Binary symmetries of tractable non-rigid structures

Paolo Marimon
Institut für Diskrete Mathematik & Geometrie
TU Wien
Vienna, Austria
0000-0001-9355-1685

Michael Pinsker
Institut für Diskrete Mathematik & Geometrie
TU Wien
Vienna, Austria
0000-0002-4727-918X

***Abstract*—We study constraint satisfaction problems of non-rigid structures in a finite and omega-categorical setting. We show that not having a binary essential polymorphism is a sufficient criterion for NP-hardness of the constraint satisfaction problem of a (model-complete) core, as long as its automorphism group is not the free action of a Boolean group. To understand the behaviour of low arity polymorphisms, we classify the possible types of minimal operations above an arbitrary permutation group. In this, we generalise a classical theorem of Rosenberg above the trivial group, and significantly improve a result of Bodirsky and Chen above the automorphism groups of omega-categorical structures. Finally, we answer three questions of Bodirsky on binary polymorphisms of infinite templates for constraint satisfaction problems.**

***Index Terms*—Constraint Satisfaction Problem, minimal operation, polymorphism, Boolean group, oligomorphic permutation group, primitive positive interpretation**

I. Introduction

A. Background: polymorphisms of non-rigid structures

For a relational structure $\mathbb{B}$, its Constraint Satisfaction Problem, $\mathrm{CSP}(\mathbb{B})$, is the computational problem of deciding, given some finite structure $\mathbb{A}$ in the same relational language, whether there is a homomorphism $h : \mathbb{A} \to \mathbb{B}$. A key insight of the algebraic approach to Constraint Satisfaction Problems (CSPs) is that differences in the computational and descriptive complexity of the CSPs of finite structures stem from structural differences captured by certain higher-arity symmetries known as *polymorphisms*. This is exemplified by the Bulatov and Zhuk's proofs [1], [2] of the Feder-Vardi Conjecture [3]: for a finite structure $\mathbb{B}$, either its polymorphisms only satisfy "trivial" (height-one) identities, in which case $\mathbb{B}$ *pp-constructs* all finite structures [4] and so $\mathrm{CSP}(\mathbb{B})$ is NP-complete, or $\mathbb{B}$ has a polymorphism satisfying a particular non-trivial identity (Siggers' identity [5]), in which case $\mathrm{CSP}(\mathbb{B})$ is in P. Many aspects of the algebraic approach also adapt to the infinite-domain setting for CSPs of certain countable structures which are characterised by very rich automorphism groups, known as *ω-categorical* structures (Definition II.15). For a subclass of these, i.e., *first-order reducts of finitely bounded homogeneous structures*, Bodirsky and Pinsker conjectured a complexity dichotomy similar to that on a finite domain (cf. [6]–[9]) and there is a substantial body of research proving specific cases [10]–[14].

When studying the CSP of a finite structure $\mathbb{B}$, we can always reduce to the case of $\mathbb{B}$ being a *core* [15], i.e., such that every endomorphism is an automorphism. In fact, we can always find a core $\mathbb{B}'$ with the same CSP as $\mathbb{B}$. When $\mathbb{B}$ is ω-categorical, we can similarly find a structure $\mathbb{B}'$ with the same CSP which is a *model-complete core* [16], i.e., such that the restriction of any endomorphism to any finite set can be extended to an automorphism. In the finite setting, another helpful reduction is that $\mathrm{CSP}(\mathbb{B}')$ is polynomial-time equivalent to $\mathrm{CSP}(\mathbb{B}'')$, where $\mathbb{B}''$ is obtained by naming all elements of $\mathbb{B}'$ [17], and so is *rigid*, i.e., with no non-trivial automorphisms.

For this reason, several tools of the algebraic approach in a finite setting work under the assumption that $\mathbb{B}$ is a rigid core. This presents some challenges in adapting these techniques to the infinitary setting, since ω-categorical structures are never rigid. Hence, work on infinite-domain CSPs studies polymorphisms of structures without assuming their rigidity, often yielding a better understanding of finite structures as a by-product [4], [18]. In this paper, we study polymorphisms of non-rigid (model-complete) cores in a finite and infinite setting. We will see that the non-rigidity assumption yields some surprisingly different behaviour from the rigid case.

We call the set of polymorphisms of a structure $\mathbb{B}$ its *polymorphism clone*, $\mathrm{Pol}(\mathbb{B})$. It is well-known that for $\mathbb{B}$ finite or ω-categorical, if all polymorphisms of $\mathbb{B}$ are non-constant and *essentially unary*, i.e., depend on at most one variable, then $\mathrm{CSP}(\mathbb{B})$ is NP-hard, as $\mathbb{B}$ *pp-interprets* all finite structures (Definition II.20) [19]–[21]. Hence, for the purposes of studying the complexity of CSPs on a given class of structures, we may assume without loss of generality that the structures we are working with have some polymorphisms which are *essential*, i.e. depend on more than one variable. Several CSP complexity dichotomies begin by studying the possible behaviours of low-arity essential polymorphisms in a given class of structures [10]–[12], [22], [23]. Often, this is part of the so-called "bottom-up" approach: first, show that for a structure $\mathbb{B}$ in the class being studied, if $\mathrm{CSP}(\mathbb{B})$ is not NP-hard (due to *pp*-constructing all finite structures),

Funded by the European Union (ERC, POCOCOP, 101071674). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This research was funded in whole or in part by the Austrian Science Fund (FWF) [I 5948].

then the polymorphisms of $\mathbb{B}$ must exhibit one of a given list of behaviours. Then, proceed by case-distinctions to show that structures whose polymorphisms exhibit such behaviours must have a tractable CSP. This is, for example, Jeavons' [19] approach in his algebraic proof of Schaefer's Theorem [24], i.e., the complexity dichotomy for CSPs on a two-element domain, and the strategy adopted by Bulatov [22] in the complexity dichotomy for CSPs on a three-element domain. Moreover, most complexity dichotomies in an ω-categorical setting rely on such a "bottom-up" approach [10]–[12], [23].

B. Motivation: three questions on binary polymorphisms

Our motivation for studying essential polymorphisms in non-rigid structures begins with three questions of Bodirsky [8] on binary polymorphisms in ω-categorical structures. A curious feature of the ω-categorical setting is that, at least in certain aspects, essential polymorphisms seem to exhibit better behaviour than what is known to be possible in arbitrary finite structures. In particular, a standard step in most infinite-domain dichotomies, including those that avoid a "bottom-up" approach [14], [25]–[27], relies on showing that if $\mathbb{B}$ has an essential polymorphism, then it also has a binary essential polymorphism [14], [28], [29]. Whilst this phenomenon is in general false for finite structures (even on a two-element domain), it seemed to be ubiquitous in an ω-categorical setting, leading Bodirsky to ask:

Question 1. (Question 14.2.6 (24) in [8]) Does every countable ω-categorical model-complete core with an essential polymorphism also have a binary essential polymorphism?

Previous authors developed techniques to find such binary essential polymorphisms [14], [28], [29]. The most general of these relies on $\mathrm{Aut}(\mathbb{B})$ satisfying a strong form of transitivity known as the *orbital extension property* (Definition IV.11). The only hitherto known transitive ω-categorical counterexamples to this property were *imprimitive* (i.e., with some non-trivial equivalence relation invariant under $\mathrm{Aut}(\mathbb{B})$). Thus, Bodirsky also asked:

Question 2. (Question 14.2.1 (2) in [8]) Does every ω-categorical structure with primitive automorphism group have the orbital extension property?

Finally, the existence of a binary essential polymorphism is often used to deduce (under mild assumptions) the existence of a binary injective one [10], [14], [28]–[30]. Binary injective polymorphisms are helpful to reduce $\mathrm{CSP}(\mathbb{B})$ to its injective variant [14], and more generally in most complexity classifications on an infinite domain [10], [14], [29]. Hence, it is natural to ask whether the existence of a binary injective polymorphism can be deduced for ω-categorical structures with sufficiently easy CSP:

Question 3. (Question 14.2.6 (27) in [8]) Does every ω-categorical structure without algebraicity that can be solved by Datalog also have a binary injective polymorphism?

C. Related work: minimal operations

We approach Bodirsky's questions by investigating the more general setting of polymorphisms of arbitrary non-rigid (model-complete) cores. When $\mathbb{B}$ is a finite core, if $\mathrm{Pol}(\mathbb{B})$ contains an essential polymorphism, it also contains an essential one which is *minimal* above $\mathrm{Aut}(\mathbb{B})$ [31] in the sense that it is of minimal arity such that every other polymorphism it generates together with $\mathrm{Aut}(\mathbb{B})$ and projections generates it back by composition with elements of $\mathrm{Aut}(\mathbb{B})$ and projections (Definition II.6). Such minimal operations can also be found whenever $\mathbb{B}$ is ω-categorical in a finite relational language [32], replacing the aforementioned notion of generation by that of *local generation* (Definition II.2). From the assumption of minimality, one can obtain a great deal of information about the behaviour of a polymorphism. In particular, a classical theorem of Rosenberg [33] classifies the possible behaviours (described in Definition II.8) of minimal operations above the trivial group acting on a finite set:

Theorem I.1 ([33]). *Let $\mathbb{B}$ be a finite rigid core, and suppose that $\mathrm{Pol}(\mathbb{B})$ contains an essential polymorphism. Then, $\mathrm{Pol}(\mathbb{B})$ contains a polymorphism which is minimal above the trivial group acting on $\mathbb{B}$ and of one of the following types:*

1) *a binary operation;*
2) *a ternary majority operation;*
3) *a ternary minority operation of the form $x + y + z$ in a Boolean group $(B, +)$;*
4) *a k-ary semiprojection for some $k \geq 3$.*

Using an easy generalisation of this result (Theorem II.13), Bodirsky and Chen [32] classify minimal operations above the automorphism groups of ω-categorical structures, yielding the following:

Theorem I.2 ([8], [32]). *Let $\mathbb{B}$ be an ω-categorical model-complete core in a finite relational language. Suppose that $\mathrm{Pol}(\mathbb{B})$ contains an essential polymorphism. Then, $\mathrm{Pol}(\mathbb{B})$ contains a polymorphism which is minimal above $\mathrm{Aut}(\mathbb{B})$ and of one of the following types:*

1) *a binary operation;*
2) *a ternary quasi-majority operation;*
3) *a k-ary quasi-semiprojection for some $3 \leq k \leq 2r - s$, where r is the number of $\mathrm{Aut}(\mathbb{B})$-orbitals (orbits under the componentwise action of $\mathrm{Aut}(\mathbb{B})$ on pairs) and s is the number of $\mathrm{Aut}(\mathbb{B})$-orbits.*

Again, these operations are described in Definition II.8. Theorem I.2 has one fewer type of operation than Theorem I.1, since it does not feature a minority-like operation. We will see (in Theorem B) that Theorem I.2 can be improved by excluding type (2) and further characterising type (3).

D. Results

We give counterexamples to all three questions of Bodirsky [8]. Regarding Question 1, we prove in Subsection IV-A that it has a positive answer for structures whose automorphism group have at most two orbits, but a negative answer in general (Theorem IV.3). In a sense, Question 1 asks whether we can find a binary essential polymorphism in every ω-categorical model-complete core whose CSP is not NP-hard for trivial reasons (namely, all polymorphisms being essentially unary). We prove in Subsection IV-B that we can find a binary essential polymorphism in most non-rigid finite cores and all ω-categorical model-complete cores whose CSP is not NP-hard due to *pp*-interpreting all finite structures:

Theorem A. *Let $\mathbb{B}$ be EITHER a finite core such that $\mathrm{Aut}(\mathbb{B})$ is not a Boolean group acting freely on B, OR an ω-categorical model-complete core. Suppose that $\mathrm{Pol}(\mathbb{B})$ contains no binary essential polymorphism. Then, $\mathbb{B}$ pp-interprets all finite structures.*

A *Boolean group* (a.k.a. *elementary Abelian* 2*-group*) is a group where every non-identity element has order 2, and we include the trivial group among them. A group action is *free* if the only group element fixing any point is the identity. Since *pp*-interpreting all finite structures implies NP-hardness of $\mathrm{CSP}(\mathbb{B})$ [20], [21], for the purposes of studying CSPs of ω-categorical structures we may assume without loss of generality that $\mathrm{Pol}(\mathbb{B})$ has a binary essential polymorphism.

