REDUCTS OF RAMSEY STRUCTURES

MANUEL BODIRSKY AND MICHAEL PINSKER

ABSTRACT. One way of studying a relational structure is to investigate functions which are related to that structure and which leave certain aspects of the structure invariant. Examples are the automorphism group, the self-embedding monoid, the endomorphism monoid, or the polymorphism clone of a structure. Such functions can be particularly well understood when the relational structure is countably infinite and has a first-order definition in another relational structure which has a finite language, is totally ordered and homogeneous, and has the Ramsey property. This is because in this situation, Ramsey theory provides the combinatorial tool for analyzing these functions – in a certain sense, it allows to represent such functions by functions on finite sets.

This is a survey of results in model theory and theoretical computer science obtained recently by the authors in this context. In model theory, we approach the problem of classifying the reducts of countably infinite ordered homogeneous Ramsey structures in a finite language, and certain decidability questions connected with such reducts. In theoretical computer science, we use the same combinatorial methods in order to classify the computational complexity for various classes of infinite-domain constraint satisfaction problems. While the first set of applications is obviously of an infinitary character, the second set concerns genuinely finitary problems – their unifying feature is that the same tools from Ramsey theory are used in their solution.

CONTENTS

1

1. Introduction

"I prefer finite mathematics much more than infinite mathematics. I think that it is much more natural, much more appealing and the theory is much more beautiful. It is very concrete. It is something that you can touch and something you can feel and something to relate to. Infinity mathematics, to me, is something that is meaningless, because it is abstract nonsense."
(Doron Zeilberger, February 2010)

"To the person who does deny infinity and says that it doesn't exist, I feel sorry for them, I don't see how such view enriches the world. Infinity may be does not exist, but it is a beautiful subject. I can say that the stars do not exist and always look down, but then I don't see the beauty of the stars. Until one has a real reason to doubt the existence of mathematical infinity, I just don't see the point."
(Hugh Woodin, February 2010)

Sometimes, infinite mathematics is not just beautiful, but also *useful*, even when one is ultimately interested in finite mathematics. A fascinating example of this type of mathematics is the recent theorem by Kechris, Pestov, and Todorcevic [32], which links Ramsey classes and topological dynamics. A class of finite structures $\mathcal{C}$ closed under isomorphisms, induced substructures, and with the joint embedding property (see [28]) is called a *Ramsey class* [38,39] (or *has the Ramsey property*) if for all $P, H \in \mathcal{C}$ and every $k \geq 2$ there is a $S \in \mathcal{C}$ such that for every coloring of the copies of P in S with k colors there is a copy H' of H in $\mathcal{C}$ such that all copies of P in H' have the same color. This is a very strong requirement — and certainly from the finite world. Proving that a class has the Ramsey property can be difficult [38], and Ramsey theory rather provides a tool box than a theory to answer this question.

Kechris, Pestov, and Todorcevic [32] provide a characterization of such classes in topological dynamics, connecting Ramsey classes with *extreme amenability* in (infinite) group theory, a concept from the 1960s [27]. The result can be used in two directions. One can use it to translate deep existing Ramsey results into proofs of extreme amenability of topological groups (and this is the main focus of the already cited article [32]). One can also use it in the other direction to obtain a more systematic understanding of Ramsey classes. A key insight for this direction is the result of Nešetřil (see [39]) which says that Ramsey classes $\mathcal{C}$ have the *amalgamation property*. Hence, by Fraïssé's theorem, there exists a countably infinite homogeneous and ω-categorical structure Γ such that a finite structure is from $\mathcal{C}$ if and only if it embeds into Γ. The structure Γ is unique up to isomorphism, and is called the *Fraïssé limit* of $\mathcal{C}$. Now let $\mathcal{D}$ be any amalgamation class whose Fraïssé limit Δ is bi-interpretable with Γ. By the theorem of Ahlbrandt and Ziegler [3], two ω-categorical structures are first-order bi-interpretable if and only if their automorphism groups are isomorphic as *(abstract) topological groups*. In addition, the above-mentioned result from [32] shows that whether or not $\mathcal{D}$ is a Ramsey class only depends on the automorphism group $\mathrm{Aut}(\Delta)$ of Δ; in fact, and much more

interestingly, it only depends on $\mathrm{Aut}(\Delta)$ viewed as a topological group (which has cardinality 2^ω). From this we immediately get our first example where [32] is used in the second direction, with massive consequences for finite structures: the Ramsey property is preserved under first-order bi-interpretations. We will see another statement of this type (Proposition 24) and more concrete applications of such statements later (in Section 5, Section 7, and Section 8).

Constraint Satisfaction. Our next example where infinite mathematics is a powerful tool comes from (finite) computer science. A *constraint satisfaction problem* is a computational problem where we are given a set of variables and a set of constraints on those variables, and where the task is to decide whether there is an assignment of values to the variables that satisfies all constraints. Computational problems of this type appear in many areas of computer science, for example in artificial intelligence, computer algebra, scheduling, computational linguistics, and computational biology.

As an example, consider the BETWEENNESS problem. The input to this problem consists of a finite set of variables V, and a finite set of triples of the form (x, y, z) where $x, y, z \in V$. The task is to find an ordering $<$ on V such that for each of the given triples (x, y, z) we have either $x < y < z$ or $z < y < x$. It is well-known that this problem is NP-complete [24, 43], and that we therefore cannot expect to find a polynomial-time algorithm that solves it. In contrast, when we want to find an ordering $<$ on V such that for each of the given triples (x, y, z) we have $x < y$ or $x < z$, then the corresponding problem can be solved in polynomial time.

Many constraint satisfaction problems can be modeled formally as follows. Let Γ be a structure with a finite relational signature. Then the *constraint satisfaction problem for* Γ, denoted by $\mathrm{CSP}(\Gamma)$, is the problem of deciding whether a given primitive positive sentence ϕ is true in Γ. By choosing Γ appropriately, many problems in the above mentioned application areas can be expressed as $\mathrm{CSP}(\Gamma)$. The BETWEENNESS problem, for instance, can be modeled as $\mathrm{CSP}((\mathbb{Q}; Betw))$ where $\mathbb{Q}$ are the rational numbers and $Betw = \{(x, y, z) \in \mathbb{Q}^3 \mid x < y < z \vee z < y < x\}$.

Note that even though the structure Γ might be infinite, the problem $\mathrm{CSP}(\Gamma)$ is always a well-defined and *discrete* problem. Since the signature of Γ is finite, the complexity of $\mathrm{CSP}(\Gamma)$ is independent of the representation of the relation symbols of Γ in input instances of $\mathrm{CSP}(\Gamma)$. The task is to decide whether there *exists* an assignment to the variables of a given instance, and we do not have to exhibit such a solution. Therefore, the computational problems under consideration are finitistic and concrete even when the domain of Γ is, say, the real numbers.

There are many reasons to formulate a discrete problem as $\mathrm{CSP}(\Gamma)$ for an infinite structure Γ. The advantages of such a formulation are most striking when Γ can be chosen to be ω-categorical. In this case, the computational complexity of $\mathrm{CSP}(\Gamma)$ is fully captured by the *polymorphism clone* of Γ; the polymorphism clone can be seen as a higher-dimensional generalization of the automorphism group of Γ. When studying polymorphism clones, we can apply techniques from universal algebra, and, as we will see here, from Ramsey theory to obtain results about the computational complexity of $\mathrm{CSP}(\Gamma)$.

Contributions and Outline. In this article we give a survey presentation of a technique how to apply Ramsey theory when studying automorphism groups, endomorphism monoids, and polymorphism clones of countably infinite structures with a first-order definition in an

ordered homogeneous Ramsey structure in a finite language – such structures are always ω-categorical. We present applications of this technique in two fields. Let Δ be a countable structure with a first-order definition in an ordered homogeneous Ramsey structure in a finite language. In model theory, our technique can be used to classify the set of all structures Γ that are first-order definable in Δ. In constraint satisfaction, it can be used to obtain a complete complexity classification for the class of all problems CSP(Γ) where Γ is first-order definable in Δ. We demonstrate this for $\Delta = (\mathbb{Q}; <)$, and for $\Delta = (V; E)$, the countably infinite random graph.

2. REDUCTS

One way to classify relational structures on a fixed domain is by identifying two structures when they *define* one another. The term "define" will classically stand for "first-order define", i.e., a structure Γ_1 has a first-order definition in a structure Γ_2 on the same domain iff all relations of Γ_1 can be defined by a first-order formula over Γ_2. When Γ_1 has a first-order definition in Γ_2 and vice-versa, then two structures are considered equivalent *up to first-order interdefinability*.

Depending on the application, other notions of definability might be suitable; such notions include syntactic restrictions of first-order definability. In this paper, besides first-order definability, we will consider the notions of *existential positive definability* and *primitive positive definability*; in particular, we will explain the importance of the latter notion in theoretical computer science in Section 8.

The structures which we consider in this article will all be countably infinite, and we will henceforth assume this property without further mentioning it. A structure is called ω-*categorical* if all countable models of its first-order theory are isomorphic. We are interested in the situation where all structures to be classified are *reducts* of a single countable ω-categorical structure in the following sense (which differs from the standard definition of a reduct and morally follows e.g. [46]).

Definition 1. Let Δ be a structure. A *reduct* of Δ is a structure with the same domain as Δ all of whose relations can be defined by a first-order formula in Δ.

When all structures under consideration are reducts of a countably infinite base structure Δ which is ω-categorical, then there are natural ways of obtaining classifications up to first-order, existential positive, or primitive positive interdefinability by means of certain sets of functions. In this section, we explain these ways, and give some examples of classifications that have been obtained in the past. In the following sections, we then observe that these results have actually been obtained in a more specific context than ω-categoricity, namely, where the structures are reducts of an *ordered Ramsey* structure Δ which has a finite relational signature and which is *homogeneous* in the sense that every isomorphism between finite induced substructures of Δ can be extended to an automorphism of Δ. We further develop a general framework for proving such results in this context.

We start with first-order definability. Consider the assignment that sends every structure Γ with domain D to its automorphism group $\mathrm{Aut}(\Gamma)$. Automorphism groups are closed sets in the convergence topology of all permutations on D, and conversely, every closed permutation group on D is the automorphism group of a relational structure with domain D. The closed permutation groups on D form a complete lattice, where the meet of a set of groups is given by their intersection. Similarly, the set of those relational structures on D which are *first-order closed*, i.e., which contain all relations which they define by a first-order formula,

forms a lattice, where the meet of a set S of such structures is the structure which has those relations that are contained in all structures in S. Now when Γ is a countable ω-categorical structure, then it follows from the proof of the theorem of Ryll-Nardzewki (see [28]) that its automorphism group $\mathrm{Aut}(\Gamma)$ still has the first-order theory of Γ encoded in it. And indeed we can, up to first-order interdefinability, recover Γ from its automorphism group as follows: For a set $\mathcal{F}$ of finitary functions on D, let $\mathrm{Inv}(\mathcal{F})$ be the structure on D which has those relations R which are *invariant under* $\mathcal{F}$, i.e., those relations that contain $f(r_1, \ldots, r_n)$ (calculated componentwise) whenever $f \in \mathcal{F}$ and $r_1, \ldots, r_n \in R$.

Theorem 2. *Let Δ be ω-categorical. Then the mapping $\Gamma \mapsto \mathrm{Aut}(\Gamma)$ is an antiisomorphism between the lattice of first-order closed reducts of Δ and the lattice of closed permutation groups containing $\mathrm{Aut}(\Delta)$. The inverse mapping is given by $\mathcal{G} \mapsto \mathrm{Inv}(\mathcal{G})$.*

This connection between closed permutation groups and first-order definability has been exploited several times in the past in order to obtain complete classifications of reducts of ω-categorical structures. For example, let Δ be the order of the rational numbers – we write $\Delta = (\mathbb{Q}; <)$. Then it has been shown in [20] that there are exactly five reducts of Δ, up to first-order interdefinability, which we will define in the following.

On the permutation side, let $\leftrightarrow$ be the function that sends every $x \in \mathbb{Q}$ to $-x$. For our purposes, we can equivalently choose $\leftrightarrow$ to be any permutation that inverts the order $<$ on $\mathbb{Q}$. For any fixed irrational real number α, let $\circlearrowleft$ be any permutation on $\mathbb{Q}$ with the property that $x < y < \alpha < u < v$ implies $\circlearrowleft(u) < \circlearrowleft(v) < \circlearrowleft(x) < \circlearrowleft(y)$, for all $x, y, u, v \in \mathbb{Q}$. We will consider closed groups *generated* by these permutations: For a set of permutations $\mathcal{F}$ and a closed permutation group $\mathcal{G}$, we say that $\mathcal{F}$ generates $\mathcal{G}$ iff $\mathcal{G}$ is the smallest closed group containing $\mathcal{F}$.

