \to D$ be defined by $(u, v) \mapsto (h(u), h(v))$; one can verify that g is an embedding of Λ into Γ, a contradiction to our assumptions.

Since $(\mathbb{Q}; \leq)$ is finitely bounded by structures of size at most three, and Λ'' does not embed into $(\mathbb{Q}; \leq)$, there exists a structure Δ'' of size at most three that embeds into Λ'' but does not embed into $(\mathbb{Q}; \leq)$. For each element x of Δ'' there must be a y such that $(x, y) \in D$ or $(y, x) \in D$. Let Δ be the structure induced by those pairs; since Δ'' has at most three elements, the same holds for Δ. Then the structure Δ is in $\mathscr{F}$ since it does not embed into Γ, but it embeds into Λ, which is what we wanted to show. ⊣

Corollary 28. *For Allen's Interval Algebra* Γ *the problem* $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$ *is decidable.*

§7. Discussion and open problems. We presented an algorithm that decides primitive positive definability in finite relational signature reducts Γ of structures that are ordered, Ramsey, homogeneous, finitely bounded, and with finite relational signature. All of those structures Γ are ω-categorical. While the condition for Γ might appear rather restrictive at first sight, it is actually quite general: we want to point out that we do not require that Γ is Ramsey, we only require that Γ is definable in a Ramsey structure. We do not know of a single homogeneous structure Γ with finite relational signature which is *not* a reduct of an ordered homogeneous Ramsey structure with finite relational signature.

Problem 29. *Does every structure which is homogeneous in a finite relational signature have a homogeneous expansion by finitely many relations such that the resulting structure is Ramsey*?

A variant of this problem is the following.

Problem 30. *Does every ω-categorical structure have an ω-categorical expansion which is Ramsey*?

Note that our method is non-constructive: the algorithm does not produce a primitive positive definition in case that there is one. It is an interesting open problem to come up with bounds on the number of existential variables that suffice for a primitive positive definition of R in Θ. For many structures Γ of practical interest, such as $(\mathbb{Q}; <)$ or the random graph, our algorithm can certainly be tuned so that $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$ becomes feasible for reasonable input size; in particular, the gigantic Ramsey constants involved in the proofs of our results do not affect the running time of our procedure.

Another important open problem is whether the method can be extended to show decidability of our computational problem for *first-order* definability instead of primitive positive, existential positive, and existential definability; we denote this computational problem by $\mathrm{Expr}_{\mathrm{pp}}(\Gamma)$. By the theorem of Ryll-Nardzewski, first-order definability is characterized by preservation under automorphisms, i.e., surjective self-embeddings. But the requirement of surjectivity is difficult to deal with in our approach.

PROBLEM 31. *Let Δ be a structure which is ordered, homogeneous, Ramsey, finitely bounded, and has a finite relational signature, and let Γ be a reduct of Δ with finite relational signature. Is the problem $\mathrm{Expr}_{\mathrm{fo}}(\Gamma)$ decidable?*

Acknowledgements. We would like to thank the referees for their comments.

REFERENCES

[1] SAMSON ADEPOJU ADELEKE and PETER M. NEUMANN, ***Relations related to betweenness: their structure and automorphisms***, vol. 131, Memoirs of the American Mathematical Society, no. 623, 1998.

[2] JAMES F. ALLEN, *Maintaining knowledge about temporal intervals*, ***Communications of the ACM***, vol. 26 (1983), no. 11, pp. 832–843.

[3] MANUEL BODIRSKY, HUBIE CHEN, and MICHAEL PINSKER, *The reducts of equality up to primitive positive interdefinability*, this JOURNAL, vol. 75 (2010), no. 4, pp. 1249–1292.

[4] MANUEL BODIRSKY and VÍCTOR DALMAU, *Datalog and constraint satisfaction with infinite templates*, ***Journal on Computer and System Sciences***, vol. 79 (2013), pp. 79–100, A preliminary version appeared in the ***Proceedings of the Symposium on Theoretical Aspects of Computer Science*** (***STACS'05***).

[5] MANUEL BODIRSKY and JAN KÁRA, *The complexity of temporal constraint satisfaction problems*, ***Journal of the ACM***, vol. 57 (2009), no. 2, pp. 1–41, an extended abstract appeared in the ***Proceedings of the Symposium on Theory of Computing*** (***STOC'08***).