Our strategy to prove Theorem A relies on a classification of the possible behaviours of polymorphisms satisfying a weak version of minimality (Theorem III.8). In this direction, we obtain general results (Theorem V.10), which for the settings relevant to CSPs boil down to the following:

Theorem B. *Let $\mathbb{B}$ be EITHER a finite non-rigid core OR an ω-categorical model-complete core in a finite relational language. Let s be the number of orbits of $\mathrm{Aut}(\mathbb{B})$. Suppose that $\mathrm{Pol}(\mathbb{B})$ contains an essential polymorphism. Then, it has a polymorphism minimal above $\mathrm{Aut}(\mathbb{B})$ which is one of:*

1) *a binary operation;*
2) *a ternary quasi-minority operation of the form $\alpha\mathfrak{q}$ for $\alpha \in \mathrm{Aut}(\mathbb{B})$, where*
   - *$\mathbb{B}$ is finite, $\mathrm{Aut}(\mathbb{B})$ is a Boolean group acting freely on $\mathbb{B}$, and $s = 2^n$ for some $n \in \mathbb{N}$;*
   - *the operation $\mathfrak{q}$ is a $\mathrm{Aut}(\mathbb{B})$-invariant Boolean Steiner 3-quasigroup;*
3) *a k-ary orbit-semiprojection for $3 \leq k \leq s$.*

The definitions of orbit-semiprojection and $\mathrm{Aut}(\mathbb{B})$-invariant Boolean Steiner 3-quasigroup are given in Definitions III.2 and V.7, and they correspond to strengthenings of being a quasi-semiprojection, and being a minority of the form $x + y + z$ in a Boolean group respectively.

A surprising feature of Theorem B is that (quasi-)majorities, item (2) in Theorems I.1 and I.2, cannot exist as minimal for non-rigid (model-complete) cores. Moreover, intriguingly, the case of $\mathrm{Aut}(\mathbb{B})$ being the free action of a Boolean group exhibits distinct behaviour from all other group actions in our study, with case (2) of Theorem B only occurring for them and when $\mathrm{Aut}(\mathbb{B})$ has 2^n orbits. Hence, for most non-rigid structures, there are only two possible types of minimal essential polymorphisms. Our results improve on Theorem I.2 from [32] in two ways: (i) we show that minimal quasi-majorities cannot occur, and (ii) for the only remaining type of operations of arity > 2 (i.e., the quasi-semiprojections), we specify the behaviour on the orbits and bound their arity by the number of orbits rather than the number of orbitals. Finally, no analogue of Theorem B was known on a finite domain.

Going back to Theorems IV.3 and A, not only do they give us general tools to find binary essential polymorphisms in ω-categorical structures when studying their CSPs, but they also go substantially beyond the scope of previous techniques, which relied on the orbital extension property. Firstly, the orbital extension property implies transitivity, and Theorem IV.3 tells us that whenever $\mathbb{B}$ is a model-complete core with transitive automorphism group, if $\mathrm{Pol}(\mathbb{B})$ contains an essential polymorphism, then it also contains a binary essential one. Moreover, in Subsection IV-C, we give counterexamples to Question 2, showing that previous techniques could not even deal with primitive ω-categorical structures:

Theorem C. *There are ω-categorical structures whose automorphism group is primitive and without the orbital extension property. Hence, Question 2 has a negative answer.*

Our counterexamples (Definition IV.16) are under the scope of the Bodirsky-Pinsker conjecture and come from Cherlin's recent classification of homogeneous 2-multitournaments [34].

Finally, in Section VI, we give a counterexample to Question 3. Hence, whilst for the purposes of CSPs of ω-categorical structures, we can in general assume that we are working with structures with binary essential polymorphisms, such polymorphisms may fail to be injective even in relatively simple contexts:

Theorem D. *There is an ω-categorical structure with no algebraicity whose CSP is solvable in Datalog, but which has no binary injective polymorphisms. Hence, Question 3 has a negative answer.*

Our counterexample (Definition VI.7) falls under the scope of the Bodirsky-Pinsker conjecture and is first-order definable in an infinite unary structure (i.e., a structure whose all relations are unary).

E. Overview of the paper

After the Preliminaries (Section II), Section III proves Theorem III.8, which is the main ingredient for Theorems A and B. Section IV is dedicated to its consequences. In Subsection IV-A, we prove Theorem IV.3, answering negatively Question 1. In Subsection IV-B, we prove Theorem IV.10, which implies Theorem A. In Subsection IV-C, we prove Proposition IV.17, which implies Theorem C. Section V sketches the proofs for Theorem V.10, which implies Theorem B. We conclude with Section VI, where we sketch the proof of Theorem VI.8, which implies Theorem D. Detailed proofs of all statements can be found in [35].

II. Preliminaries

We use uppercase letters to denote sets A, B, C, etc. We often work with relational structures and use blackboard bold capital letters such as $\mathbb{A}, \mathbb{B}, \mathbb{C}$, etc. to denote them, where $\mathbb{A}$ is a relational structure with domain A, and so on.

A. Closed function clones and minimal operations

We study minimal and almost minimal polymorphisms in arbitrary non-rigid (model-complete) cores through the algebraic analogue of polymorphism clones, i.e., closed function clones (Definition II.2). It is well-known that the latter correspond to the polymorphism clones of relational structures (Fact II.4). Below, we give some of the basic definitions and results relevant to the rest of the paper following the discussion of [31] and [32].

Let B denote a set. For $n \in \mathbb{N}$, $\mathcal{O}^{(n)}$ denotes the set B^{B^n} of functions $B^n \to B$, and $\mathcal{O}$ denotes $\bigcup_{n\in\mathbb{N}} \mathcal{O}^{(n)}$.

Definition II.1. Let B be a set. A **function clone** over B is a set $\mathcal{C} \subseteq \mathcal{O}$ such that

- $\mathcal{C}$ contains all projections: for each $1 \leq i \leq k \in \mathbb{N}$, $\mathcal{C}$ contains the k-ary projection to the ith coordinate $\pi_i^k \in \mathcal{O}^{(k)}$, given by $(x_1, \ldots, x_k) \mapsto x_i$;
- $\mathcal{C}$ is closed under composition: for all $f \in \mathcal{C} \cap \mathcal{O}^{(n)}$ and all $g_1, \ldots, g_n \in \mathcal{C} \cap \mathcal{O}^{(m)}$, $f(g_1, \ldots, g_n)$, given by
$$(x_1, \ldots, x_m) \mapsto f(g_1(x_1, \ldots, x_m), \ldots, g_n(x_1, \ldots, x_m)) \;,$$
is in $\mathcal{C} \cap \mathcal{O}^{(m)}$.

Definition II.2. Given $\mathcal{S} \subseteq \mathcal{O}$, $\langle \mathcal{S} \rangle$ denotes the smallest function clone containing $\mathcal{S}$. It consists of all functions in $\mathcal{O}$ which can be written as a term function using functions from $\mathcal{S}$ and projections. We equip $\mathcal{O}^{(n)}$ with the product topology and then $\mathcal{O}$ with the sum topology, where B was endowed with the discrete topology. We call a function clone $\mathcal{C}$ **closed**, when it is closed with respect to this topology. Meanwhile, for $\mathcal{S} \subseteq \mathcal{O}$, $\overline{\mathcal{S}}$ denotes the topological closure of $\mathcal{S}$ in $\mathcal{O}$. We have that $f \in \overline{\mathcal{S}}$ if for each finite $A \subseteq B$, there is some $g \in \mathcal{S}$ such that $g_{\restriction A} = f_{\restriction A}$. It is easy to see that for $\mathcal{S} \subseteq \mathcal{O}$, $\overline{\langle \mathcal{S} \rangle}$ is the smallest closed function clone containing $\mathcal{S}$. We say that $\mathcal{S}$ **locally generates** g if $g \in \overline{\langle \mathcal{S} \rangle}$. Frequently, we will just say that $\mathcal{S}$ generates g since we always work with local generation.

Note that over a finite set the notion of closure defined above trivialises, and so closed clones correspond to clones.

Definition II.3. Let $\mathbb{B}$ be a relational structure. We say that $f : B^n \to B$ is a **polymorphism** of $\mathbb{B}$ if it preserves all relations of $\mathbb{B}$: for any such relation R,
$$\text{if } \begin{pmatrix} a_1^1 \\ \vdots \\ a_1^k \end{pmatrix}, \ldots, \begin{pmatrix} a_n^1 \\ \vdots \\ a_n^k \end{pmatrix} \in R, \text{ then } \begin{pmatrix} f(a_1^1, \ldots, a_n^1) \\ \vdots \\ f(a_1^k, \ldots, a_n^k) \end{pmatrix} \in R \;.$$

We call the set of polymorphisms of $\mathbb{B}$ the **polymorphism clone of** $\mathbb{B}$, and denote it by $\mathrm{Pol}(\mathbb{B})$.

Fact II.4 ([36])**.** *The closed function clones on B correspond to the polymorphism clones of relational structures on B.*

Definition II.5. Let $\mathcal{D} \supsetneq \mathcal{C}$ be closed subclones of $\mathcal{O}$. We say that $\mathcal{D}$ is **minimal above** $\mathcal{C}$ if there is no closed clone $\mathcal{E}$ such that $\mathcal{C} \subsetneq \mathcal{E} \subsetneq \mathcal{D}$.

Definition II.6. Let $\mathcal{C}$ be a closed subclone of $\mathcal{O}$. We say that an operation $f \in \mathcal{O} \setminus \mathcal{C}$ is **minimal above** $\mathcal{C}$ if $\overline{\langle \mathcal{C} \cup \{f\} \rangle}$ is minimal above $\mathcal{C}$ and every operation in $\overline{\langle \mathcal{C} \cup \{f\} \rangle} \setminus \mathcal{C}$ has arity greater than or equal to that of f. For $\mathcal{S} \subseteq \mathcal{O}$, we say that $f \in \mathcal{O}$ is minimal above $\mathcal{S}$ if it is minimal above $\overline{\langle \mathcal{S} \rangle}$.

A closed function clone $\mathcal{D}$ is minimal above $\mathcal{C}$ if and only if there is some operation f on B which is minimal above $\mathcal{C}$ and such that $\overline{\langle \mathcal{C} \cup \{f\} \rangle} = \mathcal{D}$ [31]. Moreover, whenever B is finite and $\mathcal{C} \subsetneq \mathcal{D}$ are closed (necessarily, by finiteness) closed clones on B, there is a closed closed clone $\mathcal{E} \subseteq \mathcal{D}$ which is minimal above $\mathcal{C}$ [31]. This is in general not true on an infinite domain [37], but a version of this fact also holds in the context that interests us (cf. Fact II.19). Minimal clones and operations are heavily studied in universal algebra. We refer the reader to [37] and [38] for reviews of research on the topic beyond [33].