On the relational side, for $x_1, \ldots, x_n \in \mathbb{Q}$ write $\overrightarrow{x_1 \ldots x_n}$ when $x_1 < \ldots < x_n$. Then we define a ternary relation *Betw* on $\mathbb{Q}$ by $Betw := \{(x, y, z) \in \mathbb{Q}^3 \mid \overrightarrow{xyz} \vee \overrightarrow{zyx}\}$. Define another ternary relation *Cycl* by $Cycl := \{(x, y, z) \in \mathbb{Q}^3 \mid \overrightarrow{xyz} \vee \overrightarrow{yzx} \vee \overrightarrow{zxy}\}$. Finally, define a 4-ary relation *Sep* by

$$\{(x_1, y_1, x_2, y_2) \in \mathbb{Q}^4 \mid \overrightarrow{x_1 x_2 y_1 y_2} \vee \overrightarrow{x_1 y_2 y_1 x_2} \vee \overrightarrow{y_1 x_2 x_1 y_2} \vee \overrightarrow{y_1 y_2 x_1 x_2}$$
$$\vee \overrightarrow{x_2 x_1 y_2 y_1} \vee \overrightarrow{x_2 y_1 y_2 x_1} \vee \overrightarrow{y_2 x_1 x_2 y_1} \vee \overrightarrow{y_2 y_1 x_2 x_1}\}.$$

Theorem 3 (Cameron [20]). *Let Γ be a reduct of $(\mathbb{Q}; <)$. Then exactly one of the following holds:*

- *Γ is first-order interdefinable with $(\mathbb{Q}; <)$; equivalently, $\mathrm{Aut}(\Gamma) = \mathrm{Aut}((\mathbb{Q}; <))$.*
- *Γ is first-order interdefinable with $(\mathbb{Q}; Betw)$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the closed group generated by $\mathrm{Aut}((\mathbb{Q}; <))$ and $\leftrightarrow$.*
- *Γ is first-order interdefinable with $(\mathbb{Q}; Cycl)$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the closed group generated by $\mathrm{Aut}((\mathbb{Q}; <))$ and $\circlearrowleft$.*
- *Γ is first-order interdefinable with $(\mathbb{Q}; Sep)$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the closed group generated by $\mathrm{Aut}((\mathbb{Q}; <))$ and $\{\leftrightarrow, \circlearrowleft\}$.*
- *Γ is first-order interdefinable with $(\mathbb{Q}; =)$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the group of all permutations on $\mathbb{Q}$.*

Another instance of the application of Theorem 2 in the classification of reducts up to first-order interdefinability has been provided by Thomas [46]. Let $G = (V; E)$ be the random graph, i.e., the up to isomorphism unique countably infinite graph which is homogeneous and

which contains all finite graphs as induced subgraphs. It turns out that up to first-order interdefinability, G has precisely five reducts, too.

On the permutation side, observe that the graph $\bar{G}$ obtained by making two distinct vertices $x, y \in V$ adjacent iff they are not adjacent in G is isomorphic to G; let $-$ be any permutation on V witnessing this isomorphism. Moreover, for any fixed vertex $0 \in V$, the graph obtained by making all vertices which are adjacent with 0 non-adjacent with 0, and all vertices different from 0 and non-adjacent with 0 adjacent with 0, is isomorphic to G. Let sw be any permutation on V witnessing this fact.

On the relational side, define for all $k \geq 2$ a k-ary relation $R^{(k)}$ on V by

$$R^{(k)} := \{(x_1, \ldots, x_k) \mid \text{ all } x_i \text{ are distinct,}$$
$$\text{and the number of edges on } \{x_1, \ldots, x_k\} \text{ is odd}\}.$$

Theorem 4 (Thomas [46]). *Let Γ be a reduct of the random graph $G = (V; E)$. Then exactly one of the following holds:*

- *Γ is first-order interdefinable with G; equivalently, $\mathrm{Aut}(\Gamma) = \mathrm{Aut}(G)$.*
- *Γ is first-order interdefinable with $(V; R^{(3)})$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the closed group generated by $\mathrm{Aut}(G)$ and sw.*
- *Γ is first-order interdefinable with $(V; R^{(4)})$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the closed group generated by $\mathrm{Aut}(G)$ and $-$.*
- *Γ is first-order interdefinable with $(V; R^{(5)})$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the closed group generated by $\mathrm{Aut}(G)$ and $\{\mathrm{sw}, -\}$.*
- *Γ is first-order interdefinable with $(V; =)$; equivalently, $\mathrm{Aut}(\Gamma)$ equals the group of all permuations on V.*

In a similar fashion, the reducts of several prominent ω-categorical structures Δ have been classified up to first-order interdefinability by finding all closed supergroups of $\mathrm{Aut}(\Delta)$. Examples are:

- The countable homogeneous K_n-free graph, i.e., the unique countable homogeneous graph which contains precisely those finite graphs which do not contain a clique of size n as induced subgraphs, has 2 reducts up to first-order interdefinability (Thomas [46]), for all $n \geq 3$.
- The countable homogeneous k-hypergraph has $2^k + 1$ reducts up to first-order interdefinability (Thomas [47]), for all $k \geq 2$.
- The structure $(\mathbb{Q}; <, 0)$, i.e., the order of the rationals which in addition "knows" one of its points, has 116 reducts up to first-order interdefinability (Junker and Ziegler [31]).

All these examples have in common that the structures have a high degree of symmetry in the sense that they are homogeneous in a finite language – intuitively, one would expect the automorphism group of such a structure to be rather large. And indeed, Thomas conjectured in [46]:

Conjecture 5 (Thomas [46]). *Let Δ be a countable relational structure which is homogeneous in a finite language. Then Δ has finitely many reducts up to first-order interdefinability.*

It turns out that all the examples above are not only homogeneous in a finite language; in fact, they all have a first-order definition in (in other words: are themselves reducts of) an *ordered Ramsey structure* which is homogeneous in a finite language. Functions on such structures, in particular automorphisms of reducts, can be analyzed by the means of Ramsey

theory, and we will outline a general method for classifying the reducts of such structures in Sections 3 to 6.

We now turn to analogs of Theorem 2 for syntactic restrictions of first-order logic. A first-order formula is called *existential* iff it is of the form $\exists x_1 \dots \exists x_n.\ \phi$, where ϕ is quantifier-free. It is called *existential positive* iff it is existential and does not contain any negations. Now observe that similarly to permutation groups, the *endomorphism monoid* $\mathrm{End}(\Delta)$ of a relational structure Δ with domain D is always closed in the pointwise convergence topology on the space of all functions from D to D, and that every closed transformation monoid $\mathcal{M}$ acting on D is the endomorphism monoid of the structure $\mathrm{Inv}(\mathcal{M})$, i.e., the structure with domain D which contains those relations which are invariant under all functions in $\mathcal{M}$. Note also that the set of closed transformation monoids on D, ordered by inclusion, forms a complete lattice, and that likewise the set of all existential positive closed structures forms a complete lattice. The analog to Theorem 2 for existential positive definability is an easy consequence of the homomorphism preservation theorem (see [28]) and goes like this:

Theorem 6. *Let Δ be ω-categorical. Then the mapping $\Gamma \mapsto \mathrm{End}(\Gamma)$ is an antiisomorphism between the lattice of existential positive closed reducts of Δ and the lattice of closed transformation monoids containing $\mathrm{Aut}(\Delta)$. The inverse mapping is given by $\mathcal{M} \mapsto \mathrm{Inv}(\mathcal{M})$.*

All the closed monoids containing the group of all permutations on a countably infinite set D (which equals the automorphism group of the empty structure $(D; =)$) have been determined in [8], and their number is countably infinite. Therefore, every structure has infinitely many reducts up to existential positive interdefinability. In general, it will be impossible to determine all of them, but sometimes it is already useful to determine certain closed monoids, as in the following theorem about endomorphism monoids of reducts of the random graph from [14]. We need the following definitions. Since the random graph $G = (V; E)$ contains all countable graphs, it contains an infinite clique. Let e_E be any injective function from V to V whose image induces such a clique in G. Similarly, let e_N be any injection from V to V whose image induces an independent set in G.

Theorem 7 (Bodirsky and Pinsker [14]). *Let Γ be a reduct of the random graph $G = (V; E)$. Then at least one of the following holds.*

- $\mathrm{End}(\Gamma)$ *contains a constant operation.*
- $\mathrm{End}(\Gamma)$ *contains e_E.*
- $\mathrm{End}(\Gamma)$ *contains e_N.*
- $\mathrm{Aut}(\Gamma)$ *is a dense subset of $\mathrm{End}(\Gamma)$ (equipped with the topology of pointwise convergence).*

Theorem 7 states that for reducts Γ of the random graph, either $\mathrm{End}(\Gamma)$ contains a function that destroys all structure of the random graph, or it contains basically no functions except the automorphisms. This has the following non-trivial consequence. A theory T is called *model-complete* iff every embedding between models of T is elementary, i.e., preserves all first-order formulas. A structure is said to be model-complete iff its first-order theory is model-complete.

Corollary 8 (Bodirsky and Pinsker [14]). *All reducts of the random graph are model-complete.*

Proof. It is not hard to see (cf. [14]) that an ω-categorical structure Γ is model-complete if and only if $\mathrm{Aut}(\Gamma)$ is dense in the monoid of self-embeddings of Γ. Now let Γ be a reduct of G, and let $\mathcal{M}$ be the closed monoid of self-embeddings of Γ; we will show that $\mathrm{Aut}(\Gamma)$ is dense

in $\mathcal{M}$. We apply Theorem 7 to $\mathcal{M}$ (which, as a closed monoid containing $\mathrm{Aut}(G)$, is also an endomorphism monoid of a reduct Γ' of G). Clearly, Γ' and Γ have the same automorphisms, namely those permutations in $\mathcal{M}$ whose inverse is also in $\mathcal{M}$. Therefore we are done if the last case of the theorem holds. Note that $\mathcal{M}$ cannot contain a constant operation as all its operations are injective. So suppose that $\mathcal{M}$ contains e_N – the argument for e_E is analogous. Let R be any relation of Γ, and ϕ_R be its defining quantifier-free formula; ϕ_R exists since G has quantifier-elimination, i.e., every first-order formula over G is equivalent to a quantifier-free formula. Let ψ_R be the formula obtained by replacing all occurrences of E by *false*; so ψ_R is a formula over the empty language. Then a tuple a satisfies ϕ_R in G iff $e_N(a)$ satisfies ϕ_R in G (because e_N is an embedding) iff $e_N(a)$ satisfies ψ_R in G (as there are no edges on $e_N(a)$) iff $e_N(a)$ satisfies ψ_R in the substructure induced by $e_N[V]$ (since ψ_R does not contain any quantifiers). Thus, Γ is isomorphic to the structure on $e_N[V]$ which has the relations defined by the formulas ψ_R; hence, Γ is isomorphic to a structure with a first-order definition over the empty language. This structure has, of course, all injections as self-embeddings, and all permutations as automorphisms, and hence is model-complete; thus, the same is true for Γ. $\square$

It follows from [11, Proposition 19] that all reducts of the linear order of the rationals $(\mathbb{Q}; <)$ are model-complete as well. This is remarkable, since similar structures do not have this property – for example, $(\mathbb{Q}; <, 0)$ is first-order interdefinable with the structure $(\mathbb{Q}; <, [0, \infty))$ which is not model-complete.

We now turn to an even finer way of distinguishing reducts of an ω-categorical structure, namely up to *primitive positive interdefinability*. This is of importance in connection with the constraint satisfaction problem from the introduction, as we will describe in more detail in Section 8. We call a formula *primitive positive* iff it is existential positive and does not contain disjunctions. A *clone* on domain D is a set of finitary operations on D which contains all projections (i.e., functions of the form $(x_1, \ldots, x_n) \mapsto x_i$) and which is closed under composition. A clone $\mathcal{C}$ is *closed* (also called *locally closed* or *local* in the literature) iff for each $n \geq 1$, the set of n-ary functions in $\mathcal{C}$ is a closed subset of the space D^{D^n}, where D is taken to be discrete. The closed clones on D form a complete lattice with respect to inclusion – the structure of this lattice has been studied in the universal algebra literature (see [26], [44]). Similarly, the set of relational structures with domain D which are *primitive positive closed*, i.e., which contain all relations which they define by primitive positive formulas, forms a complete lattice. For a structure Γ, we define $\mathrm{Pol}(\Gamma)$ to consist of all finitary operations on the domain of Γ which preserve all relations of Γ, i.e., an n-ary function f is an element of $\mathrm{Pol}(\Gamma)$ iff for all relations R of Γ and all tuples $r_1, \ldots, r_n \in R$ the tuple $f(r_1, \ldots, r_n)$ is an element of R. It is easy to see that $\mathrm{Pol}(\Gamma)$ is always a closed clone. Observe also that $\mathrm{Pol}(\Gamma)$ is a generalization of $\mathrm{End}(\Gamma)$ to higher (finite) arities.

Theorem 9 (Bodirsky and Nešetřil [12]). *Let Δ be ω-categorical. Then the mapping $\Gamma \mapsto \mathrm{Pol}(\Gamma)$ is an antiisomorphism between the lattice of primitive positive closed reducts of Δ and the lattice of closed clones containing $\mathrm{Aut}(\Delta)$. The inverse mapping is given by $\mathcal{C} \mapsto \mathrm{Inv}(\mathcal{C})$.*

It turns out that even for the empty structure $(X; =)$, the lattice of primitive positive closed reducts is probably too complicated to be completely described – the lattice has been thoroughly investigated in [8].

Theorem 10 (Bodirsky, Chen, Pinsker [8])**.** *The structure* $(X; =)$ *(where X is countably infinite), and therefore all countably infinite structures, have $2^{\aleph_0}$ reducts up to primitive positive interdefinability.*

Fortunately, it is sometimes sufficient in applications to understand only parts of this lattice. We will see examples of this in Section 8.

3. Ramsey Classes

While Theorems 2, 6 and 9 provide a theoretical method for determining reducts of an ω-categorical structure Δ by transforming them into sets of functions on Δ, understanding these infinite objects could turn out difficult without further tools for handling them. We will now focus on structures which have the additional property that they are reducts of an *ordered Ramsey structure* that is homogeneous in a finite relational signature; such structures are ω-categorical since homogeneous structures in a finite language are ω-categorical and since reducts of ω-categorical structures are ω-categorical. This is less restrictive than it might appear at first sight: we remark that it could be the case that all homogeneous structures with a finite relational signature are reducts of ordered homogeneous Ramsey structures with a finite relational signature (that is, we do not know of a counterexample). It turns out that in this context, certain infinite functions can be represented by finite ones, making classification projects more feasible.