[6] MANUEL BODIRSKY and JAROSLAV NEŠETŘIL, *Constraint satisfaction with countable homogeneous templates*, ***Journal of Logic and Computation***, vol. 16 (2006), no. 3, pp. 359–373.

[7] MANUEL BODIRSKY and DIANA PIGUET, *Finite trees are Ramsey with respect to topological embeddings*, preprint, `arXiv:1002.1557`, 2010.

[8] MANUEL BODIRSKY and MICHAEL PINSKER, *Reducts of Ramsey structures*, ***Model theoretic methods in finite combinatorics***, Contemporary Mathematics, vol. 558, AMS, 2011, pp. 489–519.

[9] ———, *Minimal functions on the random graph*, ***Israel Journal of Mathematics***, to appear, preprint `arXiv.org/abs/1003.4030`.

[10] IVO DUENTSCH, *Relation algebras and their application in temporal and spatial reasoning*, ***Artificial Intelligence Review***, vol. 23 (2005), pp. 315–357.

[11] MARTIN GOLDSTERN and MICHAEL PINSKER, *A survey of clones on infinite sets*, ***Algebra Universalis***, vol. 59 (2008), pp. 365–403.

[12] RON L. GRAHAM, BRUCE L. ROTHSCHILD, and JOEL H. SPENCER, ***Ramsey theory***, second ed., Wiley-Interscience Series in Discrete Mathematics and Optimization, John Wiley & Sons, Inc., New York, 1990.

[13] C. WARD HENSON, *Countable homogeneous relational systems and categorical theories*, this JOURNAL, vol. 37 (1972), pp. 494–500.

[14] R. HIRSCH, *Expressive power and complexity in algebraic logic*, ***Journal of Logic and Computation***, vol. 7 (1997), no. 3, pp. 309–351.

[15] WILFRID HODGES, ***A shorter model theory***, Cambridge University Press, Cambridge, 1997.

[16] ALEXANDER KECHRIS, VLADIMIR PESTOV, and STEVO TODORCEVIC, *Fraissé limits, Ramsey theory, and topological dynamics of automorphism groups*, ***Geometric and Functional Analysis***, vol. 15 (2005), no. 1, pp. 106–189.

[17] DUGALD MACPHERSON, *A survey of homogeneous structures*, ***Discrete Mathematics***, vol. 311 (2011), no. 15, pp. 1599–1634.

[18] KEITH R. MILLIKEN, *A Ramsey theorem for trees*, ***Journal of Combinatorial Theory, Series A***, vol. 26 (1979), no. 3, pp. 215–237.

[19] JAROSLAV NEŠETŘIL, *Ramsey classes and homogeneous structures*, ***Combinatorics, Probability & Computing***, vol. 14 (2005), no. 1–2, pp. 171–189.

[20] JAROSLAV NEŠETŘIL and VOJTĚCH RÖDL, *The partite construction and Ramsey set systems*, ***Discrete Mathematics***, vol. 75 (1989), no. 1–3, pp. 327–334.

[21] MICHAEL PINSKER, *More sublattices of the lattice of local clones*, ***Order***, vol. 27 (2010), no. 3, pp. 353–364.

[22] SIMON THOMAS, *Reducts of the random graph*, this JOURNAL, vol. 56 (1991), no. 1, pp. 176–181.

[23] ROSS WILLARD, *Testing expressibility is hard*, ***Principles and practice of constraint programming - CP 2010*** (David Cohen, editor), Lecture Notes in Computer Science, vol. 6308, Springer, 2010, pp. 9–23.

LABORATOIRE D'INFORMATIQUE (LIX)
CNRS UMR 7161
ÉCOLE POLYTECHNIQUE
91128 PALAISEAU, FRANCE
E-mail: bodirsky@lix.polytechnique.fr
URL: http://www.lix.polytechnique.fr/~bodirsky/

ÉQUIPE DE LOGIQUE MATHÉMATIQUE
UNIVERSITÉ DIDEROT—PARIS 7
UFR DE MATHÉMATIQUES
75205 PARIS CEDEX 13, FRANCE
E-mail: marula@gmx.at
URL: http://dmg.tuwien.ac.at/pinsker/
E-mail: todor@math.jussieu.fr
URL: http://people.math.jussieu.fr/~todor/