Notation II.7. We use the symbol '$\approx$' to denote identities which hold universally. For example, instead of
$$\forall x, y \in B \; f(x, y) = f(y, x) \;,$$
we write
$$f(x, y) \approx f(y, x) \;. \tag{1}$$

Definition II.8. We define some types of operations relevant to the rest of the paper by the identities that they satisfy:

- a **ternary quasi-majority** operation is a ternary operation m such that
$$m(x, x, y) \approx m(x, y, x) \approx m(y, x, x) \approx m(x, x, x) \;;$$
- a **quasi-Malcev** operation is a ternary operation M such that
$$M(x, y, y) \approx M(y, y, x) \approx M(x, x, x) \;;$$
- a **ternary quasi-minority** operation is a ternary operation $\mathfrak{m}$ such that
$$\mathfrak{m}(x, x, y) \approx \mathfrak{m}(x, y, x) \approx \mathfrak{m}(y, x, x) \approx \mathfrak{m}(y, y, y) \;;$$
- a **quasi-semiprojection** is a k-ary operation f such that there is an $i \in \{1, \ldots, k\}$ and a unary operation g such that whenever $(a_1, \ldots, a_k)$ is a non-injective tuple from B,
$$f(a_1, \ldots, a_k) = g(a_i) \;.$$

For each of the operations above, we remove the prefix "quasi" when the operation is **idempotent**, i.e., satisfies $f(x, \ldots, x) \approx x$. For example, a ternary majority is a ternary quasi-majority m which also satisfies $m(x, x, x) \approx x$; in the case of a semiprojection idempotency implies $g(x) \approx x$.

Definition II.9. We say that a k-ary operation f is **essentially unary** if there is a unary operation g and $1 \leq i \leq k$ such that
$$f(x_1, \ldots, x_k) \approx g(x_i) \;.$$

We say that f is **essential** if it is not essentially unary. A function clone $\mathcal{C}$ is essentially unary if all of its operations are essentially unary. It is essential if it has an essential operation.

B. Almost minimality

We obtain Theorem B by classifying minimal operations above the clone generated by any non-trivial permutation group (Theorem V.10). Towards this result, we first classify almost minimal operations above arbitrary permutation groups (Theorems III.8, V.3, and V.6), where we write $\boldsymbol{G \curvearrowright B}$ to denote a group acting faithfully on a set B (with no assumptions on its cardinality).

Definition II.10. Let $\mathcal{C}$ be a closed function clone. The k-ary operation $f \in \mathcal{O}\setminus\mathcal{C}$ is **almost minimal** above $\mathcal{C}$ if $\overline{\langle \mathcal{C} \cup \{f\}\rangle} \cap \mathcal{O}^{(r)} = \mathcal{C} \cap \mathcal{O}^{(r)}$ for each $r < k$.

Remark II.11. Unary functions in $\mathcal{O} \setminus \mathcal{C}$ are almost minimal.

Lemma II.12. *Let $\mathcal{C} \subsetneq \mathcal{D}$ be closed function clones. Then, $\mathcal{D}$ contains some almost minimal operation above $\mathcal{C}$.*

Proof: Let $f \in \mathcal{D} \setminus \mathcal{C}$ be of minimal arity. Any function that f locally generates together with $\mathcal{C}$ belongs to $\mathcal{D}$; hence if such a function is of strictly smaller arity than f, it also belongs to $\mathcal{C}$, by the minimality of the arity of f. Thus, f is almost minimal by definition. ∎

The starting point for our classification of almost minimal operations above permutation groups is the observation that the following version of Theorem I.1 of [33] in a non-idempotent context from [8, Theorem 6.1.42] only uses the assumption of almost minimality:

Theorem II.13 ([8], [32])**.** *Let $\mathcal{C}$ be an essentially unary clone without constant operations and let f be minimal above $\mathcal{C}$. Then, up to permuting its variables, f is of one of the following five types:*

1) *a unary operation;*
2) *a binary operation;*
3) *a ternary quasi-majority operation;*
4) *a quasi-Malcev operation;*
5) *a k-ary quasi-semiprojection for some $k \geq 3$.*

Remark II.14. Inspecting the proof of Theorem II.13 in [8, §6.1.8], one can see that the only property of minimal operations being used is almost minimality. In particular, Theorem II.13 actually gives a classification into five types of the almost minimal operations above an essentially unary clone (not containing any constant operations).

C. Finite and ω-categorical structures

We work with finite and ω-categorical relational structures. Below, we list some basic definitions and facts regarding these.

Definition II.15. For B countably infinite, we say that the permutation group $G \curvearrowright B$ is **oligomorphic** if, for each $n \in \mathbb{N}$, it has finitely many orbits in its componentwise action on n-tuples. A countable structure $\mathbb{B}$ is ***ω*-categorical** if its automorphism group $\mathrm{Aut}(\mathbb{B}) \curvearrowright B$ is oligomorphic.

Examples II.16. The ordered rational numbers $(\mathbb{Q}, <)$, the random graph, and countable vector spaces over finite fields are all ω-categorical. The first two are also finitely bounded homogeneous structures (see Definition IV.13).

We refer the reader to [8] and [9] for more on ω-categorical structures and their CSPs.

Definition II.17. We say that a finite or ω-categorical relational structure $\mathbb{B}$ is a **model-complete core** if $\overline{\mathrm{Aut}(\mathbb{B})}$ is its endomorphism monoid.

A finite model-complete core is such that $\mathrm{Aut}(\mathbb{B}) = \mathrm{End}(\mathbb{B})$ and is just known as a core. For $\mathbb{B}$ finite or ω-categorical, there is always a model-complete core $\mathbb{B}'$ such that $\mathrm{CSP}(\mathbb{B}) = \mathrm{CSP}(\mathbb{B}')$ [15], [16], which is also finite or ω-categorical (and finite when $\mathbb{B}$ is finite).

Remark II.18. By Definition II.15, if $\mathbb{B}$ is ω-categorical and $\mathbb{C}$ is such that $\mathrm{Pol}(\mathbb{B}) \subseteq \mathrm{Pol}(\mathbb{C})$, then $\mathbb{C}$ is also ω-categorical. In particular, by Fact II.4, any closed function clone whose unary operations are $\overline{\mathrm{Aut}(\mathbb{B})}$ is the polymorphism clone of an ω-categorical model-complete core with automorphism group $\mathrm{Aut}(\mathbb{B})$.

A good feature of ω-categorical structures in a finite relational language is that their polymorphism clones will always contain a minimal operation above their automorphism group as long as they have any essential polymorphism:

Fact II.19 ([32])**.** *Let $\mathbb{B}$ be an ω-categorical structure in a finite relational language. Let $\overline{\langle \mathrm{Aut}(\mathbb{B})\rangle} \subsetneq \mathrm{Pol}(\mathbb{C})$. Then, $\mathrm{Pol}(\mathbb{C})$ contains a minimal operation above $\overline{\langle \mathrm{Aut}(\mathbb{B})\rangle}$.*

Below, we define the notion of *pp*-interpretability, which is central to Theorem A. Our proofs rely on an algebraic characterisation of *pp*-interpretability in finite and ω-categorical structures [21], [39], which we give in Fact IV.5.

Definition II.20. A formula ϕ is **primitive positive** if it only contains existential quantifiers, conjunctions, and atomic formulas. For $m \in \mathbb{N}$, a set $S \subseteq B^m$ is ***pp*-definable** in $\mathbb{B}$ if there is a primitive positive formula ϕ such that S equals $\phi(\mathbb{B})$, the set of tuples from $\mathbb{B}$ satisfying ϕ. We say that a structure $\mathbb{B}$ ***pp*-interprets** another structure $\mathbb{A}$ if there is a partial surjective map $h : B^d \to A$ for some $d \geq 1$ such that for $R \subseteq A^n$ a relation of $\mathbb{A}$ (including equality and all of A), the set $h^{-1}(R) \subseteq B^{nd}$ is *pp*-definable in $\mathbb{B}$.

III. A three types theorem when $G \curvearrowright B$ is not the free action of a Boolean group

The main result of this section is Theorem III.8, which classifies almost minimal operations above $\overline{\langle G\rangle}$ for $G \curvearrowright B$ being a faithful group action which is not the free action of a Boolean group. We use as a starting point Theorem II.13, which, as pointed out in Remark II.14, gives a coarse classification of almost minimal operations above an essentially unary clone (with no constant operations). Then, we explore in more detail the consequences of almost minimality for each of the types of operations of arity > 2. Often, our arguments rely on showing that certain behaviours cannot be witnessed by an almost minimal operation of arity > 1 since they would imply that it generates some non-injective unary operation together with $\overline{\langle G\rangle}$. The latter contradicts almost minimality, as every operation in $\overline{G} = \overline{\langle G \rangle} \cap \mathcal{O}^{(1)}$ is injective (since it locally agrees with a permutation).

Remark III.1. We begin by pointing out that the actions of oligomorphic permutation groups are never free, and so the results in this section apply to them. To see this, take any element $a \in B$. The action of the stabilizer of a, G_a on B still has finitely many orbits by oligomorphicity. Hence, there are two elements $b, c \in B$ such that (a, b) and (a, c) lie in the same G-orbit. In particular, this means that there is a non-identity group element fixing a.

Definition III.2. Let $G \curvearrowright B$. We say that the k-ary operation f on B is an **orbit-semiprojection** with respect to the i-th variable for $i \in \{1, \ldots, k\}$ if there is a unary operation $g \in \overline{G}$ such that for any tuple $(a_1, \ldots, a_k)$ where at least two of the a_j lie in the same G-orbit,

$$f(a_1, \ldots, a_k) = g(a_i) \ .$$

Lemma III.3. *Let $G \curvearrowright B$. For $k \geq 3$, let f be a k-ary quasi-semiprojection which is almost minimal above $\overline{\langle G \rangle}$. Then, f is an orbit-semiprojection.*

Proof: We may say without loss of generality that there is a unary operation g such that whenever $|\{a_1, \ldots, a_k\}| < k$,

$$f(a_1, \ldots, a_k) = g(a_1) \ .$$

Let $\alpha \in G$. For all $i < j$, the function $h_{i,j,\alpha}$ defined by

$$(x_1, \ldots, x_k) \mapsto f(x_1, \ldots, x_{j-1}, \alpha x_i, x_{j+1}, \ldots, x_k)$$

does not depend on x_j, and so it depends on only one variable by the almost minimality of f. Moreover, setting for any $\ell \notin \{1, j\}$ the variable x_ℓ equal to x_1 and applying $h_{i,j,\alpha}$ yields $g(x_1)$, by the property on f assumed above. Thus, the variable $h_{i,j,\alpha}$ depends on is x_1, and

$$h_{i,j,\alpha}(x_1, \ldots, x_k) \approx g(x_1) \ .$$

Whenever $(a_1, \ldots, a_k)$ is any tuple where a_i, a_j belong to the same orbit, then $a_j = \alpha a_i$ for some $\alpha \in G$. By the above, f will return $g(a_1)$ on this tuple; hence, f is an orbit-semiprojection onto x_1 as witnessed by g. ∎

Lemma III.4. *Let $G \curvearrowright B$ be such that the action of G on B is not free. Then, no almost minimal function above $\overline{\langle G \rangle}$ can be a quasi-Malcev operation.*

Proof: Since the action of G on B is not free, there is some non-trivial $\alpha \in G$ and distinct $a, b, c \in B$ such that $\alpha(a) = a, \alpha(b) = c$. Suppose by contradiction that $M(x, y, z)$ is almost minimal and quasi-Malcev, and consider $h(x, y) = M(x, \alpha x, y)$. If $h(x, y)$ depends on the first argument,

$$M(a, a, a) = h(a, a) = h(a, b) = M(a, a, b) = M(b, b, b) \ ,$$

contradicting injectivity of $M(x, x, x) \in \overline{\langle G \rangle}$. Similarly, if $h(x, y)$ depends on the second argument,

$$\begin{aligned} M(c, c, c) = M(a, a, c) &= h(a, c) = h(b, c) \\ &= M(b, c, c) = M(b, b, b) \ , \end{aligned}$$

again contradicting injectivity of $M(x, x, x) \in \overline{\langle G \rangle}$. Thus, $h(x, y)$ depends on both arguments, contradicting the almost minimality of M. ∎