Definition 11. A structure is called *ordered* iff it has a total order among its relations.

Definition 12. Let τ a relational signature. For τ-structures $\mathcal{S}, \mathcal{H}, \mathcal{P}$ and an integer $k \geq 1$, we write $\mathcal{S} \to (\mathcal{H})_k^{\mathcal{P}}$ iff for every k-coloring χ of the copies of $\mathcal{P}$ in $\mathcal{S}$ there exists a copy $\mathcal{H}'$ of $\mathcal{H}$ in $\mathcal{S}$ such that all copies of $\mathcal{P}$ in $\mathcal{H}'$ have the same color under χ.

Definition 13. A class $\mathcal{C}$ of finite τ-structures which is closed under isomorphisms, induced substructures, and with the joint embedding property (see [28]) is called a *Ramsey class* iff it is closed under substructures and for all $k \geq 1$ and all $\mathcal{H}, \mathcal{P} \in \mathcal{C}$ there exists $\mathcal{S}$ in $\mathcal{C}$ such that $\mathcal{S} \to (\mathcal{H})_k^{\mathcal{P}}$.

Definition 14. A relational structure is called *Ramsey* iff its *age*, i.e., the set of finite structures isomorphic to a finite induced substructure, is a Ramsey class.

Examples of Ramsey structures are the dense linear order $(\mathbb{Q}; <)$ and the *ordered random graph* $(V; E, <)$, i.e., the Fraïssé limit of the class of finite ordered graphs. We remark that the random graph itself is not Ramsey, but since it is a reduct of the ordered random graph, the methods we are about to expose apply as well.

We will now see that one can find regular patterns in the behavior of any function acting on an ordered Ramsey structure which is ω-categorical.

Definition 15. Let Γ be a structure. The *type* $\mathrm{tp}(a)$ of an n-tuple $a \in \Gamma$ is the set of first-order formulas with free variables $x_1, \ldots, x_n$ that hold for a in Γ.

We recall the classical theorem of Ryll-Nardzewski about the number of types in ω-categorical structures.

Theorem 16 (Ryll-Nardzewski)**.** *The following are equivalent for a countable structure Γ in a countable language.*

- *Γ is ω-categorical, i.e., any countable model of the theory of Γ is isomorphic to Γ.*

  • Γ *has for all $n \geq 1$ only finitely many different types of n-tuples.*

We also mention that moreover, as a well-known consequence of the proof of this theorem, two tuples in a countable ω-categorical structure have the same type if and only if there is an automorphism of Γ which sends one tuple to the other.

Definition 17. A *type condition* between two structures Γ_1, Γ_2 is a pair (t_1, t_2), where each t_i is a type of an n-tuple in Γ_i. A function $f : \Gamma_1 \to \Gamma_2$ *satisfies* a type condition (t_1, t_2) if for all n-tuples $(a_1, \dots, a_n)$ of type t_1, the n-tuple $(f(a_1), \dots, f(a_n))$ is of type t_2.

A *behavior* is a set of type conditions between two structures. A function *has behavior B* if it satisfies all the type conditions of the behavior B. A behavior B is called *complete* iff for all types t_1 of tuples in Γ_1 there is a type t_2 of a tuple in Γ_2 such that $(t_1, t_2) \in B$.

A function $f : \Gamma_1 \to \Gamma_2$ is *canonical* iff it has a complete behavior. If $F \subseteq \Gamma_1$, then we say that f is *canonical on F* if its restriction to F is canonical.

Observe that the function $\leftrightarrow$ of Theorem 3 is canonical for the structure $(\mathbb{Q}; <)$. The function $\circlearrowleft$ is not, but it is canonical on each of the intervals $(-\infty, \alpha)$ and (α, ∞). For the random graph, the function $-$ of Theorem 4 is canonical, while sw is canonical on $V \setminus \{0\}$. Also, sw is canonical as a function from $(V; E, 0)$ to $(V; E)$, where $(V; E, 0)$ denotes the structure obtained from $(V; E)$ by adding a new constant symbol for the element 0 by which we defined the function sw. Moreover, the constant function and e_E, e_N of Theorem 7 are canonical on $(V; E)$. We will now show that it is no coincidence that canonical functions are that ubiquitous.

Definition 18. Let Δ be a structure. A property P *holds for arbitrarily large finite substructures of Δ* iff for all finite substructures $F \subseteq \Delta$ there is a copy of F in Δ for which P holds.

The following observation is just an easy application of the definition of a Ramsey class, but crucial in understanding functions on ordered Ramsey structures.

Lemma 19. *Let Δ be ordered Ramsey and ω-categorical, and let $f : \Delta \to \Delta$. Then f is canonical on arbitrarily large finite substructures.*

The proof goes along the following lines: Let F be any finite substructure of Δ. Then the function f induces a mapping from the tuples in Δ to the set of types in Δ (each tuple is sent to the type of its image under f). If we restrict this mapping to tuples of length at most the size of F, then since Δ is ω-categorical, the range of this restriction is finite by Theorem 16, and thus is a k-coloring of tuples for some finite k. Now apply the Ramsey property once for every type of tuple that occurs in F – see [16] for details. We remark that this lemma would be false if one dropped the order assumption, which implies that coloring induced substructures and coloring tuples in Δ are one and the same thing.

The motivation for working with ordered Ramsey structures is the rough idea that all "important" functions can be assumed to be canonical. While this is simply false when stated boldly like this, it is still true for some functions when the idea is further refined, as we will show in the following. Observe that if Δ is ω-categorical, then for each $n \geq 1$ there are only finitely many possible type conditions for n-types over Δ (Theorem 16). Suppose that Δ has in addition a finite language and *quantifier elimination*, i.e., every first-order formula in the language of Δ is equivalent to a quantifier-free formula over Δ; this follows in particular from homogeneity in a finite language. Then, if $n(\Delta)$ is the largest arity of its relations, then a function $f : \Delta \to \Delta$ is canonical iff for every type t_1 of an $n(\Delta)$-tuple in Δ there is a type

t_2 in Δ such that f satisfies the type condition (t_1, t_2). In other words, the complete behavior of f is already determined by its behavior on $n(\Delta)$-types. Hence, a canonical function on Δ is essentially a function on the $n(\Delta)$-types of Δ – a finite object.

Definition 20. Let $f, g : \Delta \to \Delta$. We say that f *generates* g over Δ iff g is contained in the smallest closed monoid containing f and $\mathrm{Aut}(\Delta)$. Equivalently, for every finite subset F of Δ, there exists a term $\beta \circ f \circ \alpha_1 \circ f \circ \alpha_2 \circ \cdots \circ f \circ \alpha_n$, where $\beta, \alpha_i \in \mathrm{Aut}(\Delta)$, which agrees with g on F.

Proposition 21. *Let Δ be a structure in a finite language which is ordered, Ramsey, and homogeneous. Let $f : \Delta \to \Delta$. Then f generates a canonical function $g : \Delta \to \Delta$.*

First proof. Let $(F_i)_{i \in \omega}$ be an increasing sequence of finite substructures of Δ such that $\bigcup_{i \in \omega} F_i = \Delta$. By Lemma 19, for each $i \in \omega$ we find a copy F_i' of F_i in Δ on which f is canonical. Since there are only finitely many possibilities of canonical behavior, one behavior occurs an infinite number of times; thus, by thinning out the sequence, we may assume that the behavior is the same on all F_i'. By the homogeneity of Δ, there exist automorphisms α_i of Δ sending F_i to F_i', for all $i \in \omega$. Also, since the behavior on all the F_i' is the same, we can inductively pick automorphisms β_i of Δ 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 $\beta_i \circ f \circ \alpha_i : F_i \to \Delta$ is a canonical function on Δ. $\square$

Second proof. The identity function $\mathrm{id} : \Delta \to \Delta$ is generated by f and is canonical. $\square$

The problem with the preceding lemma is the second proof, which makes it trivial. What we really want is that f generates 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 functions. For example, we could wish that if f violates a certain relation, then so does g; or, if f is not an automorphism of Δ, we will look for a canonical function g which is not an automorphism of Δ either.

We are now going to refine our method, and fix constants $c_1, \ldots, c_n$ such that $f \notin \mathrm{Aut}(\Delta)$ is witnessed on $\{c_1, \ldots, c_n\}$. We then consider f as a function from $(\Delta, c_1, \ldots, c_n)$ to Δ, where $(\Delta, c_1, \ldots, c_n)$ denotes the expansion of Δ by the constants $c_1, \ldots, c_n$. It turns out that f is canonical on arbitrarily large substructures of $(\Delta, c_1, \ldots, c_n)$, and that it generates a canonical function $g : (\Delta, c_1, \ldots, c_n) \to \Delta$ which agrees with f on $c_1, \ldots, c_n$; in particular, g is not an automorphism of Δ, and the problem of triviality in Proposition 20 no longer occurs. 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. Topological Dynamics

We have seen in the previous section that our approach crucially relies on the fact that when an ordered homogeneous Ramsey structure is expanded by finitely many constants, the expansion is again Ramsey (it is clear that the expansion is again ordered and homogeneous). To prove this, we use a characterization in topological dynamics of those ordered homogeneous structures which are Ramsey.

Recall that a *topological group* is an (abstract) group G together with a topology on the elements of G such that $(x, y) \mapsto xy^{-1}$ is continuous from G^2 to G. In other words, we require that the binary group operation and the inverse function are continuous.

Definition 22. A topological group is *extremely amenable* iff any continuous action of the group on a compact Hausdorff space has a fixed point.

Kechris, Pestov and Todorcevic have characterized the Ramsey property of the age of an ordered homogeneous structure by means of extreme amenability in the following theorem.

Theorem 23 (Kechris, Pestov, Todorcevic [32])**.** *Let Δ be an ordered homogeneous relational structure. Then the age of Δ has the Ramsey property iff $\mathrm{Aut}(\Delta)$ is extremely amenable.*

This theorem can be applied to provide a short and elegant proof of the following.

Proposition 24 (Bodirsky, Pinsker and Tsankov [16])**.** *Let Δ be ordered, Ramsey, and homogeneous, and let $c_1, \ldots, c_n \in \Delta$. Then $(\Delta, c_1, \ldots, c_n)$ is Ramsey as well.*

When Δ is ordered, Ramsey, and homogeneous, then $\mathrm{Aut}(\Delta)$ 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 directly from the following fact – confer [16].

Lemma 25. *Let G be an extremely amenable group, and let H be an open subgroup of G. Then H is extremely amenable.*

5. Minimal Functions

The results of the preceding section provide a tool for "climbing up" the lattice of closed monoids containing the automorphism group of an ordered Ramsey structure which is homogeneous and has a finite language.

Definition 26. Let $\mathcal{C}, \mathcal{D}$ be closed clones. Then $\mathcal{D}$ is called *minimal above* $\mathcal{C}$ iff $\mathcal{D} \supseteq \mathcal{C}$ and there are no closed clones between $\mathcal{C}$ and $\mathcal{D}$.

Observe that transformation monoids can be identified with those clones which have the property that all their functions depend on only one variable. Hence, Definition 26 also provides us with a notion of a minimal closed monoid above another closed monoid.

It follows from Theorem 9 and Zorn's Lemma that if Δ is an ω-categorical structure in a finite language, then every closed clone containing $\mathrm{Pol}(\Delta)$ contains a minimal closed clone above $\mathrm{Pol}(\Delta)$. Similarly, as a consequence of Theorem 6, every closed monoid containing $\mathrm{End}(\Delta)$ contains a minimal closed monoid.

For closed permutation groups, minimality can be defined analogously. Then Theorem 2 implies that for ω-categorical structures Δ in a finite language, every closed permutation group containing $\mathrm{Aut}(\Delta)$ contains a minimal closed permutation group above $\mathrm{Aut}(\Delta)$.

Clearly, if a closed clone $\mathcal{D}$ is minimal above $\mathcal{C}$, then any function $f \in \mathcal{D} \setminus \mathcal{C}$ generates $\mathcal{D}$ with $\mathcal{C}$ (i.e., $\mathcal{D}$ is the smallest closed clone containing f and $\mathcal{C}$) – similar statements hold for monoids and groups. In the case of clones and monoids and in the setting of reducts of ordered Ramsey structures which are homogeneous in a finite language, we can standardize such generating functions. This is the contents of the coming subsections.

5.1. **Minimal unary functions.**
Adapting the proof of Lemma 21, with the use of the Proposition 24, one can show the following.

Lemma 27. *Let Δ be ordered, Ramsey, homogeneous, and of finite language. Let $f : \Delta \to \Delta$, and let $c_1, \ldots, c_n \in \Delta$. Then f together with $\mathrm{Aut}(\Delta)$ generates a function which agrees with f on $\{c_1, \ldots, c_n\}$ and which is canonical as a function from $(\Delta, c_1, \ldots, c_n)$ to Δ.*

Let Γ be a finite language reduct of a structure Δ which is ordered, Ramsey, homogeneous, and of finite language, and let $\mathcal{N}$ be a minimal closed monoid containing $\mathrm{End}(\Gamma)$. Then, setting $n(\Gamma)$ to be the largest arity of the relations of Γ, we can pick constants $c_1, \ldots, c_{n(\Gamma)} \in \Gamma$ and a function $f \in \mathcal{N} \setminus \mathrm{End}(\Gamma)$ such that $f \notin \mathrm{End}(\Gamma)$ is witnessed on $\{c_1, \ldots, c_{n(\Gamma)}\}$. By the preceding lemma, f and $\mathrm{Aut}(\Delta)$ generate a function g which behaves like f on $\{c_1, \ldots, c_{n(\Gamma)}\}$ and which is canonical as a function from $(\Delta, c_1, \ldots, c_{n(\Gamma)})$ to Δ. This function g, together with $\mathrm{End}(\Gamma)$, generates $\mathcal{N}$. Since there are only finitely many choices for the type of the tuple $(c_1, \ldots, c_{n(\Gamma)})$ and for each choice only finitely many behaviors of functions from $(\Delta, c_1, \ldots, c_{n(\Gamma)})$ to Δ, we get the following.