Lemma III.5. *Let $G \curvearrowright B$ with G not Boolean. Then, no almost minimal function above $\overline{\langle G \rangle}$ can be a quasi-Malcev operation.*

Proof: Let $\alpha \in G$ have order ≥ 3. Since $\alpha^2 \neq 1$, there is $b \in B$ such that $b, \alpha(b) = c$, and $\alpha^2(b) = d$ are all distinct. Suppose that $M(x, y, z)$ is almost minimal and quasi-Malcev. Consider $h(x, y) = M(x, y, \alpha x)$. We have that

$$h(c, d) = M(c, d, d) = M(c, c, c) = M(b, b, c) = h(b, b) \ .$$

By almost minimality, $h \in \overline{\langle G \rangle}$, and so it equals a unary injective function of either its first or its second argument. However, since b, c, d are distinct, this yields a contradiction. ∎

Lemma III.6. *Let $G \curvearrowright B$. Suppose that $m(x, y, z)$ is a quasi-majority operation which is almost minimal above $\overline{\langle G \rangle}$. Then, for any $\beta \in G \setminus \{1\}$, $h(x, y, z) := m(x, \beta y, z)$ is a quasi-Malcev operation such that*

$$h(x, x, x) \approx m(x, x, x) \ . \tag{2}$$

Proof: Note that h is also almost minimal above $\overline{\langle G \rangle}$ since $m(x, y, z) = h(x, \beta^{-1} y, z)$. Hence, the binary function given by $l(x, y) := h(x, x, y)$ must be in $\overline{\langle G \rangle}$ and so essentially unary. Observe that $l(x, y)$ cannot depend on the first variable. Otherwise,

$$\begin{aligned} m(x, x, x) \approx m(x, \beta x, x) \approx h(x, x, x) \approx l(x, x) \approx l(x, \beta x) \\ \approx h(x, x, \beta x) \approx m(x, \beta x, \beta x) \approx m(\beta x, \beta x, \beta x) \ , \end{aligned}$$

which contradicts injectivity of $m(x, x, x) \in \overline{\langle G \rangle}$. Hence, $l(x, y)$ depends on the second variable yielding

$$\begin{aligned} h(x, x, y) \approx l(x, y) \approx l(y, y) \approx h(y, y, y) \\ \approx m(y, \beta y, y) \approx m(y, y, y) \ . \end{aligned}$$

From this equation we obtain that h satisfies condition (2). By symmetry, the same argument applies for $h(y, x, x)$, yielding that

$$h(x, x, y) \approx h(y, y, y) \approx h(y, x, x) \ ,$$

and so that h is quasi-Malcev. ∎

Lemma III.7. *Let $G \curvearrowright B$. Suppose no almost minimal function above $\overline{\langle G \rangle}$ is quasi-Malcev. Then, no almost minimal function above $\overline{\langle G \rangle}$ can be a ternary quasi-majority operation.*

Proof: Suppose that m is a ternary almost minimal quasi-majority above $\overline{\langle G \rangle}$. Then, for $\beta \in G \setminus \{1\}$, $m(x, \beta y, z)$ is quasi-Malcev and almost minimal above $\overline{\langle G \rangle}$ by Lemma III.6. Since we are assuming there are no quasi-Malcev operations almost minimal above $\overline{\langle G \rangle}$, this yields a contradiction. ∎

Theorem III.8 (Three types theorem)**.** *Let $G \curvearrowright B$ be such that G is not a Boolean group acting freely on B. Let s be the (possibly infinite) number of orbits of G on B. Let f be* *an almost minimal operation above $\overline{\langle G \rangle}$. Then, f is of one of the following types:*

1) *a unary operation;*
2) *a binary operation;*
3) *a k-ary orbit-semiprojection for $3 \leq k \leq s$.*

Proof: We need to consider the five possibilities for f from Theorem II.13, which gives us a type-classification of the almost minimal operations above $\overline{\langle G \rangle}$ (Remark II.14). Lemma III.5 and Lemma III.4 tell us that f cannot be a quasi-Malcev. Moreover, f cannot be a quasi-majority from Lemma III.7. Finally, if f is a quasi-semiprojection it must be an orbit-semiprojection from Lemma III.3. Since f is an orbit-semiprojection, its arity must be $\leq s$, where s is the number of orbits of G. ■

It is easy to prove that the necessary conditions for the existence of an orbit-semiprojection as almost minimal in Theorem III.8 are also sufficient:

Proposition III.9. *Let $G \curvearrowright B$, and let s be the number of orbits of G on B. Then, for every $2 \leq k \leq s$, there is a k-ary orbit-semiprojection which is almost minimal over $\overline{\langle G \rangle}$.*

Proof sketch: It is easy to prove that an essential orbit-semiprojection is always almost minimal over $\overline{\langle G \rangle}$. So, we only need to build one. Take an enumeration $(O_r | r < s)$ of the G-orbits on B. Given $(a_1, \ldots, a_k) \in B^k$, set $f(a_1, \ldots, a_k) = a_1$ if two of the a_i are in the same orbit, and otherwise, let $f(a_1, \ldots, a_k) = a_j$, where a_j is the element appearing in the orbit with the largest index in our ordering. ■

IV. Consequences of the three types theorem for polymorphisms of non-rigid model-complete cores

In this section, we explore consequences of the three types theorem (Theorem III.8) for the polymorphism clones of non-rigid model-complete cores whose automorphism group is not the free action of a Boolean group on B. As pointed out in Remark III.1, these include the polymorphism clones of ω-categorical structures, for which we derive several results.

A. Essential polymorphisms in ω-categorical structures

In this subsection, we prove Theorem IV.3, which states that Question 1 has a positive answer whenever $\mathrm{Aut}(\mathbb{B})$ has ≤ 2 orbits and a negative answer in general.

Definition IV.1. For $G \curvearrowright B$, we say that a closed function clone $\mathcal{C}$ is a **core clone** with respect to G if $\mathcal{C} \cap \mathcal{O}^{(1)} = \overline{G}$. We simply say that $\mathcal{C}$ is a core clone if it is a core clone with respect to some $G \curvearrowright B$.

The following Corollary is a trivial consequence of Theorem III.8 and Proposition III.9.

Corollary IV.2. *Let $G \curvearrowright B$.*

- *Suppose that G has ≤ 2 orbits and it is not a Boolean group acting freely on B. Then any essential core clone with respect to G contains a binary essential operation;*
- *Suppose G has ≥ 3 orbits. Then, there is an essential core clone with respect to G with no binary essential operation.*

Proof: For the first statement, if $\mathcal{C}$ is an essential core clone with respect to G and G has ≤ 2 orbits, it must contain a binary essential operation almost minimal above $\overline{\langle G \rangle}$ by Lemma II.12 and Theorem III.8. For the second statement, if G has ≥ 3 orbits, by Proposition III.9, there is a k-ary orbit-semiprojection f almost minimal above $\overline{\langle G \rangle}$ for $k \geq 3$. Hence, $\overline{\langle G \cup \{f\} \rangle}$ is a core clone with no binary essential operation, by definition of almost minimality. ■

For $\mathbb{B}$ finite or ω-categorical, core clones with respect to $\mathrm{Aut}(\mathbb{B})$ correspond to polymorphism clones of model-complete cores with automorphism group $\mathrm{Aut}(\mathbb{B})$. Hence, from Corollary IV.2, and Remarks II.18 and III.1, we obtain:

Theorem IV.3. *Let $\mathbb{B}$ be an ω-categorical structure.*

- *Suppose $\mathrm{Aut}(\mathbb{B})$ has ≤ 2 orbits. If $\mathbb{B}$ is a model-complete core and it has an essential polymorphism, then it also has a binary essential one;*
- *Suppose $\mathrm{Aut}(\mathbb{B})$ has ≥ 3 orbits. Then, there is an ω-categorical model-complete core $\mathbb{C}$ with the same automorphism group as $\mathbb{B}$ and which has an essential polymorphism but no binary essential polymorphism. Hence, Question 1 has a negative answer.*

B. Finding binary symmetries

In this subsection, we prove Theorem IV.10, which implies Theorem A. It states that as long as $\mathcal{C}$ is a core clone with respect to a suitable permutation group $G \curvearrowright B$ without a uniformly continuous clone homomorphism to the clone $\mathcal{P}_{\{0,1\}}$ of projections on a two-element set (in the sense of Definition IV.4), it will have a binary essential polymorphism. For $\mathbb{B}$ finite or ω-categorical, $\mathrm{Pol}(\mathbb{B})$ having such a clone homomorphism is equivalent to $\mathbb{B}$ *pp*-interpreting all finite structures (cf. Fact IV.5 [21], [39]). Hence, Theorem IV.10 implies Theorem A. Moreover, Theorem IV.10 shows an interesting disanalogy with the case of idempotent clones (i.e., core clones above the trivial group), where on any domain there are clones with no essential binary operations but no clone homomorphism to $\mathcal{P}_{\{0,1\}}$ (e.g., any clone generated by a ternary majority).

Definition IV.4. Let $\mathcal{C}$ and $\mathcal{D}$ be clones with domains C and D respectively. A map $\eta : \mathcal{C} \to \mathcal{D}$ is a **clone homomorphism** if it preserves arities and universally quantified identities. In the special case where the domain D of $\mathcal{D}$ is finite, we say that η is **uniformly continuous** if there exists a finite $A \subseteq C$ such that $f_{\restriction A} = g_{\restriction A}$ implies $\eta(f) = \eta(g)$ for all $f, g \in \mathcal{C}$.

Fact IV.5. *([21], [39]) Let $\mathbb{B}$ be finite or ω-categorical and $\mathbb{A}$ finite. Then, $\mathbb{B}$ pp-interprets $\mathbb{A}$ if and only if there is a uniformly continuous clone homomorphism from $\mathrm{Pol}(\mathbb{B})$ to $\mathrm{Pol}(\mathbb{A})$. In particular, $\mathbb{B}$ pp-interprets all finite structures if and only if there is a uniformly continuous clone homomorphism from $\mathrm{Pol}(\mathbb{B})$ to the clone of projections on a two-element set $\mathcal{P}_{\{0,1\}}$.*

Definition IV.6. Let $G \curvearrowright B$. Then, by $\mathcal{S}$ we denote the closed clone generated by all orbit-semiprojections:

$$\mathcal{S} := \overline{\langle G \cup \{f \mid f \text{ is an orbit-semiprojection for } G \curvearrowright B\}\rangle} \,.$$

Lemma IV.7. *Let $G \curvearrowright B$. There is a uniformly continuous clone homomorphism from $\mathcal{S}$ to $\mathcal{P}_{\{0,1\}}$.*

Proof: Since $G \curvearrowright B$ is non-trivial, let $C \subseteq B$ be a non-trivial G-orbit (i.e., $|C| > 1$). Clearly, the map $\rho : \mathcal{S} \to \mathcal{S}_{\restriction C}$ sending each operation f to its restriction $f_{\restriction C}$ to C is a clone homomorphism since each identity satisfied by operations in $\mathcal{S}$ on B will also be satisfied on a restriction of the domain. Note next that any such restriction $f_{\restriction C}$ is essentially unary: this is clear for orbit-semiprojections and elements of G, and follows by an easy induction on terms for arbitrary operations in $\mathcal{S}$. Let $\tau : \mathcal{S}_{\restriction C} \to \mathcal{P}_{\{0,1\}}$ send each k-ary operation $f_{\restriction C}$ in $\mathcal{S}_{\restriction C}$ to the k-ary projection to the ith coordinate π_i^k, where i is the variable on which f depends. This is again easily seen to be a clone homomorphism. Thus, $\xi := \tau \circ \rho$ is a clone homomorphism $\xi : \mathcal{S} \to \mathcal{P}_{\{0,1\}}$, which moreover is uniformly continuous: let $B' = \{c, d\}$ for distinct $c, d \in C$. For any $f, g \in \mathcal{S}$, $f_{\restriction B'} = g_{\restriction B'}$ implies that f and g depend on the same variable, and whence $\xi(f) = \xi(g)$, yielding uniform continuity. ∎

Definition IV.8. Let $\mathcal{C}$ be a function clone. Let $\xi : \mathcal{C} \cap \mathcal{O}^{\leq(3)} \to \mathcal{P}_{\{0,1\}}$ be a map preserving arities. The **minor extension of** ξ is the map $\xi' : \mathcal{C} \to \mathcal{P}_{\{0,1\}}$ defined as follows:

Let $f \in \mathcal{C}$ be an n-ary operation. Let $a := (a_1, \ldots, a_n) \in \{0,1\}^n$. Write $f_a(x, y)$ for the binary operation induced by f substituting the variable x_i with x if $a_i = 0$ and with y otherwise. We then define $\xi'(f)(a) := \xi(f_a)(0, 1)$.

The requirement of ξ being defined on ternary operations rather than only binary ones stems from the following fact.

Fact IV.9 (Proposition 6.8 in [7])**.** *Let $\mathcal{C}$ be a function clone. Suppose that $\xi : \mathcal{C} \cap \mathcal{O}^{\leq(3)} \to \mathcal{P}_{\{0,1\}}$ is a partial clone homomorphism (i.e., it preserves arities and identities). Then, the minor extension $\xi' : \mathcal{C} \to \mathcal{P}_{\{0,1\}}$ is a clone homomorphism.*

From this, it follows that if a clone exhibits any structure at all in the sense that it has no clone homomorphism to $\mathcal{P}_{\{0,1\}}$, then it contains an essential operation of arity ≤ 3 [7]. In our context, we get the following:

Theorem IV.10. *Let $G \curvearrowright B$ be such that G is not a Boolean group acting freely on B. Suppose that $\mathcal{C}$ is a core clone with respect to G, and that $\mathcal{C}$ has no uniformly continuous clone homomorphism to $\mathcal{P}_{\{0,1\}}$, the clone of projections on a two-element set. Then, $\mathcal{C}$ contains a binary essential operation almost minimal above $\overline{\langle G \rangle}$.*

Proof: Note that if $f \in \mathcal{C} \cap \mathcal{O}^{(2)} \supsetneq \overline{\langle G \rangle} \cap \mathcal{O}^{(2)}$, since $\mathcal{C}$ is a core clone with respect to G, the operation f must be essential and almost minimal, and the desired conclusion follows. Hence, suppose that $\mathcal{C} \cap \mathcal{O}^{(2)} = \overline{\langle G \rangle} \cap \mathcal{O}^{(2)}$. Then, all ternary operations in $\mathcal{C} \setminus \overline{\langle G \rangle}$ are almost minimal. In particular, $\mathcal{C} \cap \mathcal{O}^{(3)}$ consists entirely of essentially unary operations and orbit-semiprojections.

From Lemma IV.7, there is a uniformly continuous clone homomorphism from $\mathcal{S}$ to $\mathcal{P}_{\{0,1\}}$ which, when restricted to $\mathcal{C} \cap \mathcal{O}^{\leq(3)}$ yields a map $\xi : \mathcal{C} \cap \mathcal{O}^{\leq(3)} \to \mathcal{P}_{\{0,1\}}$ preserving arities and identities. Hence, from Fact IV.9, the minor extension of ξ yields a clone homomorphism $\xi' : \mathcal{C} \to \mathcal{P}_{\{0,1\}}$. It is easy to see that this is uniformly continuous by observing that given B' a two-element set from a G-orbit, if f and g agree on B', then they are sent to the same projection. ∎

C. The orbital extension property and its failure

In this subsection, we answer negatively Question 2 by proving Theorem C, which follows from Proposition IV.17.

Definition IV.11. We say that $G \curvearrowright B$ has the **orbital extension property** if there is an orbital (i.e., an orbit of the componentwise action $G \curvearrowright B^2$) O such that for any $u, v \in B$ there is $z \in B$ such that $(u, z), (v, z) \in O$.

Note that having the orbital extension property implies transitivity of $G \curvearrowright B$. The automorphism groups of several transitive ω-categorical structures, such as the order of the rational numbers or the random graph, have the orbital extension property (see [14, Example 20] and [30, Lemma 3.7] for more examples). A transitive ω-categorical structure whose automorphism group does not have the orbital extension property is $K_{\omega,\omega}$, the complete bipartite graph where both parts of the partition are countable. As mentioned in the introduction, the orbital extension property was the main tool to prove existence of binary essential polymorphisms previously to Theorems IV.3 and A due to the following fact:

Fact IV.12 ([8], [14], [28], [29])**.** *Let $\mathcal{C}$ be a closed function clone with an essential operation and containing a permutation group $G \curvearrowright B$ with the orbital extension property. Then, $\mathcal{C}$ contains a binary essential operation.*

Below, we give examples of ω-categorical structures whose automorphism groups do not have the orbital extension property in spite of being primitive. Indeed, our examples are finitely bounded homogeneous structures in a binary language, and so fall under the scope of the Bodirsky-Pinsker conjecture [6], [7]. Below, we introduce some basic terminology regarding homogeneous structures (cf. [40]):

Definition IV.13. A countable relational structure is **homogeneous** if every isomorphism between its finite substructures extends to an automorphism of the whole structure. Homogeneous structures in a finite relational language are ω-categorical. For $\mathbb{B}$ a relational structure, its **age**, $\mathrm{Age}(\mathbb{B})$, is the class of finite substructures of B. The ages of homogeneous structures correspond to classes of finite structures known as Fraïssé classes which satisfy some combinatorial closure properties. In particular, given a Fraïssé class $\mathcal{F}$, there is a (unique up to isomorphism) countable homogeneous structure $\mathbb{B}$ whose age is $\mathcal{F}$, which we call its **Fraïssé limit**. For a relational language $\mathcal{L}$ and a set of finite $\mathcal{L}$-structures $\mathcal{D}$, $\mathrm{Forb}^{\mathrm{emb}}(\mathcal{D})$ denotes the class of finite $\mathcal{L}$-structures such that no structure in $\mathcal{D}$ embeds in them. For $\mathcal{L}$ finite, a homogeneous

TABLE I
FORBIDDEN SUBSTRUCTURES FOR $\widetilde{\mathbb{S}(3)}$ AND $\mathbb{S}(4)$

STRUCTURE	FORBIDDEN INSTANCES OF	
	$C(i,j,k)$	$L(i,j,k)$
$\widetilde{\mathbb{S}(3)}$	$(1,1,1),(2,2,2)$	$(1,1,1),(1,2,2),(2,1,2)$
$\mathbb{S}(4)$	$(1,1,1),(1,2,2)$	$(1,2,1),(2,1,1),(2,2,1),(2,2,2)$

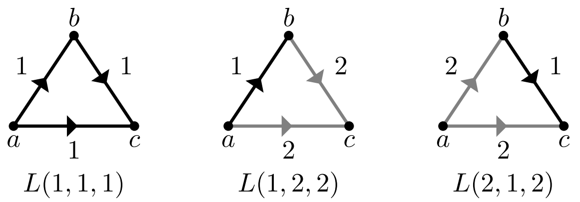

Fig. 1. Forbidden instances of $L(i,j,k)$ in $\widetilde{\mathbb{S}(3)}$.

$\mathcal{L}$-structure $\mathbb{B}$ is **finitely bounded** if there is some finite set of $\mathcal{L}$-structures $\mathcal{D}$ such that $\mathrm{Age}(\mathbb{B}) = \mathrm{Forb}^{\mathrm{emb}}(\mathcal{D})$.

The structures which will yield a negative answer to Question 2 are homogeneous 2-multitournaments:

Definition IV.14. A **2-multitournament** is a relational structure $\mathbb{C}$ in a language with two binary relations $\left\{\xrightarrow{1}, \xrightarrow{2}\right\}$ such that for any two elements $a, b \in C$, exactly one of

$$a \xrightarrow{1} b, b \xrightarrow{1} a, a \xrightarrow{2} b, \text{ or } b \xrightarrow{2} a$$

holds. We can think of a 2-multitournament as a tournament where every arc is coloured of one of two colours.

Cherlin [34, Table 18.1] recently classified the family of primitive 3-constrained homogeneous 2-multitournaments (i.e. primitive homogeneous 2-multitournaments whose age is of the form $\mathrm{Forb}^{\mathrm{emb}}(\mathcal{D})$, where every structure in $\mathcal{D}$ has size 3). We shall see that two such multitournaments do not have the orbital extension property.

Notation IV.15. For $i, j, k \in \{1, 2\}$ we let $C(i,j,k)$ denote the 2-multitournament of size 3 consisting of an oriented 3-cycle given by three vertices a, b, c such that

$$a \xrightarrow{i} b, b \xrightarrow{j} c, \text{ and } c \xrightarrow{k} a .$$

Similarly, $L(i,j,k)$ denotes the 2-multitournament of size 3 where the three vertices a, b, c are such that

$$a \xrightarrow{i} b, b \xrightarrow{j} c, \text{ and } a \xrightarrow{k} c .$$

Definition IV.16. The homogeneous 2-multitournaments $\widetilde{\mathbb{S}(3)}$ and $\mathbb{S}(4)$ are given by the Fraïssé limit of the classes of 2-multitournaments omitting embeddings of the 2-multitournaments on three vertices described by Table I.

Proposition IV.17. *The automorphism groups of the homogeneous* 2*-multitournaments* $\widetilde{\mathbb{S}(3)}$ *and* $\mathbb{S}(4)$ *do not have the orbital extension property.*

Proof: We run the proof for $\widetilde{\mathbb{S}(3)}$ since the proof for $\mathbb{S}(4)$ is virtually identical. There are four possible orbits of pairs for which the orbital extension property could hold. These orbitals of $\widetilde{\mathbb{S}(3)}$ are described by the formulas:

$$x \xrightarrow{1} y, y \xrightarrow{1} x, x \xrightarrow{2} y, \text{ and } y \xrightarrow{2} x .$$

Suppose the orbital extension property holds with respect to the orbital O. Let $u, v \in \widetilde{\mathbb{S}(3)}$ be such that $u \xrightarrow{1} v$. Then, O can be neither of the form $x \xrightarrow{1} y$ nor $y \xrightarrow{1} x$ since then there would be some $z \in \widetilde{\mathbb{S}(3)}$ such that u, v, z forms a copy of $L(1,1,1)$, which is forbidden. This can be seen by inspecting Figure 1. Similarly, O cannot be of the form $x \xrightarrow{2} y$ or $y \xrightarrow{2} x$ since $\widetilde{\mathbb{S}(3)}$ omits $L(1,2,2)$ and $L(2,1,2)$. Hence, there is no orbital for which the orbital extension property holds. ■

As pointed out earlier, $\widetilde{\mathbb{S}(3)}$ and $\mathbb{S}(4)$ are binary finitely bounded homogeneous structures with primitive automorphism group [34]. In particular, they are ω-categorical being homogeneous in a finite language. Thus, they yield a counterexample to Question 2, proving Theorem C.