Proposition 28 (Bodirsky, Pinsker, Tsankov [16]). *Let Γ be a finite language reduct of a structure Δ which is ordered, Ramsey, homogeneous, and of finite language. Then the number of minimal closed monoids above $\mathrm{End}(\Gamma)$ is finite, and each such monoid is generated by $\mathrm{End}(\Gamma)$ plus a canonical function $g : (\Delta, c_1, \ldots, c_{n(\Gamma)}) \to \Delta$, for constants $c_1, \ldots, c_{n(\Gamma)} \in \Gamma$.*

Since for every relation R of Γ we can add its negation to the language, we get the following

Corollary 29. *Let $\mathcal{M}$ be the monoid of self-embeddings of a finite-language structure Γ which is a reduct of a structure Δ which is ordered, Ramsey, homogeneous, and of finite language. Then the number of minimal closed monoids above $\mathcal{M}$ is finite, and each such monoid is generated by $\mathcal{M}$ and a canonical function $g : (\Delta, c_1, \ldots, c_{n(\Gamma)}) \to \Delta$.*

The following is an example for the random graph $G = (V; E)$. Since G is model-complete, its monoid of self-embeddings is just the topological closure $\langle \mathrm{Aut}(G) \rangle$ of $\mathrm{Aut}(G)$ in the space V^V. Therefore, the minimal closed monoids above the monoid of self-embeddings of G are just the minimal closed monoids above $\langle \mathrm{Aut}(G) \rangle$.

Theorem 30 (Thomas [47]). *Let $G = (V; E)$ be the random graph. The minimal closed monoids containing $\langle \mathrm{Aut}(G) \rangle$ are the following:*

- *The monoid generated by a constant operation with $\mathrm{Aut}(G)$.*
- *The monoid generated by e_E with $\mathrm{Aut}(G)$.*
- *The monoid generated by e_N with $\mathrm{Aut}(G)$.*
- *The monoid generated by $-$ with $\mathrm{Aut}(G)$.*
- *The monoid generated by* sw *with $\mathrm{Aut}(G)$.*

5.2. **Minimal higher arity functions.** We now generalize the concepts from unary functions and monoids to higher arity functions and clones.

Definition 31. Let Δ be a structure. For $1 \leq i \leq m$ and a tuple x in the power Δ^m, we write x_i for the i-th coordinate of x. The *type* of a sequence of tuples $a^1, \ldots, a^n \in \Delta^m$, denoted by $\mathrm{tp}(a^1, \ldots, a^n)$, is the cartesian product of the types of $(a_i^1, \ldots, a_i^n)$ in Δ.

With this definition, the notions of *type condition*, *behavior*, *complete behavior*, and *canonical* generalize in complete analogy from functions $f : \Gamma_1 \to \Gamma_2$ to functions $f : \Gamma_1^m \to \Gamma_2$, for structures Γ_1, Γ_2. It can be shown that for ordered structures, the Ramsey property is not lost when going to products; an example of a proof can be found in [16].

Proposition 32. *Let Δ be ordered and Ramsey, and let $m \geq 1$. Let moreover a number $k \geq 1$, an n-tuple $(a^1, \ldots, a^n) \in \Delta^m$, and finite $F_i \subseteq \Delta$ be given for $1 \leq i \leq m$. Then there exist finite $S_i \subseteq \Delta$ with the property that whenever the n-tuples in $S_1 \times \cdots \times S_m$ of type $\mathrm{tp}(a^1, \ldots, a^n)$ are colored with k colors, then there are copies F_i' of F_i in S_i such that the coloring is constant on $F_1' \times \cdots \times F_m'$.*

We remark that Proposition 32 does not hold in general if Δ is not assumed to be ordered – an example for the random graph can be found in [14]. Similarly to the unary case (Proposition 28), one gets the following.

Proposition 33 (Bodirsky, Pinsker, Tsankov [16])**.** *Let Γ be a finite language reduct of a structure Δ which is ordered, Ramsey, homogeneous and of finite language. Then every minimal closed clone above $\mathrm{Pol}(\Gamma)$ is generated by $\mathrm{Pol}(\Gamma)$ and a canonical function $g : (\Delta, c_1, \ldots, c_k)^m \to \Delta$, where $m \geq 1$, $k \geq 0$, and $c_1, \ldots, c_k \in \Delta$. Moreover, m only depends on the number of $n(\Gamma)$-types in Γ (and not on the clone), and k only depends on m and $n(\Gamma)$, and the number of minimal closed clones above $\mathrm{Pol}(\Gamma)$ is finite.*

In the case of minimal closed clones above an endomorphism monoid, the arity of the generating canonical functions can be further reduced as follows.

Proposition 34 (Bodirsky, Pinsker, Tsankov [16])**.** *Let Γ be a finite language reduct of a structure Δ which is ordered, Ramsey, homogeneous and of finite language. Then every minimal closed clone above $\mathrm{End}(\Gamma)$ is generated by $\mathrm{End}(\Gamma)$ and a canonical function $g : (\Delta, c_1, \ldots, c_{n(\Gamma)}) \to \Delta$, or by $\mathrm{End}(\Gamma)$ and a canonical function $g : (\Delta, c_1, \ldots, c_m)^m \to \Delta$, where m only depends on the number of 2-types in Γ (and not on the clone). In particular, the number of minimal closed clones above $\mathrm{End}(\Gamma)$ is finite.*

Using this technique, the minimal closed clones containing the automorphism group of the random graph $G = (V; E)$ have been determined. In the following, let $f : V^2 \to V$ be a binary operation; we now define some possible behaviors for f. We say that f is

- *of type p_1 iff for all $x_1, x_2, y_1, y_2 \in V$ with $x_1 \neq x_2$ and $y_1 \neq y_2$ we have $E(f(x_1, y_1), f(x_2, y_2))$ if and only if $E(x_1, x_2)$;*
- *of type* max *iff for all $x_1, x_2, y_1, y_2 \in V$ with $x_1 \neq x_2$ and $y_1 \neq y_2$ we have $E(f(x_1, y_1), f(x_2, y_2))$ if and only if $E(x_1, x_2)$ or $E(y_1, y_2)$;*
- *balanced in the first argument iff for all $x_1, x_2, y \in V$ with $x_1 \neq x_2$ we have $E(f(x_1, y), f(x_2, y))$ if and only if $E(x_1, x_2)$;*
- *balanced in the second argument iff $(x, y) \mapsto f(y, x)$ is balanced in the first argument;*
- *E-dominated in the first argument iff for all $x_1, x_2, y \in V$ with $x_1 \neq x_2$ we have that $E(f(x_1, y), f(x_2, y))$;*
- *E-dominated in the second argument iff $(x, y) \mapsto f(y, x)$ is E-dominated in the first argument.*

The *dual* of an operation $f(x_1, \ldots, x_n)$ on V is defined by $-f(-x_1, \ldots, -x_n)$.

Theorem 35 (Bodirsky and Pinsker [14])**.** *Let $G = (V; E)$ be the random graph, and let $\mathcal{C}$ be a minimal closed clone above $\langle \mathrm{Aut}(G) \rangle$. Then $\mathcal{C}$ is generated by $\mathrm{Aut}(G)$ together with one of the unary functions of Theorem 30, or by $\mathrm{Aut}(G)$ and one of the following canonical operations from G^2 to G:*

- *a binary injection of type p_1 that is balanced in both arguments;*
- *a binary injection of type* max *that is balanced in both arguments;*
- *a binary injection of type* max *that is E-dominated in both arguments;*
- *a binary injection of type p_1 that is E-dominated in both arguments;*
- *a binary injection of type p_1 that is balanced in the first and E-dominated in the second argument;*
- *the dual of one of the last four operations.*

In [15], the technique of canonical functions was applied again to climb up further in the lattice of closed clones above $\mathrm{Aut}(G)$ – we will come back to this in Section 8.

Another example are the minimal closed clones containing all permutations of a countably infinite base set X. Observe that the set $\mathcal{S}_X$ of all permutations on X is the automorphism group of the structure $(X; =)$ which has no relations.

Theorem 36 (Bodirsky and Kára [10]; cf. also [8])**.** *The minimal closed clones containing $\langle \mathcal{S}_X \rangle$ on a countably infinite set X are:*

- *The closed clone generated by $\mathcal{S}_X$ and any constant operation;*
- *The closed clone generated by $\mathcal{S}_X$ and any binary injection.*

Observe that any constant operation and any binary injection on X are canonical operations for the structure $(X; =)$.

We end this section with a last example which lists the minimal closed clones containing the self-embdeddings of the dense linear order $(\mathbb{Q}; <)$. As with the random graph and the empty structure, since $(\mathbb{Q}; <)$ is model-complete it follows that the monoid of self-embeddings of $(\mathbb{Q}; <)$ is just the closure of $\mathrm{Aut}((\mathbb{Q}; <))$ in $\mathbb{Q}^{\mathbb{Q}}$.

Let lex be a binary operation on $\mathbb{Q}$ such that $\mathrm{lex}(a, b) < \mathrm{lex}(a', b')$ iff either $a < a'$ or $a = a'$ and $b < b'$, for all $a, a', b, b' \in \mathbb{Q}$. Observe that lex is canonical as a function from $\mathbb{Q}^2$ to $\mathbb{Q}$. Next, let pp be an arbitrary binary operation on $\mathbb{Q}$ such that for all $a, a', b, b' \in \mathbb{Q}$ we have $\mathrm{pp}(a, b) \le \mathrm{pp}(a', b')$ iff one of the following cases applies:

- $a \le 0$ and $a \le a'$;
- $0 < a$, $0 < a'$, and $b \le b'$.

The name of the operation pp stands for "projection-projection", since the operation behaves as a projection to the first argument for negative first argument, and a projection to the second argument for positive first argument. Observe that pp is canonical if we add the origin as a constant to the language. Finally, define the *dual* of an operation $f(x_1, \dots, x_n)$ on $\mathbb{Q}$ by $\leftrightarrow(f(\leftrightarrow(x_1), \dots, \leftrightarrow(x_n)))$.

Theorem 37 (Bodirsky and Kára [11])**.** *Let $(\mathbb{Q}; <)$ be the order of the rationals, and let $\mathcal{C}$ be a minimal closed clone above $\langle \mathrm{Aut}((\mathbb{Q}; <)) \rangle$. Then $\mathcal{C}$ is generated by $\mathrm{Aut}((\mathbb{Q}; <))$ together with one of the following operations:*

- *a constant operation;*
- *the operation $\leftrightarrow$;*
- *the operation $\circlearrowleft$;*
- *the operation lex;*
- *the operation pp;*
- *the dual of pp.*

6. Decidability of Definability

We turn to another application of the ideas of the last sections. Consider the following computational problem for a structure Γ: Input are quantifier-free formulas $\phi_0, \dots, \phi_n$ in the language of Γ defining relations $R_0, \dots, R_n$ on the domain of Γ, and the question is whether R_0 can be *defined* from $R_1, \dots, R_n$. As in Section 2, "defined" can stand for "first-order defined" or syntactic restrictions of this notion. We denote this computational problem by $\mathrm{Expr}_{ep}(\Gamma)$ and $\mathrm{Expr}_{pp}(\Gamma)$ if we consider existential positive and primitive positive definability, respectively.

For *finite* structures Γ the problem $\mathrm{Expr}_{pp}(\Gamma)$ is in co-NEXPTIME (and in particular decidable), and has recently shown to be co-NEXPTIME-hard [49]. For infinite structures Γ, the decidability of $\mathrm{Expr}_{pp}(\Gamma)$ is not obvious. An algorithm for primitive positive definability has theoretical and practical consequences in the study of the computational complexity of CPSs (which we will consider in Section 8). It is motivated by the fundamental fact that expansions of structures Γ by primitive positive relations do not change the complexity of $\mathrm{CSP}(\Gamma)$. On a practical side, it turns out that hardness of a CSP 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 might be a useful tool to determine the computational complexity of CSPs.

Using the methods of the last sections, one can show decidability of $\mathrm{Expr}_{ep}(\Gamma)$ and $\mathrm{Expr}_{pp}(\Gamma)$ for certain infinite structures Γ. The following uses the same terminology as in [35].

Definition 38. We say that a class $\mathcal{C}$ of finite τ-structures (or a τ-structure with age $\mathcal{C}$) is *finitely bounded* if there exists a finite set of finite τ-structures $\mathcal{F}$ such for all finite τ-structures A we have that $A \in \mathcal{C}$ iff no structure from $\mathcal{F}$ embeds into A.

Theorem 39 (Bodirsky, Pinsker, Tsankov [16])**.** *Let Δ be ordered, Ramsey, homogeneous, and of finite language, and let Γ be a finite language reduct of Δ. Then $\mathrm{Expr}_{ep}(\Gamma)$ and $\mathrm{Expr}_{pp}(\Gamma)$ are decidable.*

Examples of structures Δ that satisfy the assumptions of Theorem 39 are $(\mathbb{Q}; <)$, the Fraïssé limit of ordered finite graphs (or tournaments [39]), the Fraïssé limit of finite partial orders with a linear extension [39], the homogeneous universal 'naturally ordered' C-relation [13], just to name a few. CSPs for structures that are definable in such structures are abundant in particular for qualitative reasoning calculi in Artificial Intelligence.

We want to point out that that decidability of primitive positive definability is already non-trivial when Γ is trivial from a model-theoretic perspective: for the case that Γ is the structure $(X; =)$ (where X is countably infinite), the decidability of $\mathrm{Expr}_{pp}(\Gamma)$ has been posed as an open problem in [8]. Theorem 39 solves this problem, since $(X; =)$ is isomorphic to a reduct of the structure $(\mathbb{Q}; <)$, which is clearly finitely bounded, homogeneous, ordered, and Ramsey.

The proof of Theorem 39 goes along the following lines, and is based on the results of the last sections. We outline the algorithm for $\mathrm{Expr}_{pp}(\Gamma)$; the proof for $\mathrm{Expr}_{ep}(\Gamma)$ is a subset. So the input are formulas $\phi_0, \ldots, \phi_n$ defining relations $R_0, \ldots, R_n$, and we have to decide whether R_0 has a primitive positive definition from $R_1, \ldots, R_n$. Let Θ be the structure which has $R_1, \ldots, R_n$ as its relations. By Theorem 9, R_0 is not primitive positive definable from $R_1, \ldots, R_n$ if and only if there is a finitary function $f \in \mathrm{Pol}(\Theta)$ which violates R_0. By the ideas of the last section, such a polymorphism can be chosen to be canonical as a function from $(\Delta, c_1, \ldots, c_k)^m$ to Δ, where $c_i \in \Delta$. Such canonical functions are essentially finite objects since they can be represented as functions on types. Therefore, the algorithm can then check for a given canonical function whether it is a polymorphism of Θ and whether it violates R_0. Also, k and m can be calculated from the input, and so there are only finitely many complete behaviors to be checked. Finally, the additional assumption that Δ be finitely bounded allows the algorithm to check whether a function on types really comes from a function on Δ. We refer to [16] for details.

7. Interpretability

Many ω-categorical structures can be derived from other ω-categorical structures via first-order interpretations. In this section we will discuss the fact already mentioned in the introduction that bi-interpretations can be used to transfer the Ramsey property from one structure to another. A special type of interpretations, called *primitive positive interpretations*, will become important in Section 8. The definition of interpretability we use is standard, and follows [28].

When Δ is a structure with signature τ, and $\delta(x_1, \ldots, x_k)$ is a first-order τ-formula with the k free variables $x_1, \ldots, x_k$, we write $\delta(\Delta^k)$ for the k-ary relation that is defined by δ over Δ.

Definition 40. A relational σ-structure Γ has a *(first-order) interpretation* in a τ-structure Δ if there exists a natural number d, called the *dimension* of the interpretation, and

- a τ-formula $\delta(x_1, \ldots, x_d)$ – called *domain formula*,
- for each k-ary relation symbol R in σ a τ-formula $\phi_R(\overline{x}_1, \ldots, \overline{x}_k)$ where the $\overline{x}_i$ denote disjoint d-tuples of distinct variables – called the *defining formulas*,
- a τ-formula $\phi_=(x_1, \ldots, x_d, y_1, \ldots, y_d)$, and
- a surjective map $h : \delta(\Delta^d) \to \Gamma$ – called *coordinate map*,

such that for all relations R in Γ and all tuples $\overline{a}_i \in \delta(\Delta^d)$

$$(h(\overline{a}_1), \ldots, h(\overline{a}_k)) \in R \iff \Delta \models \phi_R(\overline{a}_1, \ldots, \overline{a}_k) , \text{ and}$$
$$h(\overline{a}_1) = h(\overline{a}_2) \iff \Delta \models \phi_=(\overline{a}_1, \overline{a}_2) .$$

If the formulas δ, ϕ_R, and $\phi_=$ are all primitive positive, we say that Γ has a *primitive positive interpretation* in Δ; many primitive positive interpretations can be found in Section 8. We say that Γ is *interpretable in Δ with finitely many parameters* if there are $c_1, \ldots, c_n \in \Delta$ such that Γ is interpretable in the expansion of Δ by the singleton relations $\{c_i\}$ for all $1 \leq i \leq n$. First-order *definitions* are a special case of interpretations: a structure Γ is *(first-order) definable* in Δ if Γ has an interpretation in Δ of dimension one where the domain formula is logically equivalent to true.

Lemma 41 (see e.g. Theorem 7.3.8 in [28]). *If Δ is an ω-categorical structure, then every structure Γ that is first-order interpretable in Δ with finitely many parameters is ω-categorical as well.*

The following nicely describes interpretability between structures in terms of the (topological) automorphism groups of the structures.

Theorem 42 (Ahlbrandt and Ziegler [3]; also see Theorem 5.3.5 and 7.3.7 in [28]). *Let Δ be an ω-categorical structure with at least two elements. Then a structure Γ has a first-order interpretation in Δ if and only if there is a continuous group homomorphism $f : \mathrm{Aut}(\Delta) \to \mathrm{Aut}(\Gamma)$ such that the image of f has finitely many orbits in its action on Γ.*

Note that if Γ_2 has a d-dimensional interpretation I in Γ_1, and Γ_3 has an e-dimensional interpretation J in Γ_2, then Γ_3 has a natural ed-dimensional interpretation in Γ_1, which we denote by $J \circ I$. To formally describe $J \circ I$, suppose that the signature of Γ_i is τ_i for $i = 1, 2, 3$, and that $I = (d, \delta, (\phi_R)_{R \in \tau_2}, \phi_=, h)$ where d is the dimension, δ the domain formula, $\phi_=$ and $(\phi_R)_{R \in \tau_2}$ the interpreting relations, and h the coordinate map. Similarly, let $J = (e, \gamma, (\psi_R)_{R \in \tau_3}, \psi_=, g)$. We use the following.

Lemma 43 (Theorem 5.3.2 in [28]). *Let Γ_1, Γ_2, I as in the preceding paragraph. Then for every first-order τ_2-formula $\phi(x_1, \ldots, x_k)$ there is τ_1-formula*

$$\phi^I(x_1^1, \ldots, x_1^d, \ldots, x_k^1, \ldots, x_k^d)$$

such that for all $a_1, \ldots, a_k \in \delta((\Gamma_1)^d)$

$$\Gamma_2 \models \phi(h(a_1), \ldots, h(a_k)) \ \Leftrightarrow \ \Gamma_1 \models \phi^I(a_1, \ldots, a_k) \,.$$

We can now define the interpretation $J \circ I$ as follows: the domain formula η is γ^I, and the defining formula for $R \in \tau_3$ is $(\psi_R)^I$. The coordinate map is from $\eta((\Gamma_1)^{ed}) \to \Gamma_3$, and defined by

$$(a_1^1, \ldots, a_1^d, \ldots, a_e^1, \ldots, a_e^d) \ \mapsto \ g(h(a_1^1, \ldots, a_1^d), \ldots, h(a_e^1, \ldots, a_e^d)) \,.$$

Two interpretations of Γ in Δ with coordinate maps h_1 and h_2 are called *homotopic* if the relation $\{(\bar{x}, \bar{y}) \mid h_1(\bar{x}) = h_2(\bar{y})\}$ is definable in Δ. The *identity interpretation* of a structure Γ is the 1-dimensional interpretation of Γ in Γ whose coordinate map is the identity. Two structures Γ and Δ are called *bi-interpretable* if there is an interpretation I of Γ in Δ and an interpretation J of Δ in Γ such that both $I \circ J$ and $J \circ I$ are homotopic to the identity interpretation (of Γ and of Δ, respectively).

Theorem 44 (Ahlbrandt and Ziegler [3]). *Two ω-categorical structures Γ and Δ are bi-interpretable if and only if $\mathrm{Aut}(\Gamma)$ and $\mathrm{Aut}(\Delta)$ are isomorphic as topological groups.*

As a consequence of this result and Theorem 23 we obtain the following.

Corollary 45. *For ordered bi-interpretable ω-categorical homogeneous structures Γ and Δ, one has the Ramsey property if and only if the other one has the Ramsey property.*

We give an example. This corollary can be used to deduce that an important structure studied in temporal reasoning in artificial intelligence has the Ramsey property. For the relevance of this fact in constraint satisfaction, see Section 8.

We have already mentioned that the age of $(\mathbb{Q}; <)$ has the Ramsey property. Let Γ be the structure whose elements are pairs $(x, y) \in \mathbb{Q}^2$ with $x < y$, representing *intervals*, and which contains all binary relations R over those intervals such that the relation $\{(x, y, u, v) \mid ((x, y), (u, v)) \in R\}$ is first-order definable in $(\mathbb{Q}; <)$. Hence, Γ has a 2-dimensional interpretation I in $(\mathbb{Q}; <)$, whose coordinate map h_1 is the identity map on $D := \{(x, y) \in \mathbb{Q}^2 \mid x < y\}$.

The structure Γ is known under the name *Allen's Interval Algebra* in artificial intelligence. We claim that its age has the Ramsey property. Using the homogeneity of $(\mathbb{Q}; <)$, it is easy to show that Γ is homogeneous as well. By Corollary 45, it suffices to show that Γ and $(\mathbb{Q}; <)$ are bi-interpretable. We first show that $(\mathbb{Q}; <)$ has an interpretation J in Γ. The coordinate map h_2 of J maps $(x, y) \in D$ to x. The formula $\phi_=(a, b)$ is $R_0(a, b)$ where R_0 is the binary relation $\{((x, y), (u, v)) \mid x = u\}$ from Γ. The formula $\phi_<(a, b)$ is $R_1(a, b)$ where R_1 is the binary relation $\{((x, y), (u, v)) \mid x < u\}$.

We prove that $J \circ I$ is homotopic to the identity interpretation of $(\mathbb{Q}; <)$ in $(\mathbb{Q}; <)$. This holds since the relation $\{(x, y, u) \in \mathbb{Q}^3 \mid h_2(h_1(x, y)) = u\}$ has the first-order definition $x = u$ in $(\mathbb{Q}; <)$. To show that $I \circ J$ is homotopic to the identity interpretation, observe that the relation $\{(a, b, c) \in D^3 \mid h_1(h_2(a), h_2(b)) = c\}$ has the first-order definition $R_0(a, c) \wedge R_3(b, c)$ in Γ, where R_0 is the binary relation from Γ as defined above, and R_3 is the binary relation $\{((x, y), (u, v)) \in \Gamma^2 \mid x = v\}$ from Γ. This shows that Γ and $(\mathbb{Q}; <)$ are bi-interpretable.

8. Complexity of Constraint Satisfaction

In recent years, a considerable amount of research concentrated on the computational complexity of CSP(Γ) for *finite* structures Γ. Feder and Vardi [22] conjectured that for such Γ, the problem CSP(Γ) is either in P, or NP-complete[1]. This conjecture has been fascinating researchers from various areas, for instance from graph theory [40] and from finite model theory [4, 22, 33]. It has been discovered that complexity classification questions translate to fundamental questions in universal algebra [19, 29], so that lately also many researchers in universal algebra started to work on questions that directly correspond to questions about the complexity of CSPs.

For arbitrary infinite structures Γ it can be shown that there are problems CSP(Γ) that are in NP, but neither in P nor NP-complete, unless P=NP. In fact, it can be shown that for every computational problem $\mathcal{P}$ there is an infinite structure Γ such that $\mathcal{P}$ and CSP(Γ) are equivalent under polynomial-time Turing reductions [9]. However, there are several classes of infinite structures Γ for which the complexity of CSP(Γ) can be classified completely.

In this section we will see three such classes of computational problems; they all have the property that

- every problem in this class can be formulated as CSP(Γ) where Γ has a first-order definition in a *base structure* Δ;
- Δ is ordered homogeneous Ramsey with finite signature.

For all three classes, the classification result can be obtained by the same method, which we describe in the following two subsections.

8.1. **Climbing up the lattice.** Clearly, if we add relations to a structure Γ with a finite relational signature, then the CSP of the structure thus obtained is computationally at least as complex as the CSP of Γ. On the other hand, when we add a primitive positive definable relation to Γ, then the CSP of the resulting structure has a polynomial-time reduction to CSP(Γ). This is not hard to show, and has been observed for finite domain structures in [30]; the same proof also works for structures over an infinite domain.

Lemma 46. *Let $\Gamma = (D; R_1, \ldots, R_l)$ be a relational structure, and let R be a relation that has a primitive positive definition in Γ. Then the problems CSP(Γ) and CSP($D; R, R_1, \ldots, R_l$) are equivalent under polynomial-time reductions.*

When we study the CSPs of the reducts Γ of a structure Δ, we therefore consider the lattice of reducts of Δ which are closed under primitive positive definitions (i.e., which contain all relations that are primitive positive definable from the reduct), and describe the border between tractability and NP-completeness in this lattice. We remark that the reducts of Δ have, since we expand them by all primitive positive definable relations, infinitely many relations, and hence do not define a CSP; however, we consider Γ tractable if and only if all structures obtained from Γ by dropping all but finitely many relations have a tractable CSP. Similarly, we consider Γ hard if there exists a structure obtained from Γ by dropping all but finitely many relations that has a hard CSP. With this convention, it is interesting to determine the *maximal tractable* reducts, i.e., those reducts closed under primitive positive definitions which do not contain any hard relation and which cannot be further extended without losing this property.

---

[1]By Ladner's theorem [34], there are infinitely many complexity classes between P and NP, unless P=NP.