V. MINIMAL OPERATIONS OVER PERMUTATION GROUPS

The techniques and ideas underlying Theorem III.8 can be pushed to classifying almost minimal and minimal operations above arbitrary permutation groups. In this section, we discuss these additional results and sketch the main proofs. Our main result, Theorem V.10, which implies Theorem B, classifies minimal operations above arbitrary permutation groups.

A. *Almost minimal operations over permutation groups*

We can also classify the almost minimal functions above the clone locally generated by a Boolean group acting freely on a set. It is easy to prove that freeness of $G \curvearrowright B$ yields that $\overline{\langle G \rangle} = \langle G \rangle$, and hence we shall write only $\langle G \rangle$.

For $|G| > 2$, a new type of almost minimal function is possible, which we call a G-quasi-minority:

Definition V.1. Let $G \curvearrowright B$. We say that a ternary operation $\mathfrak{m}$ is a ***G*-quasi-minority** if for all $\beta \in G$,

$$\begin{aligned} \mathfrak{m}(y, x, \beta x) &\approx \mathfrak{m}(x, \beta x, y) \\ &\approx \mathfrak{m}(x, y, \beta x) \approx \mathfrak{m}(\beta y, \beta y, \beta y) \ . \end{aligned} \tag{3}$$

We say that a G-quasi-minority is a ***G*-minority** if it is idempotent.

Meanwhile, we can show that as long as $|G| > 2$, there cannot be almost minimal quasi-majorities:

Lemma V.2. *Let* $G \curvearrowright B$ *where* $|G| > 2$*. Then, there is no quasi-majority almost minimal above* $\overline{\langle G \rangle}$*.*

Proof: Suppose by contradiction that m is a quasi-majority almost minimal above $\overline{\langle G \rangle}$. Since $|G| > 2$, take $\alpha, \beta \in G \setminus \{1\}$ distinct. Lemma III.6 implies $m(x, \gamma x, z) \approx m(z, z, z) \approx m(z, \gamma x, x)$ for any $\gamma \in G \setminus \{1\}$. Let u be a new variable. Setting $x = u, \gamma = \alpha, z = \beta u$ in the first identity yields $m(u, \alpha u, \beta u) \approx m(\beta u, \beta u, \beta u)$; on the other hand, setting $x = \beta u, z = u, \gamma = \alpha \circ \beta^{-1}$ in the second identity yields $m(u,\alpha u,\beta u)\approx m(u,u,u)$, contradicting injectivity of the map $x\mapsto m(x,x,x)$. ■

Hence, for the case of $G \curvearrowright B$ being the free action of a Boolean group of size > 2, we obtain the following classification of almost minimal operations:

Theorem V.3 (Boolean case)**.** *Let $G \curvearrowright B$ be a Boolean group acting freely on B with s-many orbits (where s is possibly infinite) and $|G| > 2$. Let f be an almost minimal operation above $\langle G\rangle$. Then, f is of one of the following types:*

1) *a unary operation;*
2) *a binary operation;*
3) *a ternary G-quasi-minority;*
4) *a k-ary orbit-semiprojection for $3 \le k \le s$.*

Proof sketch: Given Lemma III.3, Lemma V.2, and Theorem II.13, we only need to consider the behaviour of almost minimal quasi-Malcev operations. One can prove that these must be G-quasi-minorities when $|G| > 2$. ■

In the only remaining case of the group $\mathbb{Z}_2$ acting freely on B, there are two more types of almost minimal operations; these actually generate each other.

Definition V.4. Let $G \curvearrowright B$. Then, we say that $m(x,y,z)$ is an **odd majority** if m is a quasi-majority such that for all $\gamma \in G\setminus\{1\}$,

$$m(y,x,\gamma x)\approx m(x,\gamma x,y)\approx m(x,y,\gamma x)\approx m(y,y,y)\ . \quad (4)$$

Definition V.5. Let $G \curvearrowright B$. The ternary operation M is an **odd Malcev** if M is a quasi-Malcev such that for all $\gamma \in G\setminus\{1\}$,

$$M(x,y,x)\approx M(x,x,x)\ , \quad (5)$$
$$M(y,\gamma x,x)\approx M(x,\gamma x,y)\approx M(x,x,x)\ , \quad (6)$$
$$M(x,y,\gamma x)\approx M(\gamma y,\gamma y,\gamma y)\ . \quad (7)$$

Theorem V.6. *Let $\mathbb{Z}_2 \curvearrowright B$ freely with s-many orbits (where s is possibly infinite). Let f be an almost minimal operation above $\langle\mathbb{Z}_2\rangle$. Then, f is of one of the following types:*

1) *a unary operation;*
2) *a ternary G-quasi-minority;*
3) *an odd majority;*
4) *an odd Malcev (up to permutation of variables);*
5) *a k-ary orbit-semiprojection for $2 \le k \le s$.*

Proof sketch: Again, we use Theorem II.13 as a starting point. Firstly, using Lemma III.6, one can prove that an almost minimal quasi-majority has to be an odd majority. Then, one can prove that an almost minimal quasi-Malcev operation which is not a $\mathbb{Z}_2$-quasi-minority has to be an odd Malcev. Finally, it is easy to see that any binary almost minimal operation above $\langle\mathbb{Z}_2\rangle$ has to be an orbit-semiprojection. ■

Odd majorities and odd Malcev operations above $\mathbb{Z}_2$ are in a sense dual of each other, since an odd majority $m(x,y,z)$ always generates an odd Malcev $m(x,\gamma y,z)$ for $\gamma\in\mathbb{Z}_2\setminus\{0\}$ and vice versa. Moreover, almost minimal odd majorities and odd Malcev operations always exist over $\langle\mathbb{Z}_2\rangle$.

B. Back to minimality

We can actually prove that there are no minimal odd majorities or odd Malcevs. This allows us to treat uniformly all Boolean groups acting freely on a set when classifying minimal operations above $\langle G\rangle$ (as opposed to almost minimal operations). Moreover, we are able to prove that minimal G-quasi-minorities exhibit considerable structure:

Definition V.7. Let $G \curvearrowright B$. A ***G*-invariant Boolean Steiner 3-quasigroup** is a symmetric ternary minority operation $\mathfrak{q}$ also satisfying the following conditions:

$$\mathfrak{q}(x,y,\mathfrak{q}(x,y,z))\approx z\ ; \quad \text{(SQS)}$$
$$\mathfrak{q}(x,y,\mathfrak{q}(z,y,w))\approx \mathfrak{q}(x,z,w)\ ; \quad \text{(Bool)}$$
$$\text{for all } \alpha,\beta,\gamma\in G, \mathfrak{q}(\alpha x,\beta y,\gamma z)\approx \alpha\beta\gamma\mathfrak{q}(x,y,z)\ . \quad \text{(Inv)}$$

Idempotent symmetric minorities satisfying (SQS) are called Steiner 3-quasigroups or Steiner Skeins in the literature and have been studied in universal algebra and design theory [41]–[44]. When they further satisfy (Bool), they are called **Boolean Steiner 3-quasigroups**, since they are known to be of the form $\mathfrak{q}(x,y,z)=x+y+z$ in a Boolean group $(B,+)$. The condition (Inv) is novel and adds further constraints on the behaviour of $\mathfrak{q}$: it must induce another Boolean Steiner 3-quasigroup on the G-orbits which is related to $\mathfrak{q}$ in a particular way. Indeed, we can characterise G-invariant Boolean Steiner 3-quasigroups in much detail, yielding the following theorem:

Theorem V.8 (minimal operations, Boolean case)**.** *Let $G \curvearrowright B$ be a non-trivial Boolean group acting freely on B with s-many orbits (where s is possibly infinite). Let f be a minimal operation above $\langle G\rangle$. Then, f is of one of the following types:*

1) *a unary operation;*
2) *a binary operation;*
3) *a ternary minority of the form $\alpha\mathfrak{q}$ where $\mathfrak{q}$ is a G-invariant Boolean Steiner 3-quasigroup and $\alpha\in G$;*
4) *a k-ary orbit-semiprojection for $3 \le k \le s$.*

Furthermore, all G-invariant Boolean Steiner 3-quasigroups on B are minimal, and they exist if and only if $s = 2^n$ for some $n\in\mathbb{N}$ or is infinite.

Proof sketch: Starting from Theorem V.3 and V.6, the key step is realising that if $\mathfrak{m}$ is a G-quasi-minority, then

$$\mathfrak{m}'(x,y,z) := \mathfrak{m}(x,y,\mathfrak{m}(x,y,z))$$

has to be an orbit-semiprojection (not necessarily essential) with respect to its last coordinate. Since orbit-semiprojections can only generate other orbit-semiprojections in their same arity, if $\mathfrak{m}$ is minimal, $\mathfrak{m}'$ has to be essentially unary, and we can prove that $\mathfrak{m}'(x,y,z)\approx z$, yielding (SQS). With some additional work, we can also prove (Bool) and (Inv). Hence, (3) follows, since every G-quasi-minority is of the form $\alpha\mathfrak{q}$, where $\mathfrak{q}$ is a G-minority and $\alpha\in G$. Equation (Inv) implies that $\mathfrak{m}$ induces a Boolean Steiner 3-quasigroup on the G-orbits. Since these correspond to minorities of the form $x+y+z$ on a Boolean group, s has to be either 2^n for some $n\in\mathbb{N}$

or infinite. Moreover, for such choices of s, we can always construct a G-invariant Boolean Steiner 3-quasigroup.

Finally, we need to deal with the case of $\mathbb{Z}_2$. In this context, we can prove that if m is an idempotent odd majority, then for γ the non-identity element in $\mathbb{Z}_2$,

$$\mathfrak{m}^\star(x,y,z) := m(x, \gamma m(x,y,z), m(\gamma x, y, z))$$

is a G-minority. But if m was minimal, $\mathfrak{m}^\star$ would also be minimal, and so it would be a G-invariant Boolean Steiner 3-quasigroup. We can prove that the latter cannot generate a majority together with G, yielding that neither odd majorities nor odd Malcev operations (which generate odd majorities) can be minimal. ■

Generalising a classical result of Pálfy [45] on the existence of minimal semiprojections, we also get the following:

Theorem V.9 (Pálfy's Theorem for orbit-semiprojections)**.** *Let $G \curvearrowright B$ with s-many orbits with B finite or of the form $\mathrm{Aut}(\mathbb{B}) \curvearrowright B$ for $\mathbb{B}$ ω-categorical in a finite relational language. Then, for all $2 \leq k \leq s$, there is a k-ary orbit-semiprojection minimal above $\overline{\langle G \rangle}$.*

The various results mentioned in this section are then summarised in Theorem V.10:

Theorem V.10 (Minimal operations over permutation groups)**.** *Let $G \curvearrowright B$ be a non-trivial group acting faithfully on B with s many orbits (where s is possibly infinite). Let f be a minimal operation above $\overline{\langle G \rangle}$. Then, f is of one of the following types:*

1) *a unary operation;*
2) *a binary operation;*
3) *a ternary quasi-minority operation of the form $\alpha\mathfrak{q}$ for $\alpha \in G$, where*
   - *G is a Boolean group acting freely on B;*
   - *the operation $\mathfrak{q}$ is a G-invariant Boolean Steiner 3-quasigroup;*
4) *a k-ary orbit-semiprojection for $3 \leq k \leq s$.*

Moreover, in case (3), such operations exist if and only if $s = 2^n$ for some $n \in \mathbb{N}$ or is infinite. In case (4), if B is finite or $G = \mathrm{Aut}(\mathbb{B})$ for $\mathbb{B}$ ω-categorical in a finite relational language, we have that a minimal k-ary orbit semiprojection exists for each $2 \leq k \leq s$.