Recall the notion of a *clone* from Section 2. By Theorem 9, the lattice of primitive positive closed reducts of Δ and the lattice of closed clones containing $\text{Aut}(\Delta)$ are antiisomorphic via the mappings $\Gamma \mapsto \text{Pol}(\Gamma)$ (for reducts Γ) and $\mathcal{C} \mapsto \text{Inv}(\mathcal{C})$ (for clones $\mathcal{C}$). We refer to the introduction of [8] for a detailed exposition of this well-known connection. Therefore, the maximal tractable reducts correspond to *minimal tractable* clones, which are precisely the clones of the form $\text{Pol}(\Gamma)$ for a maximal tractable reduct.

The proof strategy of the classification results presented in Sections 8.3, 8.4, and 8.5 is as follows. We start by proving that certain reducts Γ have an NP-hard CSP. How to show this, and how to find those 'basic hard reducts' will be the topic of the next subsection. Let R be one of the relations from those hard reducts. If R does not have a primitive positive definition in Γ, then Theorem 9 implies that Γ has a polymorphism f that does not preserve R. We are now in a similar situation as in Section 5. Introducing constants, we can show that f generates an operation g that still does not preserve R but is canonical with respect to the expansion of Γ by constants. There are only finitely many canonical behaviours that g might have, and therefore we can start a combinatorial analysis. In the three classifications that follow, this strategy always leads to polymorphisms that imply that $\text{CSP}(\Gamma)$ can be solved in polynomial time.

8.2. **Primitive positive interpretations, and adding constants.** Surprisingly, in all the classification results that we present in Sections 8.3, 8.4, and 8.5, there is a single condition that implies that a CSP is NP-hard. Recall that an interpretation is called *primitive positive* if all formulas involved in the interpretation (the domain formula, the formulas ϕ_R and $\phi_=$) are primitive positive. The relevance of primitive positive interpretations in constraint satisfaction comes from the following fact, which is known for finite domain constraint satisfaction, albeit not using the terminology of primitive positive interpretations [19]. In the present form, it appears first in [6].

Theorem 47. *Let Γ and Δ be structures with finite relational signatures. If there is a primitive positive interpretation of Γ in Δ, then there is a polynomial-time reduction from $\text{CSP}(\Gamma)$ to $\text{CSP}(\Delta)$.*

All hardness proofs presented later can be shown via primitive positive interpretations of Boolean structures (i.e., structures with the domain $\{0, 1\}$) with a hard CSP. In fact, in all such Boolean structures the relation NAE defined as

$$\text{NAE} = \{0, 1\}^3 \setminus \{(0, 0, 0), (1, 1, 1)\}$$

is primitive positive definable. This fact has not been stated in the original publications; however, it deserves to be mentioned as a unifying feature of all the classification results presented here. It is often more convenient to interpret other Boolean structures than $(\{0, 1\}; \text{NAE})$, and to then apply the following Lemma. An operation $f : D^k \to D$ is called *essentially a permutation* if there exists an i and a bijection $g : D \to D$ so that $f(x_1, \ldots, x_k) = g(x_i)$ for all $(x_1, \ldots, x_k) \in D^k$.

Lemma 48. *Let Δ be a structure that interprets a Boolean structure Γ such that all polymorphisms of Γ are essentially a permutation. Then the structure $(\{0, 1\}; \text{NAE})$ has a primitive positive interpretation in Δ, and $\text{CSP}(\Delta)$ is NP-hard.*

Proof. Since the polymorphisms of Γ preserve the relation NAE, and by the well-known finite analog of Theorem 9 (due to [25] and independently, [17]), NAE is primitive positive definable in Γ. When ϕ is such a primitive positive definition, by substituting all relations in

ϕ by their defining relations in Δ we obtain an interpretation of $(\{0, 1\}; \text{NAE})$ in Δ. Hardness of $\text{CSP}(\Delta)$ follows from the NP-hardness of $\text{CSP}((\{0, 1\}; \text{NAE}))$ (this problem is called NOT-ALL-3-EQUAL-3SAT in [24]) and Theorem 47. $\qquad\square$

Typical Boolean structures Γ such that all polymorphisms of Γ are essentially a permutation are the structure $(\{0, 1\}; \{(t_1, t_2, t_3, t_4) \in \{0, 1\} \mid t_1 + t_2 + t_3 + t_4 = 2\}$, the structure $(\{0, 1\}; 1\text{IN}3)$, or the structure $(\{0, 1\}; \text{NAE})$ itself.

Sometimes it is not possible to give a primitive positive interpretation of the structure $(\{0, 1\}; \text{NAE})$ in Γ, but it is possible after expanding Γ with constants. Under an assumption about the endomorphism monoid of Γ, however, introducing constants does not change the computational complexity of Γ. More precisely, we have the following.

Theorem 49 (Theorem 19 in [5]). *Let Γ be an ω-categorical structure with a finite relational signature such that $\text{Aut}(\Gamma)$ is dense in $\text{End}(\Gamma)$. Then for any finite number of elements $c_1, \ldots, c_k$ of Γ there is a polynomial-time reduction from $\text{CSP}((\Gamma, \{c_1, \}, \ldots, \{c_k\}))$ to $\text{CSP}(\Gamma)$.*

8.3. **Reducts of equality.** One of the most fundamental classes of ω-categorical structures is the class of all reducts of $(X; =)$, where X is an arbitrary countably infinite set. Up to isomorphism, this is exactly the class of countable structures that are preserved by all permutations of their domain. The other two classes of ω-categorical structures that we will study here both contain this class.

We go straight to the statement of the complexity classification in terms of primitive positive interpretations. This is essentially a reformulation of a result from [10] which has been formulated without primitive positive interpretations. It turns out that when Γ is preserved by the operations from one of the minimal clones above the clone generated by all the permutations of X, then $\text{CSP}(\Gamma)$ can be solved in polynomial time, and otherwise $\text{CSP}(\Gamma)$ is NP-hard.

Theorem 50 (essentially from [10]). *Let Γ be a reduct of $(X; =)$. Then exactly one of the following holds.*

- *Γ has a constant endomorphism. In this case, $\text{CSP}(\Gamma)$ is trivially in P.*
- *Γ has a binary injective polymorphism. In this case, $\text{CSP}(\Gamma)$ is in P.*
- *All relations with a first-order definition in $(X; =)$ have a primitive positive definition in Γ. Furthermore, the structure $(\{0, 1\}; \text{NAE})$ has a primitive positive interpretation in Γ, and $\text{CSP}(\Gamma)$ is NP-complete.*

Proof. It has been shown in [10] that $\text{CSP}(\Gamma)$ is in P when Γ has a constant or a binary injective polymorphism. Otherwise, by Theorem 36, every polymorphism of Γ is generated by the permutations of X. Hence, every relation R with a first-order definition in $(X; =)$ is preserved by all polymorphisms of Γ, and it follows from Theorem 9 that every relation is primitive positive definable in Γ.

This holds in particular for the relation E_6 defined as follows.

$$E_6 = \{(x_1, x_2, y_1, y_2, z_1, z_2) \in X^6 \mid (x_1 = x_2 \wedge y_1 \neq y_2 \wedge z_1 \neq z_2)$$
$$\vee\ (x_1 \neq x_2 \wedge y_1 = y_2 \wedge z_1 \neq z_2)$$
$$\vee\ (x_1 \neq x_2 \wedge y_1 \neq y_2 \wedge z_1 = z_2)\}$$

We now show that the structure $(\{0, 1\}; 1\text{IN}3)$ has a primitive positive interpretation in $(X; E_6)$, which by Lemma 48 also shows that $(\{0, 1\}; \text{NAE})$ has a primitive positive interpretation in $(X; E_6)$ and that $\text{CSP}(\Gamma)$ is NP-hard.

The dimension of the interpretation is 2, and the domain formula is *'true'*. The formula $\phi_{1\text{IN3}}(x_1, x_2, y_1, y_2, z_1, z_2)$ is $E_6(x_1, x_2, y_1, y_2, z_1, z_2)$, and The formula $\phi_=(x_1, x_2, y_1, y_2)$ is

$$\exists a_1, a_2, u_1, u_2, u_3, u_4, z_1, z_2. \; a_1 = a_2 \wedge E_6(a_1, a_2, u_1, u_2, u_3, u_4)$$
$$\wedge E_6(u_1, u_2, x_1, x_2, z_1, z_2) \wedge E_6(u_3, u_4, z_1, z_2, y_1, y_2).$$

Note that the primitive positive formula $\phi_=(x_1, x_2, y_1, y_2)$ is equivalent to $x_1 = x_2 \Leftrightarrow y_1 = y_2$. The map h maps (a_1, a_2) to 1 if $a_1 = a_2$, and to 0 otherwise.   □

Note that both the constant and the binary injective operation are canonical as functions over $(X; =)$.

8.4. **Reducts of the dense linear order.** An extension of the result in the previous subsection has been obtained in [11]; there, the complexity of the CSP for all reducts of $(\mathbb{Q}; <)$ has been classified. By a theorem of Cameron, those reducts are (again up to isomorphism) exactly the structures that are *highly set-transitive* [20], i.e., structures Γ such that for any two finite subsets A, B with $|A| = |B|$ of the domain there is an automorphism of Γ that maps A to B.

The corresponding class of CSPs contains many computational problems that have been studied in Artificial Intelligence, in particular in temporal reasoning [18, 37, 48], but also in scheduling [36] or general theoretical computer science [23, 43]. The following theorem is a consequence of results from [11]. Again, we show that the hardness proofs in this class are captured by interpreting Boolean structures with few polymorphisms via primitive positive interpretations with finitely many parameters; this has not appeared in [11], so we provide the proof. The central arguments in the classification follow the reduct classification technique based on Ramsey theory that we present in this survey; see Figure 1 for an illustration of the bottom of the lattice of reducts of $(\mathbb{Q}; <)$, and the border of tractability for such reducts.

Theorem 51 (essentially from [11]). *Let Γ be a reduct of $(\mathbb{Q}; <)$. Then exactly one of the following holds.*

- *Γ has one out of 9 binary polymorphisms (for a detailed description of those see [11]), and $\mathrm{CSP}(\Gamma)$ is in P.*
- *$\mathrm{Aut}(\Gamma)$ is dense in $\mathrm{End}(\Gamma)$, and the structure $(\{0, 1\}; \mathrm{NAE})$ has a primitive positive interpretation with finitely many parameters in Γ. In this case, $\mathrm{CSP}(\Gamma)$ is NP-complete.*

Before we derive Theorem 51 from what has been shown in [11], we would like to point to Figure 1 for an illustration of the clones that correspond to maximal tractable reducts. The diagram also shows the constraint languages that just contain one of the important relations *Betw* (introduced in the introduction), *Cycl*, *Sep* (*Cycl* and *Sep* already appeared in Section 2), E_6 (which appeared earlier in this section), T_3, and $-T_3$. Here, T_3 stands for the relation

$$\{(x, y, z) \in \mathbb{Q}^3 \mid (x = y < z) \vee (x = z < y)\},$$

and when $R \subseteq \mathbb{Q}^k$, then $-R$ denotes $\{(-t_1, \ldots, -t_k) \mid (t_1, \ldots, t_k) \in R\}$.

The importance of those relations comes from the fact (shown in [11]) that unless Γ has one out of the 9 binary polymorphisms mentioned in Theorem 51 then there is a primitive positive definition of at least one of the relations *Betw*, *Cycl*, *Sep*, E_6, T_3, or $-T_3$.

Proof of Theorem 51. It has been shown in [11] that unless Γ has a constant endomorphism, $\mathrm{Aut}(\Gamma)$ is dense in $\mathrm{End}(\Gamma)$. We have already seen that there is a primitive positive interpretation of $(\{0, 1\}; \mathrm{NAE})$ in structures isomorphic to $(\mathbb{Q}; E_6)$.
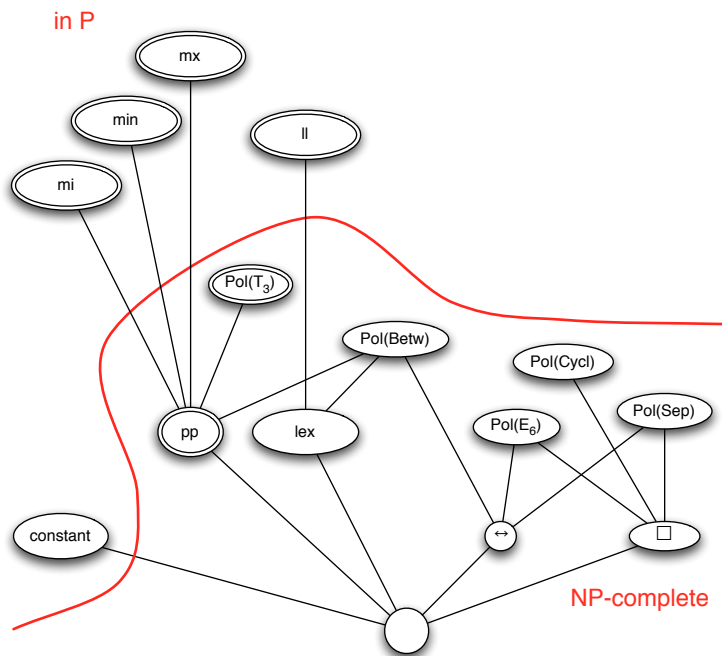
FIGURE 1. An illustration of the classification result for temporal constraint languages. Double-circles mean that the corresponding operation has a dual generating a distinct clone which is not drawn in the figure. For the definition of mi, min, mx, and ll, see [11].