Given Fact II.19, Theorem B is simply a special case of Theorem V.10.

Example V.11. Theorem V.10 can help us classify the minimal operations above permutation groups in concrete examples. Consider the free action $\mathbb{Z}_2 \curvearrowright \{0,1\}$. From Theorem V.6, in this case the only possible minimal operations are either unary or quasi-minorities. The two constant operations on $\{0,1\}$ are clearly minimal above $\langle \mathbb{Z}_2 \rangle$. Meanwhile, both non-constant unary operations on $\{0,1\}$ are in $\langle \mathbb{Z}_2 \rangle$ and so cannot be minimal above it. Finally, note that $x+y+z$ (modulo 2) is the only $\mathbb{Z}_2$-invariant Boolean Steiner 3-quasigroup. Hence, this operation and its translate $(1-(x+y+z))$ are the only minimal quasi-minorities above $\langle \mathbb{Z}_2 \rangle$.

VI. Solvability in Datalog without binary injections

In this final section, we give a counterexample to Question 3 by sketching the proof of Theorem VI.8, which implies Theorem D. Datalog is a logic programming language which captures many standard polynomial-time algorithms used in the context of CSPs. Indeed, solvability in Datalog of a CSP corresponds to solvability by k-consistency algorithms for some k [46], [47]. We refer the reader to [47] for a discussion of Datalog in an ω-categorical setting and for the definition of a Datalog program. In this section, we use an algebraic characterisation of solvability in Datalog for a particular class of ω-categorical structures [48], thus avoiding the need to formally introduce the syntax and semantics of Datalog. Question 3 asks whether, for ω-categorical structures, having sufficiently easy CSP (i.e., CSP solvable in Datalog) implies the existence of a binary injective polymorphism under tame assumptions. Without the restriction to Datalog, some first-order reducts of $(\mathbb{Q},<)$ have CSP solvable in polynomial time (though not in Datalog) but have no binary injective polymorphisms (cf. Theorem 1.4 in [49] and Claim 4.8 in [50]).

Our counterexample (Definition VI.7) is first-order definable in an infinite unary structure. The CSPs of infinite unary structures (and structures definable in them) have been studied in detail with a complexity dichotomy proved in [51, Theorems 1.1 & 1.2] and a descriptive complexity analysis given in [48, Theorem 5.1]. The latter characterises solvability in Datalog for first-order reducts of unary structures (including structures first-order definable in unary structures) as being witnessed by a particular class of (canonical) polymorphisms known as pseudo-WNU operations (Definition VI.4). We will use this description for our proof of Theorem D.

The idea behind our counterexample is that any structure on the Boolean domain $\{0,1\}$ with a minimum operation among its polymorphisms has CSP solvable in Datalog (because it contains WNU operations in all arities [52], [53]). So, we "blow-up" the points of the original structure to the structure $\mathbb{D}$ consisting of disjoint unary predicates P_0 and P_1 with infinite domain. At this point we want to consider a function m acting as a minimum on these predicates. The challenge is balancing the requirement that m should generate a very rich clone (so that structures that have it as a polymorphism are solvable in Datalog), whilst also ensuring it does not generate any non-injective unary operation or any injective binary operation (which would break the requirements for a counterexample). This leads us to the set of operations $\mathfrak{M}$ defined below.

Definition VI.1. By $\mathbb{D}$ we denote the first-order structure with domain the disjoint union of two copies of $\mathbb{N}$ in the language with two unary predicates $\{P_0, P_1\}$, where each predicate names one of the two disjoint copies (which we also call P_0 and P_1).

Definition VI.2. For $i \in \{0,1\}$, let Inj_i be the set of injections $P_i \to P_i$ and BInj_i be the set of binary injections $P_i^2 \to P_i$. Define $\mathfrak{M}$ to be the set of binary operations $\mathrm{m} : D^2 \to D$

such that, for some $g_0 \in \mathrm{BInj}_0, g_1 \in \mathrm{BInj}_1$, and $\alpha, \beta \in \mathrm{Inj}_0$ such that α, β, and g_0 have disjoint images, we have,

$$\mathrm{m}(x,y) := \begin{cases} g_i(x,y) \text{ if } x,y \in P_i \text{ for } i \in \{0,1\}; \\ \alpha x \text{ if } x \in P_0, y \in P_1; \\ \beta y \text{ if } y \in P_0, x \in P_1. \end{cases}$$

So, m acts as a binary injection on each $\mathrm{Aut}(\mathbb{D})$-orbit and if its inputs are from different orbits, it selects the one in P_0 (and moves it by an appropriate injection of P_0).

As mentioned earlier, solvability in Datalog for structures first-order definable in unary structures is witnessed by canonical pseudo-WNU operations. Below we define these terms.

Definition VI.3. Let $G \curvearrowright B$. We say that an n-ary operation $f : B^n \to B$ is **canonical** with respect to G (or G-canonical) if for all $k \geq 1$, all k-tuples $\overline{a}_1, \ldots, \overline{a}_n \in B^k$, and all $\alpha_1, \ldots, \alpha_n \in G$, there is $\beta \in G$ such that

$$f(\alpha_1 \overline{a}_1, \ldots, \alpha_n \overline{a}_n) = \beta f(\overline{a}_1, \ldots, \overline{a}_n) \;, \tag{8}$$

where f and α are applied componentwise.

Note that the definition of a canonical function in (8) says that the G-orbit of the k-tuple $f(\overline{a}_1, \ldots, \overline{a}_n)$ is entirely determined by the G-orbits of the k-tuples $\overline{a}_1, \ldots, \overline{a}_n$. In particular, writing $\mathcal{T}_k$ for the space of G-orbits of k-tuples on B, the G-canonical function f induces an n-ary operation $\xi_k^G(f)$ on $\mathcal{T}_k$ defined as follows [6]: for $O_1, \ldots, O_n \in \mathcal{T}_k$, let $\xi_k^G(f)(O_1, \ldots, O_n)$ be given by the orbit of $f(\overline{a}_1, \ldots, \overline{a}_n)$ for some/all $\overline{a}_i \in O_i$.

Definition VI.4. An n-ary operation $w : B^n \to B$ is called a **weak near-unanimity operation (WNU)** if for any two tuples $\overline{z}_1, \overline{z}_2 \in \{x,y\}^n$ containing exactly one instance of y, we have

$$w(\overline{z}_1) \approx w(\overline{z}_2) \;.$$

So, for example, a ternary WNU operation w satisfies

$$w(x,x,y) \approx w(x,y,x) \approx w(y,x,x) \;.$$

Definition VI.5. Let $G \curvearrowright B$. We say that $w : B^n \to B$ is a **pseudo-WNU operation modulo $\overline{G}$** if for any $\overline{z}_1, \overline{z}_2 \in \{x,y\}^n$ containing exactly one instance of y there are $\alpha_{\overline{z}_1}, \beta_{\overline{z}_2} \in \overline{G}$ such that $\alpha_{\overline{z}_1} w(\overline{z}_1) \approx \beta_{\overline{z}_2} w(\overline{z}_2)$.

Lemma VI.6. *Let $\mathbb{C}$ be a model-complete core in a finite language with automorphism group $\mathrm{Aut}(\mathbb{D})$ and with $\mathrm{Pol}(\mathbb{C}) \cap \mathfrak{M} \neq \emptyset$. Then, $\mathrm{CSP}(\mathbb{C})$ is solvable in Datalog.*

Proof sketch: First, one needs to prove that any $\mathrm{m} \in \mathfrak{M}$ is $\mathrm{Aut}(\mathbb{D})$-canonical. Then, letting $g := \xi_2^{\mathrm{Aut}(\mathbb{D})}(\mathrm{m})$, it is easy to prove that $g(x, g(x,y)) \approx g(x,y)$ on $\mathcal{T}_2$. For $n \geq 3$, consider the operation on $\mathcal{T}_2^n$ given by

$$w_n(x_1, \ldots, x_n) := g(x_1, g(x_2, \ldots, g(x_{n-1}, x_n) \ldots)) \;.$$

The identity $g(x, g(x,y)) \approx g(x,y)$ implies that w_n is a weak near unanimity operation. Now, suppose that $\mathrm{m} \in \mathfrak{M} \cap \mathrm{Pol}(\mathbb{C})$, and let $\mathcal{C}$ be the clone of canonical functions in $\mathrm{Pol}(\mathbb{C})$. We just proved that $\xi_2^{\mathrm{Aut}(\mathbb{D})}(\mathcal{C})$ contains WNU operations of all arities ≥ 3. From [6, Proposition 6.6], this implies that $\mathrm{Pol}(\mathbb{C})$ contains pseudo-WNU canonical polymorphisms modulo $\overline{\mathrm{Aut}(\mathbb{D})}$ for all $n \geq 3$. From [48, Theorem 5.1], this is equivalent to $\mathrm{CSP}(\mathbb{C})$ being solvable in Datalog. ■

Definition VI.7. We define the following relations on $\mathbb{D}$:

$$r(x,y,z) := (P_1(x) \wedge P_1(y)) \to P_1(z);$$
$$p(x,y) := (x = y \wedge P_0(x)) \vee (P_1(x) \wedge P_1(y)) \;.$$

Consider the first-order structure $\mathbb{D}' := (D; P_0, P_1, r, p, \neq)$.

We say that an ω-categorical structure $\mathbb{B}$ has **algebraicity** if for some finite $A \subseteq B$, some orbit of the stabilizer of A, $\mathrm{Aut}(\mathbb{B})_A$, on $B \setminus A$ is finite.

Theorem VI.8. *The structure $\mathbb{D}'$ satisfies the following:*

1) *$\mathbb{D}'$ is a finitely bounded homogeneous structure with no algebraicity and $\mathrm{Aut}(\mathbb{D}') = \mathrm{Aut}(\mathbb{D})$;*
2) *$\mathbb{D}'$ is a model-complete core;*
3) *$\mathfrak{M} \subseteq \mathrm{Pol}(\mathbb{D}')$, and so $\mathrm{CSP}(\mathbb{D}')$ is solvable in Datalog;*
4) *$\mathbb{D}'$ has no binary injective polymorphisms.*

Proof sketch: Since $\mathbb{D}'$ is obtained by expanding $\mathbb{D}$ by definable relations, $\mathrm{Aut}(\mathbb{D}') = \mathrm{Aut}(\mathbb{D})$ and (1) holds. Meanwhile, (2) follows from the endomorphisms of $\mathbb{D}'$ preserving $\neq$ and the P_i (for $i \in \{0,1\}$). For (3), one needs to check that $\neq, p$, and r are preserved by all functions in $\mathfrak{M}$. As $\mathfrak{M} \subseteq \mathrm{Pol}(\mathbb{D}')$, solvability in Datalog follows from Lemma VI.6. Finally, (4) can be deduced from the fact that $\mathrm{Pol}(\mathbb{D})$ preserves P_0, P_0, r, and p. ■

Theorem VI.8 implies Theorem D since $\mathbb{D}'$ is ω-categorical being homogeneous in a finite relational language.

VII. Conclusion

Adapting terminology introduced by Bodirsky and Bodor [54], we call a finite or oligomorphic permutation group $G \curvearrowright B$ **stubborn** if every (model-complete) core with automorphism group G *pp*-interprets all finite structures, and so has an NP-hard CSP in virtue of its automorphisms alone. A stronger property than stubbornness is being **collapsing**: $\overline{\langle G \rangle}$ is the only function clone whose unary operations are exactly $\overline{G}$. We know that several finite permutation groups satisfying strong forms of primitivity, including $S_n \curvearrowright \{1, \ldots, n\}$ for $n \geq 3$, are collapsing (cf. [55]–[57]). However, stubbornness has not been investigated systematically so far. We therefore suggest the following question at the intersection of group theory and clone theory:

Question VII.1. Which (finite or oligomorphic) permutation groups are stubborn?