Now suppose that T_3 is primitive positive definable in Γ. We give below a primitive positive interpretation of the structure $(\{0,1\}; 1\text{IN}3)$ in $\Delta = (\mathbb{Q}; T_3, 0)$. Hence, there is also a primitive positive definition of $(\{0,1\}; 1\text{IN}3)$ in the expansion of Γ by the constant 0. Expansions by constants do not change the computational complexity of $\text{CSP}(\Gamma)$ since $\text{Aut}(\Gamma)$ is dense in $\text{End}(\Gamma)$. Thus, Lemma 48 shows NP-hardness of $\text{CSP}(\Gamma)$, and that $(\{0,1\}; \text{NAE})$ has a primitive positive interpretation in $(\Gamma, 0)$.

The interpretation of $(\{0,1\}; 1\text{IN}3)$ in Δ

- has dimension 2;
- the domain formula $\delta(x_1, x_2)$ is $T_3(0, x_1, x_2)$;
- the formula $\phi_{1\text{IN}3}(x_1, x_2, y_1, y_2, z_1, z_2)$ is

$$\exists u.\ T_3(u, x_1, y_1) \wedge T_3(0, u, z_1)\ ;$$

- the formula $\phi_=(x_1, x_2, y_1, y_2)$ is $T_3(0, x_1, y_2)$;
- the coordinate map $h : \delta(\Delta^2) \to \{0,1\}$ is defined as follows. Let (b_1, b_2) be a pair of elements of Δ that satisfies δ. Then exactly one of b_1, b_2 must have value 0, and the other element is strictly greater than 0. We define $h(b_1, b_2)$ to be 1 if $b_1 = 0$, and to be 0 otherwise.

To see that this is the intended interpretation, let $(x_1, x_2), (y_1, y_2), (z_1, z_2) \in \delta(\Delta^2)$, and suppose that $t := (h(x_1, x_2), h(y_1, y_2), h(z_1, z_2)) = (1, 0, 0) \in 1\text{IN}3$. We have to verify that $(x_1, x_2, y_1, y_2, z_1, z_2)$ satisfies $\phi_{1\text{IN}3}$ in Δ. Since $h(x_1, x_2) = 1$, we have $x_1 = 0$, and similarly we

get that $y_1, z_1 > 0$. We can then set u to 0 and have $T_3(u, x_1, y_1)$ since $0 = u = x_1 < y_1$, and we also have $T_3(0, u, z_1)$ since $0 = u < z_1$. The case that $t = (0, 1, 0)$ is analogous. Suppose now that $t = (0, 0, 1) \in 1\text{IN}3$. Then $x_1, y_1 > 0$, and $z_1 = 0$. We can then set u to $\min(x_1, y_1)$, and therefore have $T_3(u, x_1, y_1)$, and $T_3(0, u, z_1)$ since $0 = z_1 < u$. Conversely, suppose that $(x_1, x_2, y_1, y_2, z_1, z_2)$ satisfies $\phi_{1\text{IN}3}$ in Δ. Since $T_3(0, u, z_1)$, exactly one out of u, z_1 equals 0. When $u = 0$, then because of $T_3(u, x_1, y_1)$ exactly one out of x_1, y_1 equals 0, and we get that $(h(x_1, x_2), h(y_1, y_2), h(z_1, z_2)) \in \{(0, 1, 0), (1, 0, 0)\} \subseteq 1\text{IN}3$. When $u > 0$, then $x_1 > 0$ and $y_1 > 0$, and so $(h(x_1, x_2), h(y_1, y_2), h(z_1, z_2)) = (0, 0, 1) \in 1\text{IN}3$.

An interpretation of $(\{0, 1\}; 1\text{IN}3)$ in $(\mathbb{Q}; -T_3, 0)$ can be obtained in a dual way.

Next, suppose that $Betw$ is primitive positive definable in Γ. We will give a primitive positive interpretation of $(\{0, 1\}; \text{NAE})$ in $(\mathbb{Q}; Betw, 0)$. Hence, when $Betw$ has a primitive positive definition in Γ, then by Theorem 49 (since $\text{Aut}(\Gamma)$ is dense in $\text{End}(\Gamma)$) and Lemma 48 we obtain NP-hardness of $\text{CSP}(\Gamma)$.

The dimension of the interpretation is one, and the domain formula is $x \neq 0$, which is clearly equivalent to a primitive positive formula over $(\mathbb{Q}; Betw, 0)$. The map h maps positive points to 1, and all other points from $\mathbb{Q}$ to 0. The formula $\phi_=(x_1, y_1)$ is

$$\exists z.\ Betw(x_1, 0, z) \wedge Betw(z, 0, y_1)$$

Note that the primitive positive formula $\phi_=$ is over $(\mathbb{Q}; Betw, 0)$ equivalent to $(x_1 > 0 \Leftrightarrow y_1 > 0)$. Finally, $\phi_{\text{NAE}}(x_1, y_1, z_1)$ is

$$\exists u.\ Betw(x_1, u, y_1) \wedge Betw(u, 0, z_1)\ .$$

If Sep has a primitive positive definition in Γ, then the statement follows easily from the previous argument since $Betw(x, y, z)$ has a 1-dimensional primitive positive interpretation in $(\mathbb{Q}; Sep)$ (the formula $\phi_{Betw}(x, y, z)$ is $\exists u.Sep(u, x, y, z)$).

Finally, if $Cycl$ is primitive positive definable in Γ, we give a 3-dimensional primitive positive interpretation of the structure $(\{0, 1\}; R, \neg)$ where $R = \{0, 1\}^3 \setminus \{(0, 0, 0)\}$ and $\neg = \{(0, 1), (1, 0)\}$. The idea of the interpretation is inspired by the NP-hardness proof of [23] for the 'Cyclic ordering problem' (see [24]).

The dimension of our interpretation is three, and the domain formula $\delta(x_1, x_2, x_3)$ is $x_1 \neq x_2 \wedge x_2 \neq x_3 \wedge x_3 \neq x_1$, which clearly has a primitive positive definition in $(\mathbb{Q}; Cycl)$. The coordinate map h sends (x_1, x_2, x_3) to 0 if $Cycl(x_1, x_2, x_3)$, and to 1 otherwise.

Let $\phi(x_1, x_2, x_3, y_1, y_2, y_3)$ be the formula

$$Cycl(x_1, y_1, x_2) \wedge Cycl(y_1, x_2, y_2) \wedge Cycl(x_2, y_2, x_3)$$
$$Cycl(y_2, x_3, y_3) \wedge Cycl(x_3, y_3, x_1) \wedge Cycl(y_3, x_1, y_1)\ .$$

When $(a_1, \ldots, a_6)$ satisfies ϕ, we can imagine $a_1, \ldots, a_6$ as points that appear clockwise in this order on the unit circle. In particular, we then have that $Cycl(a_1, a_3, a_5)$ holds if and only if $Cycl(a_2, a_4, a_6)$ holds. The formula $\phi_=(x_1, x_2, x_3, y_1, y_2, y_3)$ is

$$\exists u_1^1, \ldots, u_3^4.\ \phi(x_1, x_2, x_3, u_1^1, u_2^1, u_3^1) \wedge$$
$$\bigwedge_{i=1}^{3} \phi(u_1^i, u_2^i, u_3^i, u_1^{i+1}, u_2^{i+1}, u_3^{i+1}) \wedge \phi(u_1^4, u_2^4, u_3^4, y_1, y_2, y_3)\ ,$$

which is equivalent to

$$\delta(x_1, x_2, x_3) \wedge \delta(y_1, y_2, y_3) \wedge (Cycl(x_1, x_2, x_3) \Leftrightarrow Cycl(y_1, y_2, y_3))\ ;$$

this is tedious, but straightforward to verify, and we omit the proof.

The formula $\phi_\neg(x_1, x_2, x_3, y_1, y_2, y_3)$ is $\phi_=(x_1, x_2, x_3, z_1, z_3, z_2)$.
The formula $\phi_R(x_1, x_2, x_3, y_1, y_2, y_3, z_1, z_2, z_3)$ is

$$\exists a, b, c, d, e, f, g, h, i, j, k, l, m, n. \; Cycl(a, c, j) \wedge Cycl(b, j, k) \wedge Cycl(c, k, l)$$
$$\wedge \; Cycl(d, f, j) \wedge Cycl(e, j, l) \wedge Cycl(f, l, m)$$
$$\wedge \; Cycl(g, i, k) \wedge Cycl(h, k, m) \wedge Cycl(i, m, n)$$
$$\wedge \; Cycl(n, m, l) \wedge \phi_=(x_1, x_2, x_3, a, b, c)$$
$$\wedge \; \phi_=(y_1, y_2, y_3, d, e, f) \wedge \phi_=(z_1, z_2, z_3, g, h, i)$$

The proof that for all tuples $\bar{a}_1, \bar{a}_2, \bar{a}_3 \in \mathbb{Q}^3$

$$(h(\bar{a}_1), h(\bar{a}_3), h(\bar{a}_3)) \in R \; \Leftrightarrow \; (\mathbb{Q}; Cycl) \models \phi_R(\bar{a}_1, \bar{a}_2, \bar{a}_3)$$

follows directly the correctness proof of the reduction presented in [23]. $\qquad\square$

8.5. **Reducts of the random graph.** The full power of the technique that is developed in this paper can be used to obtain a full complexity classification for all reducts of the random graph $G = (V; E)$ [15]. Again, the result can be stated in terms of primitive positive interpretations – this is not obvious from the statement of the result in [15], therefore we provide the proofs.

Theorem 52 (essentially from [15]). *Let Γ be a reduct of the countably infinite random graph G. Then exactly one of the following holds.*

- *Γ has one out of 17 at most ternary canonical polymorphisms (for a detailed description of those see [15]), and $\mathrm{CSP}(\Gamma)$ is in P.*
- *Γ admits a primitive positive interpretation of $(\{0, 1\}; 1\mathrm{IN}3)$. In this case, $\mathrm{CSP}(\Gamma)$ is NP-complete.*

Proof. It has been shown in [15] that Γ has one out of 17 at most ternary canonical polymorphisms, and $\mathrm{CSP}(\Gamma)$ is in P, or one of the following relations has a primitive positive definition in Γ: the relation E_6, or the relation T, H, or $P^{(3)}$, which are defined as follows. The 4-ary relation T holds on $x_1, x_2, x_3, x_4 \in V$ if x_1, x_2, x_3, x_4 are pairwise distinct, and induce in G either

- a single edge and two isolated vertices,
- a path with two edges and an isolated vertex,
- a path with three edges, or
- a complement of one of the structures stated above.

To define the relation H, we write $N(u, v)$ as a shortcut for $E(u, v) \wedge u \neq v$. Then $H(x_1, y_1, x_2, y_2, x_3, y_3)$ holds on V if

$$\bigwedge_{i,j \in \{1,2,3\}, i \neq j, u \in \{x_i, y_i\}, v \in \{x_j, y_j\}} N(u, v)$$
$$\wedge \left(\left((E(x_1, y_1) \wedge N(x_2, y_2) \wedge N(x_3, y_3)) \right. \right.$$
$$\vee \; (N(x_1, y_1) \wedge E(x_2, y_2) \wedge N(x_3, y_3))$$
$$\left. \left. \vee \; (N(x_1, y_1) \wedge N(x_2, y_2) \wedge E(x_3, y_3)) \right) \right.$$

The ternary relation $P^{(3)}$ holds on x_1, x_2, x_3 if those three vertices are pairwise distinct and do not induce a clique or an independent set in G.

Suppose first that T is primitive positive definable in Γ. Let R be the relation $\{(t_1, t_2, t_3, t_4) \in \{0,1\} \mid t_1 + t_2 + t_3 + t_4 = 2\}$. We have already mentioned that all polymorphisms of $(\{0,1\}; R)$ are essentially permutations. To show that $(\{0,1\}; \mathrm{NAE})$ has a primitive positive interpretation in Γ, we can therefore use Lemma 48 and it suffices to show that there is a primitive positive interpretation of the structure $(\{0,1\}; R)$ in $(V; T)$. For a finite subset S of V, write $\#S$ for the parity of edges between members of S. Now we define the relation $L \subseteq V^6$ as follows.

$$L := \big\{ x \in V^6 \mid \text{the entries of } x \text{ are pairwise distinct, and}$$
$$\#\{x_1, x_2, x_3\} = \#\{x_4, x_5, x_6\} \big\}$$

It has been shown in [15] that the relation L is pp-definable in $(V; T)$. We therefore freely use the relation L (and similarly $\neq$, the disequality relation) in primitive positive formulas over $(V; T)$.

Our primitive positive interpretation of $(\{0,1\}; R)$ has dimension three. The domain formula $\delta(x_1, x_2, x_3)$ is $x_1 \neq x_2 \ \wedge \ x_1 \neq x_3 \ \wedge \ x_2 \neq x_3$. The formula $\phi_R(x_1^1, x_2^1, x_3^1, \ldots, x_1^4, x_2^4, x_3^4)$ of the interpretation is

$$\exists y_1, y_2, y_3, y_4.\ T(y_1, \ldots, y_4)$$
$$\wedge L(x_1^1, x_2^1, x_3^1, y_2, y_3, y_4)$$
$$\wedge L(x_1^2, x_2^2, x_3^2, y_1, y_3, y_4)$$
$$\wedge L(x_1^3, x_2^3, x_3^3, y_1, y_2, y_4)$$
$$\wedge L(x_1^4, x_2^4, x_3^4, y_1, y_2, y_3)$$

The formula $\phi_=$ is $L(x_1, x_2, x_3, y_1, y_2, y_3)$. Finally, the coordinate map sends a tuple (a_1, a_2, a_3) for pairwise distinct a_1, a_2, a_3 to 1 if $P^{(3)}(a_1, a_2, a_3)$, and to 0 otherwise.

Next, suppose that H is primitive positive definable in Γ. We give a 2-dimensional interpretation of $(\{0,1\}; \mathrm{1IN3})$ in Γ. The domain formula is '*true*'. The formula $\phi_=(x_1, x_2, y_1, y_2)$ is

$$\exists z_1, z_2, u_1, u_2, v_1, v_2.\ H(x_1, x_2, u_1, u_2, z_1, z_2) \wedge N(u_1, u_2)$$
$$\wedge\ H(z_1, z_2, v_1, v_2, y_1, y_2) \wedge N(v_1, v_2)\,.$$

This formula is equivalent to a primitive positive formula over Γ since $N(x, y)$ is primitive positive definable by H. The formula $\phi_{\mathrm{1IN3}}(x_1, x_2, y_1, y_2, z_1, z_2)$ is

$$\exists x_1', x_2', y_1', y_2', z_1', z_2'.\ H(x_1', x_2', y_1', y_2', z_1', z_2')$$
$$\wedge\ \phi_=(x_1, x_2, x_1', x_2') \wedge \phi_=(x_1, x_2, x_1', x_2') \wedge \phi_=(x_1, x_2, x_1', x_2')\,.$$

The coordinate map sends a tuple (x_1, x_2) to 1 if $E(x_1, x_2)$ and to 0 otherwise.

Finally, suppose that $P^{(3)}$ has a primitive positive definition in Γ. We give a 2-dimensional primitive positive interpretation of $(\{0,1\}; \mathrm{NAE})$. For $k \geq 3$, let $Q^{(k)}$ be the k-ary relation that holds for a tuple $(x_1, \ldots, x_k) \in V^k$ iff $x_1, \ldots, x_k$ are pairwise distinct, and $(x_1, \ldots, x_k) \notin P^{(k)}$. It has been shown in [15] that the relation $Q^{(4)}$ is primitive positive definable by the relation $P^{(3)}$. Now, the formula $\phi_=(x_1, x_2, y_1, y_2)$ is $\exists z_1, z_2. Q^{(4)}(x_1, x_2, z_1, z_2) \wedge Q^{(4)}(z_1, z_2, y_1, y_2)$. The

formula $\phi_{\text{NAE}}(x_1, x_2, y_1, y_2, z_1, z_2)$ is

$$\exists u, v, w. \; P^{(3)}(u, v, w) \wedge Q^{(4)}(x_1, x_2, u, v)$$
$$\wedge Q^{(4)}(y_1, y_2, v, w) \wedge Q^{(4)}(z_1, z_2, w, u) \, .$$

The coordinate map sends a tuple (x_1, x_2) to 1 if $E(x_1, x_2)$ and to 0 otherwise. $\qquad\square$

9. Concluding Remarks and Further Directions

We have outlined an approach to use Ramsey theory for the classification of reducts of a structure, considered up to existential positive, or primitive positive interdefinability. The central idea in this approach is to study functions that preserve the reduct, and to apply structural Ramsey theory to show that those functions must act regularly on large parts of the domain. This insight makes those functions accessible to combinatoral arguments and classification.

Our approach has been illustrated for the reducts of $(\mathbb{Q}; <)$, and the reducts of the random graph $(V; E)$. One application of the results is complexity classification of constraint satisfaction problems in theoretical computer science. Interestingly, the hardness proofs in those classifications all follow a common pattern: they are based on primitive positive interpretations. In particular, we proved complete complexity classifications without the typical computer science hardness proofs – rather, the hardness results follow from mathematical statements about primitive positive interpretability in ω-categorical structures.

There are many other natural and important ω-categorical structures besides $(\mathbb{Q}; <)$ and $(V; E)$ where this approach seems promising. We have listed some of the simplest and most basic examples in Figure 2. In this table, the first column specifies the 'base structure' Δ, and we will be interested in the class of all structures definable in Δ. The second column lists what is known about this class, considered up to first-order interdefinability. The third column describes the corresponding Ramsey result, when Δ is equipped with an appropriate linear order. The fourth column gives the status with respect to complexity classification of the corresponding class of CSPs. The fifth class indicates in which areas in computer science those CSPs find applications.

References

[1] Samson Adepoju Adeleke and Dugald Macpherson. Classification of infinite primitive Jordan permutation groups. *Proceedings of the London Mathematical Society*, 72(1):63–123, 1996.

[2] Samson Adepoju Adeleke and Peter M. Neumann. *Relations related to betweenness: their structure and automorphisms*, volume 623 of *Memoirs of the AMS 131*. American Mathematical Society, 1998.

[3] Gisela Ahlbrandt and Martin Ziegler. Quasi-finitely axiomatizable totally categorical theories. *Annals of Pure and Applied Logic*, 30(1):63–82, 1986.

[4] Albert Atserias, Andrei A. Bulatov, and Anuj Dawar. Affine systems of equations and counting infinitary logic. *Theoretical Computer Science*, 410(18):1666–1683, 2009.

[5] Manuel Bodirsky. Cores of countably categorical structures. *Logical Methods in Computer Science*, 3(1):1–16, 2007.

[6] Manuel Bodirsky. Constraint satisfaction problems with infinite templates. In Heribert Vollmer, editor, *Complexity of Constraints (a collection of survey articles)*, pages 196–228. Springer, LNCS 5250, 2008.

[7] Manuel Bodirsky and Hubert Chen. Oligomorphic clones. *Algebra Universalis*, 57(1):109–125, 2007.

[8] Manuel Bodirsky, Hubie Chen, and Michael Pinsker. The reducts of equality up to primitive positive interdefinability. *Journal of Symbolic Logic*, 75(4):1249–1292, 2010.

[9] Manuel Bodirsky and Martin Grohe. Non-dichotomies in constraint satisfaction complexity. In *Proceedings of ICALP'08*, pages 184–196, 2008.

Reducts of	First-order Reducts	Ramsey Class	CSP Dichotomy	Application, Motivation
$(X; =)$	Trivial	Ramsey's theorem	Yes	Equality Constraints
$(\mathbb{Q}; <)$	Cameron [21]	Ramsey's theorem	Yes	Temporal Reasoning
$(V; E)$	Thomas [46]	Nešetřil + Rödl [42]	Yes	Schaefer's theorem for graphs
Homogeneous universal poset	?	Nešetřil + Rödl [41]	?	Temporal Reasoning
Homogeneous C-relation	Adeleke, Macpherson, Neumann [1, 2]	Deuber, Miliken	?	Phylogeny Reconstruction
Countable atomless Boolean algebra	?	Graham, Leeb, Rothschild (see [32])	?	Set Constraints
Allen's Interval Algebra	?	This paper, Section 7	?	Temporal Reasoning

FIGURE 2. A diagram suggesting future research.

[10] Manuel Bodirsky and Jan Kára. The complexity of equality constraint languages. *Theory of Computing Systems*, 3(2):136–158, 2008. A conference version appeared in the proceedings of CSR'06.

[11] Manuel Bodirsky and Jan Kára. The complexity of temporal constraint satisfaction problems. *Journal of the ACM*, 57(2), 2009. An extended abstract appeared in the proceedings of STOC'08.

[12] Manuel Bodirsky and Jaroslav Nešetřil. Constraint satisfaction with countable homogeneous templates. *Journal of Logic and Computation*, 16(3):359–373, 2006.

[13] Manuel Bodirsky and Diana Piguet. Finite trees are ramsey with respect to topological embeddings. Preprint, arXiv:1002:1557, 2008.

[14] Manuel Bodirsky and Michael Pinsker. Minimal functions on the random graph. Preprint, arXiv:1003.4030, 2010.

[15] Manuel Bodirsky and Michael Pinsker. Schaefer's theorem for graphs. In *Proceedings of STOC'11*, 2011. Preprint of the long version available from http://www.dmg.tuwien.ac.at/pinsker/.

[16] Manuel Bodirsky, Michael Pinsker, and Todor Tsankov. Decidability of definability. In *Proceedings of LICS'11*, 2011.

[17] V. G. Bodnarčuk, L. A. Kalužnin, V. N. Kotov, and B. A. Romov. Galois theory for post algebras, part I and II. *Cybernetics*, 5:243–539, 1969.

[18] Mathias Broxvall and Peter Jonsson. Point algebras for temporal reasoning: Algorithms and complexity. *Artificial Intelligence*, 149(2):179–220, 2003.

[19] Andrei Bulatov, Andrei Krokhin, and Peter G. Jeavons. Classifying the complexity of constraints using finite algebras. *SIAM Journal on Computing*, 34:720–742, 2005.

[20] Peter J. Cameron. Transitivity of permutation groups on unordered sets. *Mathematische Zeitschrift*, 148:127–139, 1976.

[21] Peter J. Cameron. The random graph. *R. L. Graham and J. Nešetřil, Editors, The Mathematics of Paul Erdös*, 1996.

[22] Tomás Feder and Moshe Vardi. The computational structure of monotone monadic SNP and constraint satisfaction: A study through Datalog and group theory. *SIAM Journal on Computing*, 28:57–104, 1999.

[23] Zvi Galil and Nimrod Megiddo. Cyclic ordering is NP-complete. *Theoretical Computer Science*, 5(2):179–182, 1977.
[24] Michael Garey and David Johnson. *A guide to NP-completeness*. CSLI Press, Stanford, 1978.
[25] David Geiger. Closed systems of functions and predicates. *Pacific Journal of Mathematics*, 27:95–100, 1968.
[26] Martin Goldstern and Michael Pinsker. A survey of clones on infinite sets. *Algebra Universalis*, 59:365–403, 2008.
[27] Edmond E. Granirer. Extremely amenable semigroups 2. *Mathematica Scandinavica*, 20:93–113, 1967.
[28] Wilfrid Hodges. *Model theory*. Cambridge University Press, 1993.
[29] Pawel M. Idziak, Petar Markovic, Ralph McKenzie, Matthew Valeriote, and Ross Willard. Tractability and learnability arising from algebras with few subpowers. In *Proceedings of LICS'07*, pages 213–224, 2007.
[30] Peter Jeavons, David Cohen, and Marc Gyssens. Closure properties of constraints. *Journal of the ACM*, 44(4):527–548, 1997.
[31] Markus Junker and Martin Ziegler. The 116 reducts of $(\mathbb{Q}, <, a)$. *Journal of Symbolic Logic*, 74(3):861–884, 2008.
[32] Alexander Kechris, Vladimir Pestov, and Stevo Todorcevic. Fraissé limits, Ramsey theory, and topological dynamics of automorphism groups. *Geometric and Functional Analysis*, 15(1):106–189, 2005.
[33] Phokion G. Kolaitis and Moshe Y. Vardi. Conjunctive-query containment and constraint satisfaction. In *Proceedings of PODS'98*, pages 205–213, 1998.
[34] Richard E. Ladner. On the structure of polynomial time reducibility. *Journal of the ACM*, 22(1):155–171, 1975.
[35] Dugald Macpherson. A survey of homogeneous structures. *Discrete Mathematics*, 311(15):1599–1634, 2011.
[36] Rolf H. Möhring, Martin Skutella, and Frederik Stork. Scheduling with and/or precedence constraints. *SIAM Journal on Computing*, 33(2):393–415, 2004.
[37] Bernhard Nebel and Hans-Jürgen Bürckert. Reasoning about temporal relations: A maximal tractable subclass of Allen's interval algebra. *Journal of the ACM*, 42(1):43–66, 1995.
[38] Jaroslav Nešetřil. Ramsey theory. *Handbook of Combinatorics*, pages 1331–1403, 1995.
[39] Jaroslav Nešetřil. Ramsey classes and homogeneous structures. *Combinatorics, Probability & Computing*, 14(1-2):171–189, 2005.
[40] Jaroslav Nešetřil and Pavol Hell. Colouring, constraint satisfaction, and complexity. *Computer Science Review 2*, pages 143–163, 2008.
[41] Jaroslav Nešetřil and Vojtech Rödl. Combinatorial partitions of finite posets and lattices. *Algebra Universalis*, 19:106–119, 1984.
[42] Jaroslav Nešetřil and Vojtech Rödl. *Mathematics of Ramsey Theory*. Springer, Berlin, 1998.
[43] Jaroslav Opatrny. Total ordering problem. *SIAM Journal on Computing*, 8(1):111–114, 1979.
[44] Michael Pinsker. More sublattices of the lattice of local clones. *Order*, 27(3):353–364, 2010.
[45] Ivo G. Rosenberg. Minimal clones I: the five types. *Lectures in Universal Algebra (Proc. Conf. Szeged, 1983), Colloq. Math. Soc. J. Bolyai*, 43:405–427, 1986.
[46] Simon Thomas. Reducts of the random graph. *Journal of Symbolic Logic*, 56(1):176–181, 1991.
[47] Simon Thomas. Reducts of random hypergraphs. *Annals of Pure and Applied Logic*, 80(2):165–193, 1996.
[48] Peter van Beek. Reasoning about qualitative temporal information. *Artificial Intelligence*, 58:297–326, 1992.
[49] Ross Willard. Testing expressibility is hard. In *Proceedings of CP'10*, 2010.

Laboratoire d'Informatique (LIX), CNRS UMR 7161, École Polytechnique, 91128 Palaiseau, France

E-mail address: bodirsky@lix.polytechnique.fr
URL: http://www.lix.polytechnique.fr/~bodirsky/

Équipe de Logique Mathmatique, Université Denis Diderot - Paris 7, UFR de Mathématiques, 75205 Paris Cedex 13, France

E-mail address: marula@gmx.at
URL: http://dmg.tuwien.ac.at/pinsker/