Due to Theorem A, an investigation of Question VII.1 should go through the study of the possible behaviours of binary minimal operations above different permutation groups (cf. [58]). In the meantime, we ask:

Question VII.2. Which permutation groups are such that the only binary operations minimal above them are orbit-semiprojections?

References

[1] A. A. Bulatov, "A dichotomy theorem for nonuniform CSPs," in *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2017, pp. 319–330.

[2] D. Zhuk, "A proof of the CSP dichotomy conjecture," *Journal of the ACM (JACM)*, vol. 67, no. 5, pp. 1–78, 2020.

[3] T. Feder and M. Y. Vardi, "Monotone monadic SNP and constraint satisfaction," in *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, 1993, pp. 612–622.

[4] L. Barto, J. Opršal, and M. Pinsker, "The wonderland of reflections," *Israel Journal of Mathematics*, vol. 223, no. 1, pp. 363–398, 2018.

[5] M. H. Siggers, "A strong mal'cev condition for locally finite varieties omitting the unary type," *Algebra universalis*, vol. 64, no. 1, pp. 15–20, 2010.

[6] M. Bodirsky, M. Pinsker, and A. Pongrácz, "Projective clone homomorphisms," *Journal of Symbolic Logic*, vol. 86, no. 1, pp. 148–161, 2021.

[7] L. Barto and M. Pinsker, "Topology is irrelevant (in a dichotomy conjecture for infinite domain constraint satisfaction problems)," *SIAM Journal on Computing*, vol. 49, no. 2, pp. 365–393, 2020. DOI: 10.1137/18M1216213. eprint: https://doi.org/10.1137/18M1216213. [Online]. Available: https://doi.org/10.1137/18M1216213.

[8] M. Bodirsky, *Complexity of infinite-domain constraint satisfaction*. Cambridge University Press, 2021, vol. 52.

[9] M. Pinsker, "Current challenges in infinite-domain constraint satisfaction: Dilemmas of the infinite sheep," in *2022 IEEE 52nd International Symposium on Multiple-Valued Logic (ISMVL)*, IEEE, 2022, pp. 80–87.

[10] M. Bodirsky and J. Kára, "The complexity of temporal constraint satisfaction problems," *Journal of the ACM (JACM)*, vol. 57, no. 2, pp. 1–41, 2010.

[11] M. Bodirsky and M. Pinsker, "Schaefer's theorem for graphs," *Journal of the ACM (JACM)*, vol. 62, no. 3, pp. 1–52, 2015.

[12] M. Kompatscher and T. Van Pham, "A complexity dichotomy for poset constraint satisfaction," *Journal of Applied Logics*, vol. 5, no. 8, p. 1663, 2018.

[13] M. Bodirsky, F. Madelaine, and A. Mottet, "A proof of the algebraic tractability conjecture for monotone monadic SNP," *SIAM journal on computing*, vol. 50, no. 4, pp. 1359–1409, 2021.

[14] A. Mottet and M. Pinsker, "Smooth approximations and CSPs over finitely bounded homogeneous structures," in *Proceedings of the 37th Annual ACM/IEEE Symposium on Logic in Computer Science*, 2022, pp. 1–13.

[15] P. Hell and J. Nešetřil, "The core of a graph," *Discrete Mathematics*, vol. 109, no. 1-3, pp. 117–126, 1992.

[16] M. Bodirsky, "Cores of countably categorical structures," *Logical Methods in Computer Science*, vol. 3, 2007.

[17] A. Bulatov, P. Jeavons, and A. Krokhin, "Classifying the complexity of constraints using finite algebras," *SIAM journal on computing*, vol. 34, no. 3, pp. 720–742, 2005.

[18] L. Barto, B. Bodor, M. Kozik, A. Mottet, and M. Pinsker, "Symmetries of structures that fail to interpret something finite," *arXiv preprint arXiv:2302.12112*, 2023.

[19] P. Jeavons, "On the algebraic structure of combinatorial problems," *Theoretical Computer Science*, vol. 200, no. 1-2, pp. 185–204, 1998.

[20] A. A. Bulatov, A. A. Krokhin, and P. G. Jeavons, "Classifying the complexity of constraints using finite algebras," *SIAM Journal on Computing*, vol. 34, pp. 720–742, 2005.

[21] M. Bodirsky and M. Pinsker, "Topological Birkhoff," *Transactions of the American Mathematical Society*, vol. 367, pp. 2527–2549, 2015.

[22] A. A. Bulatov, "A dichotomy theorem for constraint satisfaction problems on a 3-element set," *Journal of the ACM (JACM)*, vol. 53, no. 1, pp. 66–120, 2006.

[23] M. Bodirsky, B. Martin, M. Pinsker, and A. Pongrácz, "Constraint satisfaction problems for reducts of homogeneous graphs," *SIAM Journal on Computing*, vol. 48, no. 4, pp. 1224–1264, 2019.

[24] T. J. Schaefer, "The complexity of satisfiability problems," in *Proceedings of the tenth annual ACM symposium on Theory of computing*, 1978, pp. 216–226.

[25] A. Mottet, T. Nagy, and M. Pinsker, "An order out of nowhere: A new algorithm for infinite-domain CSPs," in *51th International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, 2024.

[26] R. Feller and M. Pinsker, "An algebraic proof of the graph orientation problem dichotomy for forbidden tournaments," *arXiv:2405.20263*, 2024.

[27] Z. Bitter and A. Mottet, "Generalized completion problems with forbidden tournaments," in *49th International Symposium on Mathematical Foundations of Computer Science, MFCS 2024*, Dagstuhl Publishing, vol. 306, 2024.

[28] M. Bodirsky and J. Kára, "The complexity of equality constraint languages," *Theory of Computing Systems*, vol. 43, pp. 136–158, 2008.

[29] M. Bodirsky and M. Pinsker, "Minimal functions on the random graph," *Israel Journal of Mathematics*, vol. 200, no. 1, pp. 251–296, 2014.

[30] M. Bodirsky and J. Greiner, "The complexity of combinations of qualitative constraint satisfaction problems," *Logical Methods in Computer Science*, vol. 16, 2020.

[31] Á. Szendrei, "Clones in universal algebra," *Les presses de L'universite de Montreal*, 1986.

[32] M. Bodirsky and H. Chen, "Oligomorphic clones," *Algebra Universalis*, vol. 57, pp. 109–125, 2007.

[33] I. G. Rosenberg, "Minimal clones I: The five types," in *Lectures in universal algebra*, Elsevier, 1986, pp. 405–427.
[34] G. Cherlin, *Homogeneous ordered graphs, metrically homogeneous graphs, and beyond*. Cambridge University Press, 2022, vol. 2, 3-Multi-Graphs and 2-Multi-Tournaments.
[35] P. Marimon and M. Pinsker, "Minimal operations over permutation groups," *arXiv:2410.22060*, 2024.
[36] I. Rosenberg and D. Schweigert, *Locally maximal clones*. Universität Kaiserslautern. Fachbereich Mathematik, 1999.
[37] B. Csákány, "Minimal clones—a minicourse," *Algebra Universalis*, vol. 54, no. 1, pp. 73–90, 2005.
[38] Á. Szendrei, "Ivo G. Rosenberg's work on maximal clones and minimal clones," *arXiv:2406.15184*, 2024.
[39] G. Birkhoff, "On the structure of abstract algebras," in *Mathematical proceedings of the Cambridge philosophical society*, Cambridge University Press, vol. 31, 1935, pp. 433–454.
[40] D. Macpherson, "A survey of homogeneous structures," *Discrete Mathematics*, vol. 311, no. 15, pp. 1599–1634, 2011, Infinite Graphs: Introductions, Connections, Surveys, ISSN: 0012-365X. DOI: https://doi.org/10.1016/j.disc.2011.01.024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0012365X11000422.
[41] R. W. Quackenbush, "Algebraic aspects of Steiner quadruple systems," in *Proceedings of the conference on algebraic aspects of combinatorics, University of Toronto*, 1975, pp. 265–268.
[42] M. H. Armanious, *Algebraische theorie der quadrupelsysteme*. Fachbereich der Mathematik der Technischen Hochschle Darmstadt. Federal Republic of Germany, 1980.
[43] R. W. Quackenbush, "Nilpotent block designs I: Basic concepts for Steiner triple and quadruple systems," *Journal of Combinatorial Designs*, vol. 7, no. 3, pp. 157–171, 1999.
[44] C. C. Lindner and A. Rosa, "Steiner quadruple systems-a survey," *Discrete Mathematics*, vol. 22, no. 2, pp. 147–181, 1978.
[45] P. P. Pálfy, "The arity of minimal clones," *Acta Scientiarum Mathematicarum*, vol. 50, pp. 331–333, 1986.
[46] T. Feder and M. Y. Vardi, "The computational structure of monotone monadic SNP and constraint satisfaction: A study through datalog and group theory," *SIAM Journal on Computing*, vol. 28, no. 1, pp. 57–104, 1998.
[47] M. Bodirsky and V. Dalmau, "Datalog and constraint satisfaction with infinite templates," *Journal of Computer and System Sciences*, vol. 79, no. 1, pp. 79–100, 2013.
[48] A. Mottet, T. Nagy, M. Pinsker, and M. Wrona, "Collapsing the bounded width hierarchy for infinite-domain constraint satisfaction problems: When symmetries are enough," *SIAM Journal on Computing*, vol. 53, no. 6, pp. 1709–1745, 2024, This paper also appeared at the 48th International Colloquium on Automata, Languages, and Programming (ICALP 2021) with title "Smooth Approximations and Relational Width Collapses". DOI: 10.1137/22M1538934. eprint: https://doi.org/10.1137/22M1538934. [Online]. Available: https://doi.org/10.1137/22M1538934.
[49] M. Bodirsky, J. Greiner, and J. Rydval, "Tractable combinations of temporal CSPs," *Logical Methods in Computer Science*, vol. 18, 2022.
[50] M. Bodirsky and J. Rydval, "On the descriptive complexity of temporal constraint satisfaction problems," *Journal of the ACM*, vol. 70, no. 1, pp. 1–58, 2022.
[51] M. Bodirsky and A. Mottet, "A dichotomy for first-order reducts of unary structures," *Logical Methods in Computer Science*, vol. 14, 2018.
[52] M. Maróti and R. McKenzie, "Existence theorems for weakly symmetric operations," *Algebra universalis*, vol. 59, no. 3-4, pp. 463–489, 2008.
[53] L. Barto and M. Kozik, "Constraint satisfaction problems solvable by local consistency methods," *Journal of the ACM (JACM)*, vol. 61, no. 1, pp. 1–19, 2014.
[54] M. Bodirsky and B. Bodor, "Structures preserved by primitive actions of S_ω," *arXiv:2501.03789*, 2025.
[55] P. P. Pálfy and Á. Szendrei, "Unary polynomials in algebras, II," in *Contributions to general algebra*, Proceedings of the Conference in Klagenfurt, vol. 2, Verlag Hölder-Pichler-Tempsky, Wien, 1982, pp. 273–290.
[56] K. A. Kearnes and Á. Szendrei, "Collapsing permutation groups," *Algebra Universalis*, vol. 45, no. 1, pp. 35–51, 2001.
[57] L. Haddad and I. G. Rosenberg, "Finite clones containing all permutations," *Canadian Journal of Mathematics*, vol. 46, no. 5, pp. 951–970, 1994. DOI: 10.4153/CJM-1994-054-1.
[58] Z. Brady, "Coarse classification of binary minimal clones," *arXiv:2301.12631*, 2023.